Biased Technological Change and Employment Reallocation[*]

Zsófia L. Bárány[†]   Christian Siegel[‡]

February 14, 2018

Abstract

To study the drivers of the employment reallocation across sectors and occupations between 1960 and 2010 in the US we propose a model where technology evolves at the sector-occupation cell level. This framework allows us to quantify the bias of technology across sectors and across occupations. We implement a novel method to extract changes in sector-occupation cell productivities from the data. Using a factor model we find that occupation and sector factors jointly explain 74-87 percent of cell productivity changes, with the occupation component being by far the most important. While in our general equilibrium model both factors imply similar reallocations of labor across sectors and occupations, quantitatively the bias in technological change across occupations is much more important than the bias across sectors.

Keywords: biased technological change, structural change, employment polarization

JEL codes: O41, O33, J24

1 Introduction

Over recent decades the labor markets in most developed countries have experienced substantial changes. There has been structural change, the reallocation of labor across sectors, while at the occupational level labor markets experienced polarization; employment shifted out of middle-earning routine jobs to low-earning manual and high-earning abstract jobs. Each of these patterns separately is typically explained through non-neutral productivity growth. Technological change biased across sectors is the leading explanation for structural change, whereas technological change that is biased across occupations is a prominent explanation for job polarization.[1] The goal of this paper is twofold: to identify the nature of technological change, and to assess what type of technological change is quantitatively relevant for the two phenomena.

We develop a new approach in assessing the nature of technological change that is driving these reallocations. We specify a flexible model in which technological change can be biased towards workers in specific sector-occupation cells. Drawing on key equations of the production side of this model together with data from the US Census and from the U.S. Bureau of Economic Analysis between 1960 and 2010 we extract the evolution of productivity for each sector-occupation cell. We use a factor model to quantify the bias in technology by decomposing the change in the cell level productivities into a neutral, a sector, and an occupation component. The neutral component captures general purpose technologies affecting all workers, the sector component captures productivity changes that are common to all workers within a sector (linked to the output an industry or firm produces), while the occupation component captures changes that are common to workers within an occupation (linked to the task content of an occupation). Having extracted these components, first we evaluate their importance in the observed productivity changes at the sector-occupation cell level, and second using our model we evaluate their role in the reallocation of employment across sectors and occupations, as well as in the evolution of occupational wages and of sectoral prices.

---

[1]See for example Kongsamut, Rebelo, and Xie (2001), Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008) for structural change, Autor, Katz, and Kearney (2006), Goos and Manning (2007), Autor and Dorn (2013), Goos, Manning, and Salomons (2014), Michaels, Natraj, and Van Reenen (2014) for polarization.

One strength of our fully flexible approach is that it does not impose any particular type of technological progress. This is especially important due to some trends in occupation and sector employment that have not received much attention in the literature. First, the goods sector has the highest share of routine workers; by far most of the decline in routine employment occurred in the goods sector, and conversely almost all of the contraction in goods sector 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. From these patterns of sector-occupation employment it is clear that the sectoral and the occupational reallocation of employment are closely linked, suggesting a connection between structural change and polarization.[2] It is precisely this link that makes it difficult to identify the true nature of technological change.

The approach in this paper 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. In Bárány and Siegel (2018) we show that forces behind structural change, i.e. differences in productivity growth across sectors, lead to polarization of wages and employment at the sectoral level, which in turn imply polarization in occupational outcomes. In that paper we took the unbalanced sectoral labor productivity growth from the data and analyzed its implications, whereas in this paper we aim to quantify to what extent technological progress is truly specific to sectors and to occupations. Conversely, 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

---

[2]Appendix A.1 lists the classification of occupations and industries and we show these employment patterns in Figure 8.

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.

The close link in the data between the sectoral and occupational reallocation of labor explains why models, such as those mentioned above, 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. Beyond the theoretical, academic interest in understanding what drives changes in occupational composition within and across sectors, it also has profound policy implications. Recently much of the political debate has focused on active labor market policies (such as training programs), and on protectionist policies aiming at maintaining certain industries of the home economy. While in our model there are no frictions or externalities which would justify these policies, we view our model as an important and useful first step in evaluating such policies. Our general equilibrium model focuses on the technology side of the economy, but could be extended to include a frictional labor market or job specific human capital accumulation.

In the analysis we rely on our flexible model to infer in each period cell productivities. Note that without a model it is impossible to quantify technological change that is biased towards a particular factor in production. While observing factor inputs and output allows the computation of a neutral productivity,[3] it is not possible to infer factor-specific productivities without making assumptions about the structure of production. We assume a CES production function over the different types of occupational labor in each sector. This is similar to Katz and Murphy (1992) who assume that skilled and unskilled labor produce a final good according to a CES production function and estimate the elasticity of substitution between these education groups and a constant trend in skill-biased technological change. Krusell, Ohanian, Ríos-Rull, and Violante (2000) specify a nested CES production function between capital equipment,

---

[3]This is how total factor productivity (for instance at the sectoral level) is extracted; note that changes in measured TFP might actually be driven by technological change augmenting only an individual factor of production. For example Greenwood, Hercowitz, and Krusell (1997) take the observed equipment price index as a measure of investment-specific (factor-biased) technological change and use a Cobb-Douglas production function to back out neutral productivity growth.

unskilled and skilled labor, and estimate the various substitution elasticities using the equipment price index as a measure of capital-equipment biased technological change. While our model does not feature any form of capital as inputs to production and we focus on labor productivities throughout, capital-driven technological change, biased or neutral, is captured in our cell productivities. Both these papers impose a specific form of factor-biased technical change, and it is this what allows them to infer the elasticity of substitution. Our approach is very different, while we do not estimate the elasticities, we also do not restrict the form of technological change, but we extract factor productivities from the data for a wide range of elasticities.

In the model productivity growth is specific to the sector-occupation cell. It is true that in such a setup one in general has to identify more parameters than in a setup which has sector-specific and/or occupation-specific technological change (but not their interaction). However, given this flexible setup we pin down these cell productivities exactly from the data using the production side of the model, rather than approximately matching fewer targets.[4] Moreover, we believe that technologies are linked to the job that workers do, where a job is defined by both the sector and the occupation of the worker (i.e. the factor of production is the sector-occupation labor input). Note, assuming that productivity evolves at the sector-occupation cell level does not rule out common trends in technological change, for instance at the sector or the occupation level.[5] To put it differently, the framework is flexible enough to allow for technological change to be specific to certain occupations, which is the key driver emphasized in the polarization literature, or to certain sectors, the mechanism often

---

[4]The targets we use for O occupations and S sectors are: $S \cdot (O-1)$ occupational labor income shares within sectors, $S - 1$ relative sectoral prices, and the overall growth rate of the economy.

[5]It is easy to conceive that some technologies improve a given occupation's productivity in a similar way regardless of the sector of work. For example an accountant's productivity has increased by the advent of computers, though potentially more so in sectors characterized by larger firms. There are also occupations which – though similar – perform different tasks depending on the sector of work. For instance, think of a cleaner working in a law firm versus in a production plant. The productivity of a cleaner working in a production plant presumably increased quite significantly in the last decades due to the introduction of specialized cleaning equipment for the production plant, whereas the productivity of the cleaner in a law firm probably has stayed constant since the introduction of vacuum cleaners. Regarding sector-specific productivities, Ford's Model T is a good example: by introducing the moving assembly line in production, rather than the then usual hand crafting, the productivity of workers directly producing the car increased, leading to a spillover on other workers in Ford, and later to workers in other car producers. In this sense the introduction of assembly lines in car manufacturing can be regarded as a sector-specific productivity change.

stressed in the structural change literature.

We take these sector-occupation cell productivity changes and use a factor model to decompose them into a neutral, a sector, an occupation and a residual component, where the last component captures productivity growth that is specific to the sector-occupation cell. From these components we construct counterfactual cell productivity series to quantify their importance in the observed cell productivity growth. We conduct this decomposition for a wide range of the elasticity of substitution between different occupations. We find that for values typically considered in the literature (an elasticity between 0.5 and 0.9), around a quarter of cell productivity growth is specific to the sector-occupation cell, between 66 and 80 percent can be attributed to occupation-specific productivity growth, 3-9 percent to sector-specific components, and 0-2 percent of productivity growth is neutral (i.e. common to all cells). We interpret this as evidence that most of productivity changes are not neutral, but biased across occupations and to a lesser extent across sectors, and that a significant part of technology is specific to the sector-occupation cell. Factor models have been used for instance in Stockman (1988), Ghosh and Wolf (1997) and Koren and Tenreyro (2007). While all these papers run a factor model at the country-sector level, they use their estimates to decompose the volatility of a series at a higher level of aggregation. We, however, not only study a very different question, but build counterfactual cell productivity series based on our factor model estimates.

Finally, to run counterfactual experiments we close the model by assuming that a representative household chooses sectoral consumption in order to maximize a non-homothetic CES utility and that individuals optimally choose their occupation subject to idiosyncratic entry costs. We first confirm that the model with the baseline cell productivities matches the data well. We then feed the counterfactual productivity paths into the model to determine how important each component is in explaining various outcomes of interest. We conclude that while qualitatively both the sector and the occupation productivity components generate employment and wage paths in line with the data, quantitatively the occupation component gets much closer. To explain the evolution of sectoral prices, both sector and occupation components are needed. For occupational income shares within sectors and employment shares at the cell level,

6

the sector-only component has almost no effect, whereas the productivity component specific to the sector-occupation cell has a significant role.

Overall our results suggest that to study the joint evolution of employment, wages and prices one needs to consider technological change at the sector-occupation cell level. However, to study labor market outcomes at the sector or at the occupation level, it is sufficient to model occupation-specific technological change. An implication of these findings is that policies targeting workers' occupational choice might be better at improving labor market outcomes than industrial policies.

The paper proceeds as follows: section 2 introduces the model and section 3 presents the data and the model parameterization. Section 4 shows the decomposition of productivity growth into components, and section 5 analyzes what the identified components imply for economic outcomes. The final section concludes.

2   Model

We assume that there is a continuum of measure one of heterogeneous workers in the economy. Workers optimally select their occupation and can freely choose which sector of the economy to supply their labor in. This implies that in equilibrium there is a single wage rate in each occupation which is common across sectors. We further assume that the different types of labor are imperfect substitutes in the production process in each sector, and that each sector values these types of workers differently in production.

The three types of workers are organized into a stand-in household, which derives utility from consuming all types of goods and services, and maximizes its utility subject to its budget constraint. The economy is in a decentralized equilibrium at all times: firms operate under perfect competition, prices and wages are such that all markets clear. We use this parsimonious static model to pin down how the productivity of the different occupations in each sector changes over time.

2.1 Sectors and production

There are three sectors in the economy which respectively produce low-skilled services (L), goods (G), and high-skilled services (H). All goods and services are produced in perfect competition. Each sector uses only labor as input in its production, but each combines all three types of occupations (manual, routine and abstract), with the following CES production function:

$$Y_J = \left[(\alpha_{mJ} l_{mJ})^{\frac{\eta-1}{\eta}} + (\alpha_{rJ} l_{rJ})^{\frac{\eta-1}{\eta}} + (\alpha_{aJ} l_{aJ})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \qquad \text{for } J \in \{L, G, H\}, \qquad (1)$$

where l_{oJ} is occupation o labor used in sector J, $\alpha_{oJ} > 0$ is an *sector-occupation* specific labor augmenting technology term for occupation $o \in \{m, r, a\}$ in sector J,[6] and $\eta \in [0, \infty]$ is the elasticity of substitution between the different types of labor.[7] In the initial year α_{oJ} 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. This formulation of the production function is very flexible and 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.[8] We use the model to calculate from the data the sector-occupation specific productivity terms, which we then decompose into common factors, as described in section 4.

Each firm takes prices and wages as given, and firms' first order conditions pin

---

[6]An alternative, isomorphic way of writing the production function in (1) is $Y_J = \left[x_{mJ}(A_{mJ} l_{mJ})^{\frac{\eta-1}{\eta}} + x_{rJ}(A_{rJ} l_{rJ})^{\frac{\eta-1}{\eta}} + x_{aJ}(A_{aJ} l_{aJ})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$, where x_{oJ} are constant weights and A_{oJ} are cell productivities that can change over time. The two formulations are equivalent since one can rewrite $\alpha_{oJ} = x_{oJ}^{\frac{\eta}{\eta-1}} A_{oJ}$. In this sense the α_{oJ} terms comprise of a fixed weight and a changing sector specific occupational labor augmenting technology. In our quantitative work we are interested in changes in productivity over time, which – due to the weights being constant – are equal in the two formulations, $\Delta \log \alpha_{oJ} = \Delta \log A_{oJ}$.

[7]We assume the same elasticity of substitution in all sectors since we do not want to confound changes in productivity that are specific to sectors with potential differences in elasticities.

[8]Given the close link between the sectoral and the occupational reallocation of employment, which we discussed in the introduction, had we set up the production function allowing only for sector-specific or only for occupation-specific terms we would potentially have attributed changes to this one factor which are actually due to the other factor. Our approach circumvents this problem as we do not impose any a priori restrictions.

down the optimal relative labor use as:

$$\frac{l_{mJ}}{l_{rJ}} = \left(\frac{w_r}{w_m}\right)^{\eta} \left(\frac{\alpha_{mJ}}{\alpha_{rJ}}\right)^{\eta-1}, \tag{2}$$

$$\frac{l_{aJ}}{l_{rJ}} = \left(\frac{w_r}{w_a}\right)^{\eta} \left(\frac{\alpha_{aJ}}{\alpha_{rJ}}\right)^{\eta-1}. \tag{3}$$

It is optimal to use more manual labor relative to routine labor in all sectors if the relative routine wage, w_r/w_m, is higher. Additionally, if in sector J the term $\left(\frac{\alpha_{mJ}}{\alpha_{rJ}}\right)^{\eta-1}$ is larger then it is optimal to use relatively more manual labor in that sector. So for example *routinization*, i.e. the replacement of routine workers by certain technologies, would be captured by an increase in $\left(\frac{\alpha_{mJ}}{\alpha_{rJ}}\right)^{\eta-1}$ and in $\left(\frac{\alpha_{aJ}}{\alpha_{rJ}}\right)^{\eta-1}$ in all sectors J.

The firm first order conditions also pin down the price of sector J output in terms of wage rates:

$$p_J = \left[\alpha_{mJ}^{\eta-1} \frac{1}{w_m^{\eta-1}} + \alpha_{rJ}^{\eta-1} \frac{1}{w_r^{\eta-1}} + \alpha_{aJ}^{\eta-1} \frac{1}{w_a^{\eta-1}}\right]^{\frac{1}{1-\eta}}. \tag{4}$$

Finally using (2), (3) and (4) to express sector J output, optimal sectoral labor use can be expressed as:

$$l_{mJ} = \left[\frac{p_J \alpha_{mJ}}{w_m}\right]^{\eta} \frac{Y_J}{\alpha_{mJ}}, \tag{5}$$

$$l_{rJ} = \left[\frac{p_J \alpha_{rJ}}{w_r}\right]^{\eta} \frac{Y_J}{\alpha_{rJ}}, \tag{6}$$

$$l_{aJ} = \left[\frac{p_J \alpha_{aJ}}{w_a}\right]^{\eta} \frac{Y_J}{\alpha_{aJ}}. \tag{7}$$

2.2 Households – occupational choice and demand for goods

The economy is populated by a unit measure of workers, who each have an idiosyncratic cost for entering each occupation, but can freely move between the three sectors, low-skilled services, goods, or high-skilled services, implying that in equilibrium, occupational wage rates must equalize across sectors. The cost that individuals pay for entering an occupation is redistributed in a lump-sum fashion. Since the consumption decisions are taken by the stand-in household, individuals choose the occupation that

9

provides them with the highest income. Thus an individual i chooses occupation o if

$$w_o - \chi_o^i \geq w_k - \chi_k^i \quad \text{for any} \quad k \neq o, \quad k, o \in \{m, r, a\},$$

where w_o is the unit wage in occupation o and χ_o^i is individual i's cost of entering occupation o. Since only the cost differences matter, we define $\chi_1^i \equiv \chi_r^i - \chi_m^i$ and $\chi_2^i \equiv \chi_a^i - \chi_m^i$. The optimal occupational choice is summarized in Figure 1.
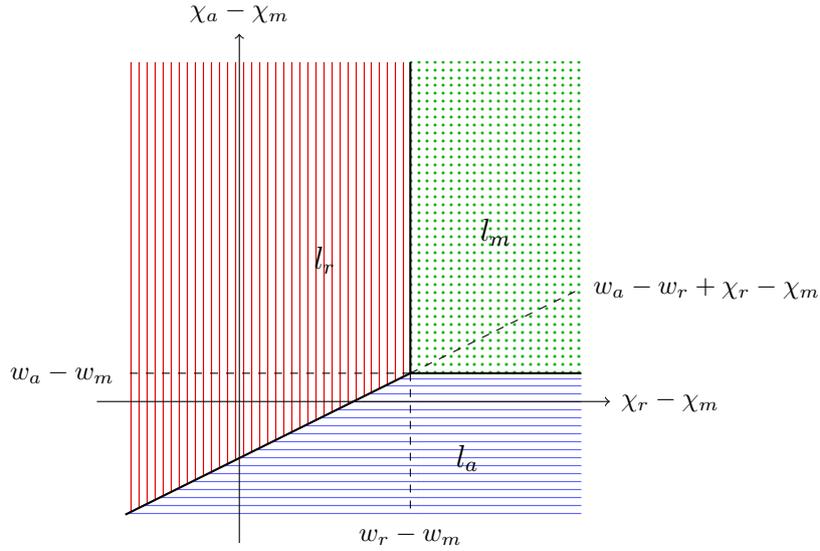

Figure 1: Optimal occupational choice

Notes: The graph shows the optimal selection of individuals into manual, routine and abstract occupations in terms of their idiosyncratic occupational cost differences as a function of occupational unit wages w_m, w_r, w_a.

Given the optimal occupational choice the fraction of labor supplied in the three occupations is given by:

$$l_m = \int_{w_r - w_m}^{\infty} \int_{w_a - w_m}^{\infty} f(\chi_1, \chi_2) d\chi_1 d\chi_2, \tag{8}$$

$$l_r = \int_{-\infty}^{w_r - w_m} \int_{w_a - w_r + \chi_1}^{\infty} f(\chi_1, \chi_2) d\chi_1 d\chi_2, \tag{9}$$

$$l_a = \int_{0}^{\infty} \int_{-\infty}^{\min\{w_a - w_r + \chi_1, w_a - w_m\}} f(\chi_1, \chi_2) d\chi_1 d\chi_2, \tag{10}$$

where $f(\chi_1, \chi_2)$ is the joint probability density function of the occupational cost differences.

The workers are organized into a stand-in household, which collects all income,

and makes utility maximizing choices in terms of sectoral consumption. The stand-in household solves the following problem:

$$\max_{c_L, c_G, c_H} \left(a_L (c_L + \bar{c}_L)^{\frac{\varepsilon-1}{\varepsilon}} + a_G c_G^{\frac{\varepsilon-1}{\varepsilon}} + a_H (c_H + \bar{c}_H)^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

$$\text{s. t.} \quad p_L c_L + p_G c_G + p_H c_H \leq l_m w_m + l_r w_r + l_a w_a$$

where $\bar{c}_L$ and $\bar{c}_H$ allow for non-homotheticity in consumption demands, and $\varepsilon < 1$, implying that goods and services are complements in consumption. We further assume that $a_L + a_G + a_H = 1$. The price of low-skilled services is denoted by p_L, that of goods is denoted by p_G, while that of high-skilled services by p_H. Assuming that the household is rich enough to consume all types of goods and services (i.e. an interior solution), optimality implies the following demand schedule:

$$C_L = \left(\frac{a_L}{p_L} \right)^\varepsilon \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^\varepsilon p_L^{1-\varepsilon} + a_G^\varepsilon p_G^{1-\varepsilon} + a_H^\varepsilon p_H^{1-\varepsilon}} - \bar{c}_L, \tag{11}$$

$$C_G = \left(\frac{a_G}{p_G} \right)^\varepsilon \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^\varepsilon p_L^{1-\varepsilon} + a_G^\varepsilon p_G^{1-\varepsilon} + a_H^\varepsilon p_H^{1-\varepsilon}}, \tag{12}$$

$$C_H = \left(\frac{a_H}{p_H} \right)^\varepsilon \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^\varepsilon p_L^{1-\varepsilon} + a_G^\varepsilon p_G^{1-\varepsilon} + a_H^\varepsilon p_H^{1-\varepsilon}} - \bar{c}_H. \tag{13}$$

2.3 Equilibrium

There are six markets in this economy: three labor markets, that of manual, routine and abstract labor; and three goods markets, that of low-skilled services, goods, and high-skilled services. There are six corresponding prices, out of which we normalize one without loss of generality, $w_r = 1$. The equilibrium is then defined as a set of prices, w_m, w_a, p_L, p_G, p_H, for which all markets clear.

Goods market clearing requires that $Y_L = C_L$, $Y_G = C_G$, and $Y_H = C_H$. Note that sectoral prices, p_J, and sectoral demands, C_J, depend on the endogenous occupational wage rates, w_m, w_r and w_a, as given in (4) and (11), (12), and (13). Then from (5), (6), or (7) optimal occupation o labor use in sector J can be expressed as a function of manual

and abstract wage rates:

$$l_{oJ}(w_m, w_a) = \left[\frac{p_J \alpha_{oJ}}{w_o}\right]^\eta \frac{C_J}{\alpha_{oJ}} \quad \text{for } o \in \{m, r, a\} \quad \text{and } J \in \{L, G, H\}.$$

The equilibrium then boils down to finding wage rates w_m and w_a such that the labor markets clear:[9]

$$l_{mL}(w_m, w_a) + l_{mG}(w_m, w_a) + l_{mH}(w_m, w_a) = l_m,$$

$$l_{rL}(w_m, w_a) + l_{rG}(w_m, w_a) + l_{rH}(w_m, w_a) = l_r.$$

3 Calibration

We need to calibrate the sectoral production functions, the distribution of the costs of entering the different occupations, and the utility function. In our model setup, there is a dichotomy that allows to back out the sector-occupation cell productivities from the data using only the production side. We therefore proceed in the following steps, similarly to Buera, Kaboski, and Rogerson (2015). First, we compute cell productivities taking as given the occupational wage rates and employment shares, as well as the sectoral income shares, in order to match in each period the income share of different occupations within each sector, the relative sectoral prices, and the overall growth rate of the economy. Second, we calibrate the distribution of costs such that it allows us to match occupational employment shares and wages in the initial and final period. Finally, we calibrate the utility function such that the model matches the sectoral income shares in the initial and final period.

3.1 Calibration targets

We use US Census and American Community Survey (ACS) data between 1960 and 2010 from IPUMS, provided by Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek (2010), to calculate occupational wage rates and occupational labor income

---

[9]If the manual and the routine markets clear, then the market for abstract labor clears as well due to Walras' law.

shares within sectors, as well as each sector's share in labor income.[10] For these calculations, we categorize workers into our three sectors based on their industry code (*ind1990*), and into our three occupations based on a harmonized and balanced panel of occupational codes as in Autor and Dorn (2013) and Bárány and Siegel (2018).

We calculate manual and abstract wage rates as the average hourly wage of a narrowly defined group – 25 to 29 year old men – in the given occupation divided by that in the routine occupation. This is in line with our normalization of $w_r = 1$. We rely on this measure – rather than on the average hourly wage of all workers within an occupation – to limit the potential influence of compositional changes, for example due to differential changes in the demographic composition of workers across occupations.[11] The occupational wage rate targets are calculated as:

$$w_m \equiv \frac{\text{average hourly wage of 25–29 year old men in manual jobs}}{\text{average hourly wage of 25–29 year old men in routine jobs}},$$

$$w_a \equiv \frac{\text{average hourly wage of 25–29 year old men in abstract jobs}}{\text{average hourly wage of 25–29 year old men in routine jobs}}.$$

We calculate the labor income share of occupation o in sector J as the ratio of total labor income of workers in occupation o and sector J relative to the total labor income of all workers in sector J:

$$\theta_{oJ} \equiv \frac{\text{earnings of occupation } o \text{ workers in sector } J}{\text{earnings of sector } J \text{ workers}}.$$

We can express occupational labor supply shares as:

$$l_o \equiv \frac{\frac{\text{earnings of workers in occupation } o}{w_o}}{\sum_{\tilde{o}} \frac{\text{earnings of workers in occupation } \tilde{o}}{w_{\tilde{o}}}},$$

these are equivalent to occupational labor supplies in the model, as total labor supply

---

[10]In our model the sectoral labor income shares are equal to value added shares as there are no other factors of production.

[11]This is similar to Buera et al. (2015), and it implies that all differences within an occupational group in hourly wages are due to differences in the endowment of efficiency units of labor. Given that we do not explicitly model heterogeneity in efficiency labor across individuals, the way we model selection implies that selection into occupations is orthogonal to efficiency.

is normalized to one. Finally, we calculate sectoral income shares as

$$\Psi_J \equiv \frac{\text{earnings of workers in sector } J}{\text{total earnings}}.$$

We use data from the U.S. Bureau of Economic Analysis (BEA) between 1960 and 2010 to get sectoral prices and the growth rate of GDP per worker between periods.

Table 4 in the Appendix contains all the calibration targets, and these are also plotted along with the model outcomes in section 5.

3.2 Extracting sector-occupation cell productivities

As mentioned before, given the structure of the model we can infer the productivity parameters directly from the data, without having to rely on a parameterization of the model's household side. We can do this conditional on a value for the elasticity of substitution in production between different types of labor.

We calculate the nine cell-specific productivity parameters, the αs, in each period. We back these out directly from nine targets: the labor income share of different occupations within each sector, the relative sectoral prices, and the overall growth of the economy. We calibrate these taking as given occupational wage rates, occupational labor supplies, and the sectoral distribution of income. Our model allows us to express the cell-specific productivity parameters as a function of the above data targets and the elasticity of substitution in production.

In particular, given occupational wages, the labor income share of different occupations within a sector pin down the ratios of αs within sectors in each period from the firm's optimality conditions (2) and (3):

$$\frac{\alpha_{mJ}}{\alpha_{rJ}} = \left(\frac{\theta_{mJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_m}{w_r}, \tag{14}$$

$$\frac{\alpha_{aJ}}{\alpha_{rJ}} = \left(\frac{\theta_{aJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_a}{w_r}. \tag{15}$$

The sectoral relative prices pin down (from (4)) the αs across sectors within each

14

period, again given occupational wages:

$$\frac{\alpha_{mJ}}{\alpha_{mK}} = \frac{p_K}{p_J}\left(\frac{\theta_{mJ}}{\theta_{mK}}\right)^{\frac{1}{\eta-1}},$$ (16)

where we also used the expressions for the relative αs within sectors.

Finally, the overall growth rate of output per worker pins down the evolution of the αs over time, given the distribution of income across sectors and occupational labor supplies. Appendix A.2 shows the full derivations.

To our knowledge, there is no consensus in the literature on the value of the elasticity of substitution in production between the different types of occupational labor within sectors. For this reason, we calibrate our model for a whole range of elasticities in $[0.1, 1.9]$. We do not calibrate the model for $\eta = 1$, as for that elasticity the model – in contrast to the data – would predict constant occupational labor income shares within sectors and our strategy for extracting the cell-level productivities which exploits variation in these shares would not go through. Extracting the cell level productivity series for η close to one is not a problem. However, notice that when η is close to one, in order to replicate the observed variations in occupational income shares within sectors and in sectoral prices, cell productivities are required to vary hugely both within and across sectors, as equations (14), (15), and (16) show. This huge variation in productivities would make the calibration of the preferences numerically very difficult for values of η close to 1.

3.3 Calibration of the cost distribution and of the consumption side

To close the model we need to parameterize the household side. It is important to note that these choices matter only for model simulations but not for assessing the contributions of sector and occupation factors to productivity growth at the sector-occupation cell level.

To calibrate the distribution of cost differences, we assume that $f(\chi_1, \chi_2)$ is a time-invariant bivariate normal distribution,[12] and we fix the correlation parameter to be

---

[12]For simplicity we assume that the distribution is time invariant. Allowing for changing costs (for example as in Caselli and Coleman (2001)) would require more parameters to be calibrated, and

0.4. Given this correlation, we calibrate the two means and the diagonal elements of the variance-covariance matrix such that in the initial and final period for given unit wages the cost distribution is able to match the employment shares. This calibration procedure by construction limits the importance of the correlation parameter, as the initial and final period outcomes are guaranteed to be the same regardless of its value. Our robustness checks on the value of the correlation parameter confirm that it is neither qualitatively, nor quantitatively important; see Appendix A.4.

Finally we calibrate the preference parameters of the model. Following Ngai and Pissarides (2007), we set the elasticity of substitution in consumption between the different sectoral outputs to $\varepsilon = 0.2$, implying that goods and the two types of services are complements. Given all the production side parameters, and the distribution of costs we calibrate $\bar{c}_L$, $\bar{c}_H$, a_L, and a_H (with $a_G = 1 - a_L - a_H$) to match the distribution of the sectoral income shares in the initial and final year, i.e. in 1960 and 2010. This also guarantees that the relative occupational wages in 1960 and 2010 are met in equilibrium. Again, this procedure by construction limits the importance of the value of the elasticity of substitution in consumption, our robustness checks in Appendix A.4 support this.

Table 1: Calibrated parameters

	Description	Value			
ε	elasticity of substitution in consumption	0.2			
ρ	correlation of cost differences	0.4			
μ_1, μ_2	mean of cost distribution	(-0.01, 0.53)			
σ_1^2, σ_2^2	variance of cost distribution	(0.03, 0.30)			
η	elasticity of substitution in production	**0.3**	**0.6**	**1.4**	**1.7**
$\bar{c}_L$	non-homotheticity term in L	0.0445	0.0036	458.38	36.887
$\bar{c}_H$	non-homotheticity term in H	0.0717	0.0058	738.19	59.404
a_L	weight on L	0.0916	0.0916	0.0916	0.0916
a_H	weight on H	0.9076	0.9076	0.9076	0.9076

Notes: The top panel shows calibrated parameters that are common across all values of the elasticity of substitution in production (η), whereas the bottom panel shows the parameters that vary with η for selected values.

Table 1 contains all the calibrated parameters of the model for selected (but for sake of readability not all) values of the elasticity of substitution in production. These,

it would not affect the sector-occupation cell productivities, neither their decomposition into various components, nor the results from the baseline model.

together with the evolution of the αs as backed out from the data fully specify the calibrated model.

It is important to note that the occupational wages and the sectoral income shares are only matched by the model in the initial and final period; in between they are not matched, i.e. also the occupational labor income shares and relative prices are not perfectly matched, as these were not targeted in the calibration. However, the model does reasonably well in matching these statistics in all periods, see Figures 5b and 7 (for $\eta = 0.6$).

4   Factor model decomposition

We set up a factor model to decompose the productivity growth of sector-occupation specific productivities – identified in the previous section – to a sector, an occupation, and a neutral component, as well as a residual. In particular we regress the log difference in the cell productivities, defined as $\Delta \ln \alpha_{oJ,t} = \ln \alpha_{oJ,t} - \ln \alpha_{oJ,t-1}$ on a (potentially time-varying) sector effect ($\gamma_{J,t}$), an occupation effect ($\delta_{o,t}$), and a time effect (β_t) in the following way

$$\Delta \ln \alpha_{oJ,t} = \beta_t + \gamma_{J,t} + \delta_{o,t} + \varepsilon_{oJ,t}, \tag{17}$$

where we use weights $\omega_{oJ,t}$ to reflect the relative importance of the sector-occupation cell. In our baseline specification we use the average labor income share of each cell between period $t - 1$ and t.[13] The sector effect $\gamma_{J,t}$ captures sector-wide innovations that affect the labor productivity of all workers in that sector equally regardless of their occupation. Productivity changes that are common to workers of a given occupation, but are independent from the sector, are assigned to $\delta_{o,t}$. Productivity changes common to all cells are captured by β_t, which can be interpreted as technological advances due to general purpose technologies, whereas $\varepsilon_{oJ,t}$ is the residual reflecting productivity changes idiosyncratic to workers in a sector-occupation cell. The sector dummy that is omitted from the regression is the one for the low-skilled service sector and the

---

[13]To be precise, we use as weights $\omega_{oJ,t} = (\Psi_{J,t}\theta_{oJ,t} + \Psi_{J,t-1}\theta_{oJ,t-1})/2$, where the values are given in Appendix Table 4. The results are very robust to alternatives, such as using employment shares, or using year $t - 1$ or year t shares, rather than averages.

omitted occupation dummy is the manual one. The estimated coefficients are therefore productivity growth rates of the sector-occupation cells relative to the one formed by the goods sector and manual occupation.

All factors together predict the following productivity series:

$$\widehat{\ln \alpha}_{oJ,0} = \ln \alpha_{oJ,0},$$

$$\widehat{\ln \alpha}_{oJ,t} = \widehat{\ln \alpha}_{oJ,t-1} + \widehat{\beta}_t + \widehat{\gamma}_{J,t} + \widehat{\delta}_{o,t}. \tag{18}$$

This productivity series contains both the sector- and the occupation-components. Its difference from the actual productivity series is due to the productivity component that is specific to the sector-occupation cell, $\varepsilon_{oJ,t}$. From now on we refer to this series as the 'full factor' prediction.

To gauge how much of the variation in cell productivities the sector- and occupation-specific components explain jointly, we calculate the R^2 of this prediction. We also want to quantify the importance of the sector and the occupation component separately. To do this we generate a sector-only productivity series, where on the one hand we shut down all cell-level productivity growth differences that come from the occupation-component, on the other hand we keep the average growth stemming from the occupation-component to ensure that overall growth is in line with the data. We therefore generate the following 'sector-only' productivity series, where within a sector productivity growth is the same for all occupations:

$$\widehat{\ln \alpha}_{oJ,0}^{sec} = \ln \alpha_{oJ,0},$$

$$\widehat{\ln \alpha}_{oJ,t}^{sec} = \widehat{\ln \alpha}_{oJ,t-1}^{sec} + \widehat{\beta}_t + \widehat{\gamma}_{J,t} + \frac{\omega_{r,t}\widehat{\delta}_{r,t} + \omega_{a,t}\widehat{\delta}_{a,t}}{\omega_{m,t} + \omega_{r,t} + \omega_{a,t}}, \tag{19}$$

where $\omega_{\tilde{o},t} = \sum_{\tilde{J}} \omega_{\tilde{o}\tilde{J},t}$. Across sectors cell productivity growth can differ only due to $\widehat{\gamma}_{J,t}$, the sector-specific component, as $\frac{\omega_{r,t}\widehat{\delta}_{r,t} + \omega_{a,t}\widehat{\delta}_{a,t}}{\omega_{m,t} + \omega_{r,t} + \omega_{a,t}}$ is the same for all cells and effectively acts like a neutral productivity growth component.[14] An alternative would be to run a naive factor model with only time and sector dummies, and construct the

---

[14]Notice that $\omega_{m,t}\widehat{\delta}_{m,t}$ does not appear in the numerator. This is because the omitted occupation is the manual one, hence there is no $\widehat{\delta}_{m,t}$; it is subsumed in the time effect $\widehat{\beta}_t$.

sector-only productivity component from it. This method however, by a priori restricting productivity growth to be biased only across sectors, would also pick up some of the differences in growth rates across occupations as sectors use occupations at different intensities. Therefore it would not be informative about the true sector-specific component of cell level productivity growth. We show the results of such a factor decomposition, which we refer to as 'naive', in Appendix A.5.

Similarly we generate an 'occupation-only' productivity series as:

$$\widehat{\ln \alpha}_{oJ,0}^{occ} = \ln \alpha_{oJ,0},$$

$$\widehat{\ln \alpha}_{oJ,t}^{occ} = \widehat{\ln \alpha}_{oJ,t-1}^{occ} + \widehat{\beta}_t + \widehat{\delta}_{o,t} + \frac{\omega_{G,t}\widehat{\gamma_{G,t}} + \omega_{H,t}\widehat{\gamma_{H,t}}}{\omega_{L,t} + \omega_{G,t} + \omega_{H,t}}, \tag{20}$$

where $\omega_{\tilde{J},t} = \sum_{\tilde{o}} \omega_{\tilde{o}\tilde{J},t}$. The last term assigns the average year-t sector component across all cells, and thus shuts down differences in cell productivity growth across sectors within occupations, while ensuring that overall growth is in line with the data. Finally, we generate a 'neutral' productivity change series as:

$$\widehat{\ln \alpha}_{oJ,0}^{neut} = \ln \alpha_{oJ,0}, \tag{21}$$

$$\widehat{\ln \alpha}_{oJ,t}^{neut} = \widehat{\ln \alpha}_{oJ,t-1}^{neut} + \widehat{\beta}_t + \frac{\omega_{r,t}\widehat{\delta_{r,t}} + \omega_{a,t}\widehat{\delta_{a,t}}}{\omega_{m,t} + \omega_{r,t} + \omega_{a,t}} + \frac{\omega_{G,t}\widehat{\gamma_{G,t}} + \omega_{H,t}\widehat{\gamma_{H,t}}}{\omega_{L,t} + \omega_{G,t} + \omega_{H,t}},$$

where besides the time effect $\widehat{\beta}_t$ we also include the average occupation and the average sector component. This technological change affects all cells equally, i.e. it is a neutral technological change.

We can use these predictions to evaluate the quantitative importance of each factor. Table 2 shows how much of the change in cell productivities is explained by each factor for various values of the elasticity of substitution between different occupations, η.[15] This table shows the R^2 of the factor model regression (17) in the first column for a range of η values. The explanatory power of sector- and occupation-specific compo-

---

[15]The αs themselves change as we back out cell productivities conditional on the value of the elasticity, but it is important to bear in mind that this series is independent of any other part of the model. A different η implies a different parametrization of the household side to match the 1960 and 2010 data (see the various columns in the lower panel of Table 1), but that part of the model does not affect the analysis of cell productivities.

η	full factor	sector	occupation	neutral
0.1	0.806	0.307	0.486	0.047
0.2	0.791	0.249	0.527	0.040
0.3	0.776	0.192	0.571	0.033
0.4	0.763	0.140	0.616	0.025
0.5	**0.752**	**0.094**	**0.661**	**0.018**
0.6	**0.744**	**0.060**	**0.703**	**0.012**
0.7	**0.739**	**0.038**	**0.742**	**0.007**
0.8	**0.738**	**0.032**	**0.773**	**0.003**
0.9	**0.741**	**0.041**	**0.797**	**0.001**
1.1	0.759	0.100	0.819	0.001
1.2	0.771	0.145	0.817	0.003
1.3	0.785	0.196	0.808	0.006
1.4	0.800	0.251	0.795	0.010
1.5	0.815	0.307	0.777	0.015
1.6	0.829	0.362	0.757	0.020
1.7	0.843	0.414	0.736	0.024
1.8	0.856	0.464	0.714	0.029
1.9	0.869	0.510	0.692	0.034

Table 2: R^2 of the decomposition

nents jointly is between 74 and 87 percent, implying that between 13 and 26 percent of the variation is due to effects idiosyncratic to the sector-occupation cell. The second and the third column show the variation in cell productivity changes explained by the sector-only and respectively by the occupation-only component, while last column shows the variation explained by neutral technological progress. The occupation component is the most important, followed by the sector component, while the neutral component has the lowest explanatory power. Goos et al. (2014), estimate, Duernecker and Herrendorf (2016) and Lee and Shin (2017) calibrate the elasticity of substitution to be between 0.5 and 0.9. In this range, our decomposition shows that 24-25 percent of cell productivity growth is idiosyncratic to the sector-occupation cell, 66-80 percent can be attributed to the occupation-specific component, 3-9 percent to the sector-specific component, while 0-2 percent of productivity growth is neutral. Our interpretation of this decomposition is that most productivity changes are biased across occupations, across sector-occupation cells, and to some extent also across sectors.[16]

---

[16]We also have conducted this decomposition using Current Population Survey (CPS) data over 1968–2016. The overall explanatory power of the full factor model is smaller in that dataset, but in terms of the role of the sector and the occupation components the results are very similar. We prefer to

When evaluating the explanatory power of the sector-only and the occupation-only predictions, it is important to bear in mind that the series in (19) and (20) are not equivalent to the predictions of a factor model with a time and a sector or respectively a time and an occupation component only.[17] As already mentioned, those productivity series would pick up differential productivity growth across sectors (or occupations) that originates from the sectors using occupations at different intensities (or the occupations being used at different intensities across sectors).[18]
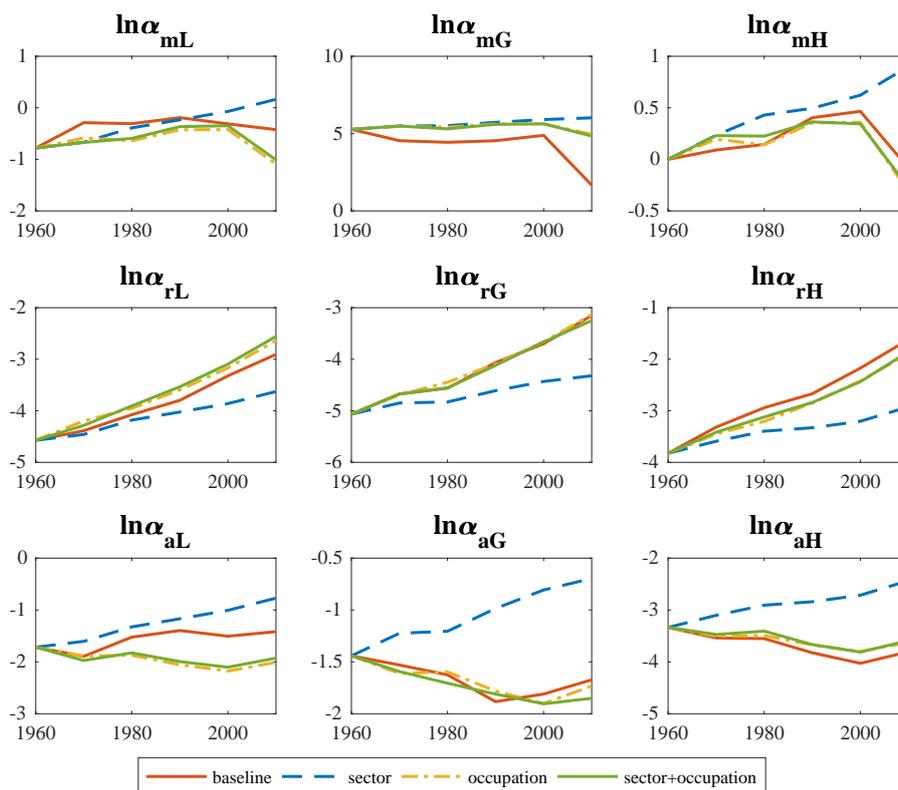

Figure 2: Baseline and counterfactual cell productivity

In section 5 we study the role of these components in the evolution of various outcomes. For now we illustrate, in Figure 2, the path of cell productivities as extracted

---

use the Census for our analysis as it starts earlier, is a larger dataset, and occupational codes have been harmonized in previous studies (Autor and Dorn (2013) and Bárány and Siegel (2018)). The CPS results are available on request.

[17]Since the sector-only and the occupation-only productivity series are not generated as partial predictions, their R^2 are not limited by that of the full factor model.

[18]The R^2 from such factor models carry information on how much of the cell productivity growth could be captured by 'naive' models, which ignore the potential existence of the sector or of the occupation component and load them on the respective other component. Table 5 in the Appendix shows these values.

from the data, as well as the different predicted productivities based on the factor model for $\eta = 0.6$. Note however, that as the extracted αs themselves change with η, so do the different components and their explanatory power. The red solid lines show the baseline cell productivities, while the counterfactual cell productivity series are the following: the green solid lines are based on the sector and the occupation components $(\widehat{\ln \alpha}_{oJ,t})$, the blue dashed lines are based on the differences in the sector components only $(\widehat{\ln \alpha}_{oJ,t}^{sec})$, while the yellow dashed-dotted lines are based on the occupation components only $(\widehat{\ln \alpha}_{oJ,t}^{occ})$. For this value of η the figure shows that the full factor and the occupation-only predictions are quite close to each other and to the baseline, whereas the sector-only predictions are further away. This is reflected in the R^2, for $\eta = 0.6$ the occupation-only and the full factor are both around 70 percent, whereas the R^2 of the sector-only component is only 6 percent.

5   Model vs data: the role of sector- and occupation-components

In this section we quantify the role of the different components of cell productivity growth on various outcomes of interest. Using the baseline cell productivities in our model and comparing the generated paths to those observed in the data informs us of how good the model is in describing the evolution of the economy. By feeding in the various counterfactual cell productivities generated from the decomposition in the previous section, we aim to measure the importance of the different components of productivity growth. The predictions from the full factor model (as in (18)) relative to the baseline allow us to measure the importance of the productivity growth component that is idiosyncratic to the cell, captured by $\varepsilon_{oJ,t}$ in the factor model (17).

By contrasting the predictions based on the sector-only component (as in (19)) with those of the full factor model we can measure the importance of productivity growth differences across occupations within a sector. It is worth to reiterate that the sector-only component is the part of cell productivity growth that is not originating from any differential change in productivities at the occupation or at the cell level. Even though sectors might differ (and indeed they do in the data) in the intensities at which they use particular occupations, this sector-only component will not pick up differ-

ential productivity across occupations, but is a measure of productivity growth after controlling for occupation effects.

Conversely, contrasting the predictions based on the occupation-only component (as in (20)) with those of the full factor model allows us to infer the role of productivity growth differences within occupations across sectors. Similarly to the sector-only component, the occupation-only component we construct does not pick up differential productivity growth across occupations that originates from them being used at intensities that vary across sectors.

5.1   Overall fit for different production elasticities

In this section we calculate for various outcomes the mean squared distance between the observed changes in the data and the model (under different cell productivity series) relative to the variance of the changes in the data, for a whole range of production elasticities.[19] This measure, while imperfect, summarizes the goodness of fit in a single number which can be easily compared across the alternative models and across different values of the elasticity of substitution in the production function. A value of zero implies a model that perfectly fits the data, and a larger value implies that model predictions are further away from the data.

We calculate the distance measure of the various models for a wide range of production elasticities, $\eta \in [0.1, 1.9]$.[20] Since there is no consensus about the value of this elasticity in the literature, we evaluate the role of the different components for various values. In what follows it becomes clear that some findings are extremely robust across different values of η, while for others whether η is larger than a certain value seems to be important.

---

[19]More precisely we define our distance measure as:

$$\frac{\sum_k \sum_{t=1960}^{2010} (\Delta x_{k,t}^d - \Delta x_{k,t}^m)^2}{\sum_k \sum_{t=1960}^{2010} (\Delta x_{k,t}^d - \Delta \overline{x}^d)^2} \geq 0,$$

where $\Delta x_{k,t}^i = x_{k,t}^i - x_{k,1960}^i$, $i = d$ denotes the data, and $i = m$ denotes the model prediction. The squared distance and variance is calculated for occupational measures pooled across $k = \{m, r, a\}$, for sectoral outcomes across $k \in \{L, G, H\}$, and for cell outcomes across all 9 cells $k \in \{m, r, a\} \times \{L, G, H\}$.

[20]As discussed in section 3.2, we do not consider the case of $\eta = 1$ since in this case the implied occupational income shares within sectors would be constant over time which is at odds with the data and invalidates our identification of baseline cell productivities.
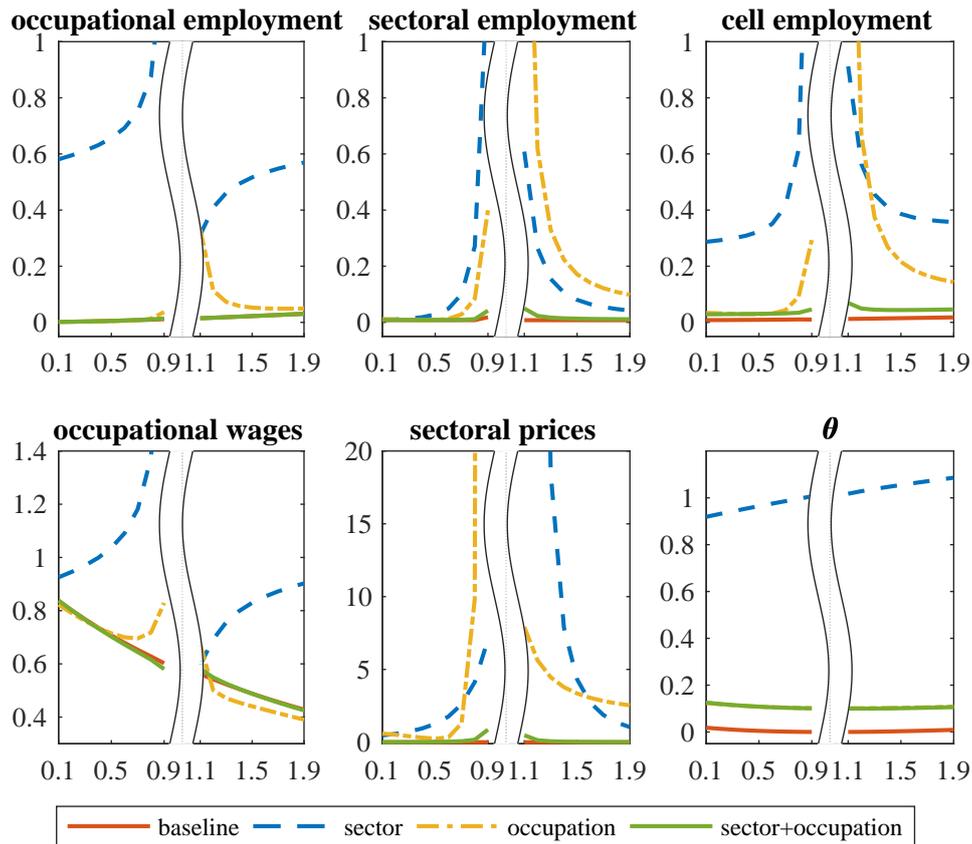
Figure 3: Productivity components' ability to replicate the data

The first thing to note in Figure 3 is that our baseline model does very well in matching the data. It is important to recall that we extracted the baseline productivities from the data using the production side of our model, taking as given the evolution of occupational wages and employment, and sectoral labor income shares. We then calibrated the time-invariant parameters of the model to match these values in the initial and final period perfectly. However, in the interim periods there could be differences between the model and the data. These differences turn out to be small for almost all outcomes of interest, as the distance measure is basically zero for all values of η, implying that our baseline model matches the data almost perfectly in all periods. This is not the case for occupational wages, where our baseline model's predictions in the interim periods deviate from the data. Our model's failure to match these paths can be understood by looking at Figure 4b below, where the dark gray solid line shows the data and the red solid line the values predicted by our baseline model (for $\eta = 0.6$). The data, as our model, displays a strong upward trend in both manual and abstract

wages relative to routine. In the data, however the 1980 values of these relative wages seem to be outliers, which might correspond to the compression of the skill premium during the 1970s. Our model stays silent about what generated these.

The second thing to note is that the model based on the productivity growth predictions of the full factor model does almost as well as our baseline model. The only difference between these two models is that the latter does not contain the productivity growth that is idiosyncratic to the sector-occupation cell. The fact that these two models perform equally well for occupational employment and wages, for sectoral employment and for sectoral prices (except for values of η close to 1) implies that the productivity growth component idiosyncratic to the cell is not the key driver of these outcomes. This is not the case for cell employment shares (l_{oJ}) and for occupational income shares within sectors (θ_{oJ}). Comparing the distance measure of the baseline and of the full factor model, it is clear that productivity growth that is idiosyncratic to the cell plays an important role in generating the path of these in the baseline model and hence in the data. We return to this in the next subsection.

The picture is less clear in terms of the role of the occupation-only and the sector-only component of productivity growth. In particular, a different message emerges for low values of η and for values of η close to and above 1. For low values of the elasticity of substitution the model based on the occupation-only component of productivity growth matches the predictions from the full factor model almost perfectly (and therefore it also matches the baseline model and the data very well). This suggests that if the elasticity of substitution between occupations is low, the sectoral component of productivity growth has almost no role except in the evolution of sectoral prices. Corroborating this, Figure 3 also shows that the fit of the model based on the sector-only component of growth is substantially worse than of all other models in terms of occupational employment and wages, and occupational income shares within sectors. However, for higher values of η the distance measure of the model generated from the occupation-only and from the sector-only component of productivity growth for sectoral employment and prices and for cell employment are similar. This implies that if η is relatively high, both the occupation and the sector component of productivity growth play an important role in these outcomes. Finally, notice that for occupational

25

income shares within sectors the distance of the full factor model and the occupation-only component are virtually the same, implying that the sector component has a very limited role in the evolution of θ.

Our results so far suggest that the ability of the sector-only or the occupation-only component of cell productivity growth to explain the outcomes of interest, especially sectoral employment and prices, depends on the production elasticity of substitution. To better understand this dependence, we compute the average sectoral labor productivity growth implied by our model for each cell productivity series.[21] Table 3 shows these for various values of η for our benchmark model as well as for the constructed series from our decomposition. For reference we show in the last column annual sectoral labor productivity growth calculated from BEA data. While we have not used these productivity measures in the calibration of our model, the model predicted productivity growth rate for the goods and for the high-skilled service sectors are remarkably close to the BEA values. However, in the low-skilled service sector there is a discrepancy between the productivity growth implied by the model and the one measured in the data. This difference is because we inferred productivity growth to match the sectoral labor income shares calculated from the Census data, rather than to match value-added shares from the BEA. These two series have diverged over the period of our analysis.[22] As the focus of our paper is to understand sectoral and occupational labor market trends, we need data to distinguish not only the sector of workers, but also their occupation therein. We therefore cannot use only BEA data, and we inform our model calibration by the detailed data from the Census. In what follows we base our discussion on the comparison of sectoral productivity growth from the baseline model and the various counterfactuals.

As is well known, average productivity growth in the goods sector exceeds the one in the (two types of) service sectors according to the BEA data. This also holds in our baseline model, which predicts for all values of η the same sectoral average

---

[21]To be precise, we define sector J's labor productivity in period t as $\frac{Y_{J,t}}{l_{mJ,t}+l_{rJ,t}+l_{aJ,t}}$, compute its growth over 1960–2010 and report the annualized rate in Table 3.

[22]While we do not aim at explaining this divergence between income shares from the Census and value added shares in the BEA data, possible explanations for this might be sectoral differences in the labor share of income (i.e. differences in capital intensity across sectors), or differences in the gap between hourly wages and hourly labor costs, both of which might possibly vary over time.

Table 3: Implied sectoral average labor productivity growth (in percent)

η	0.3	0.6	1.4	1.7	BEA
			baseline		
L	2.23	2.23	2.23	2.23	0.89
G	2.30	2.30	2.30	2.30	2.53
H	1.11	1.11	1.11	1.11	1.36
			full factor		
L	2.11	1.99	2.53	2.41	0.89
G	2.36	2.38	2.30	2.32	2.53
H	1.14	1.18	0.97	1.01	1.36
			sector component		
L	2.06	1.91	2.63	2.48	0.89
G	1.86	1.50	3.19	2.82	2.53
H	1.46	1.82	0.15	0.51	1.36
			occupation component		
L	1.79	1.82	1.65	1.69	0.89
G	2.25	2.62	0.88	1.25	2.53
H	1.42	1.11	2.58	2.27	1.36

growth rates. The model based on the full factor productivities robustly predicts that goods sector productivity growth is faster than high-skilled services, but the magnitude of the difference depends on the elasticity of substitution. Nonetheless it seems that the higher goods sector productivity growth is not driven by productivity components that are idiosyncratic to the sector-occupation cell. Our exercise can shed light on the origin of the higher labor productivity growth in the goods sector. It allows us to disentangle whether it is driven by the sector-specific component of cell productivity growth or whether it is driven by productivity components specific to individual occupations which are used at different intensities across sectors. Our results suggest that as long as occupations are complements in production, i.e. if $\eta < 1$, the occupation-component of productivity by itself implies that the goods sector has a higher productivity growth than the two types of services, for instance because it uses routine-labor most intensively and technological change is routine-biased. In fact in this range of the elasticity, the sector-specific component of productivity growth in goods might not even be the highest amongst all sectors. Conversely, if $\eta > 1$ the roles are reversed and the sector-only component is the driver behind the observed differences in sectoral productivity. However, $\eta < 1$ seems more plausible, as occupations

are in general characterized by different tasks which are believed to be complements to each other. In this case our results suggest that the observed technological change that is biased at the sector level[23], is in fact the outcome of productivity changes at the occupation level (as for instance argued by Duernecker and Herrendorf (2016)).

Given that from our distance measure, as long as it is not equal to zero, it is impossible to ascertain whether cell-productivity differences stemming solely from the sector-component, or solely from the occupation-component, or ignoring only the component idiosyncratic to the cell generate patterns qualitatively in line with the data. In the next section we show the predicted time path from our baseline model and the three counterfactual models against the data for one particular value of η. It is also worth noting that the conclusions we have drawn about the various productivity components' distance measure with respect to the various model outcomes is robust to alternative values of the elasticity of substitution in production (ε) and the correlation in the cost differences (ρ); see Appendix A.4.

5.2   Over time fit for a specific production elasticity

To illustrate in greater detail how each component of productivity affects the various model outcomes over time, we now fix the elasticity of substitution across occupational labor inputs. We set this parameter to $\eta = 0.6$ for two reasons. First, the literature typically assumes that the various occupations are complements to each other in production implying a value less than one, and this value is in the range used in previous literature.[24] The second reason for focusing on $\eta = 0.6$ is that it is around this value of the elasticity that the distance measure of the occupation-only model starts to diverge from that of the full factor model, especially in sectoral employment and prices, and hence analyzing the evolution of the various models for this value might

---

[23]This has been suggested to be the driver of the cost disease of services (e.g. Baumol (1967)) and structural change across sectors (e.g. Ngai and Pissarides (2007)).

[24]To our knowledge the only estimate of this elasticity is in Goos et al. (2014), who estimate it to be 0.53, 0.66, and 0.8 depending on the specification and the sample of countries; it is worth to note, however, that they estimate in partial equilibrium not taking into account aggregate effects. Duernecker and Herrendorf (2016) calibrate a value of 0.56, while Lee and Shin (2017) calibrate a value of 0.704 for this same parameter, though with a different set of occupations. The papers which introduced task biased technical change typically assume a Cobb-Douglas production function, i.e. an elasticity of 1 for simplicity (Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011), Autor and Dorn (2013)).

be more informative than for a very low value of η.
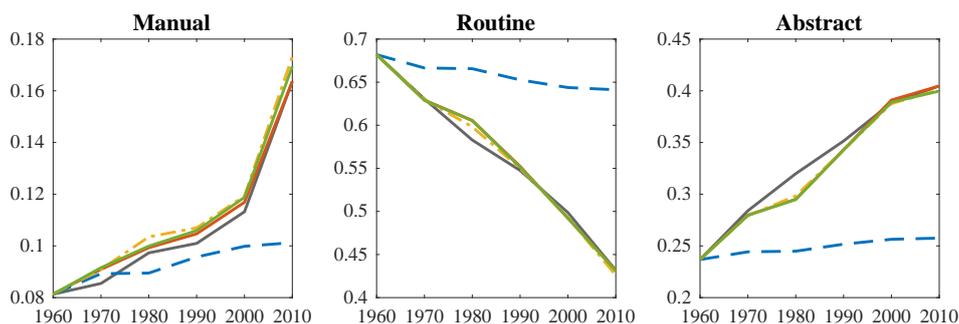
(a) Occupational employment
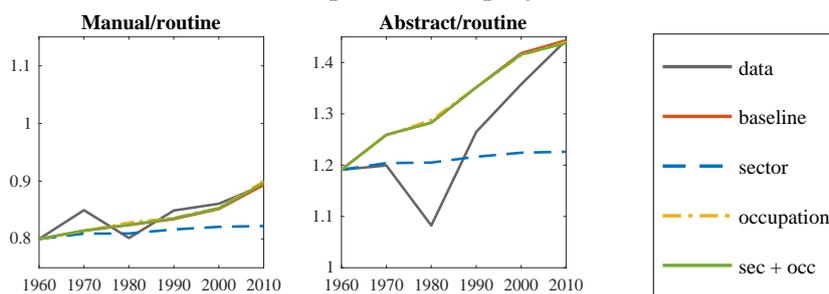

(b) Occupational relative wages

Figure 4: The evolution of occupational outcomes

Figures 4–7 show the evolution of the data in solid gray, contrasted with the economy's evolution for the various productivity paths color coded as before.[25] These figures confirm what the distance measures in Figure 3 showed. The baseline model's predictions are extremely close to the data except for occupational relative wages, where they do not pick up the drastic drop in the 1980 values. Figure 4 shows (i) that the predictions of the full factor and the occupation-only models are very close to each other and to the data for all occupational outcomes, and (ii) that while the sector-only model's predictions are qualitatively in line with the data, quantitatively they fall short. Figure 5 shows (i) that the full factor model does almost as well as our baseline model, (ii) that for sectoral employment both the occupation-only and the sector-only models are qualitatively and quantitatively close to the data, and (iii) that for sectoral prices neither the occupation-only nor the sector-only model does very well.

Figure 6 shows the predicted evolution of employment shares by sector-occupation

---

[25]The baseline cell productivities in solid red, the full factor model predictions in solid green, the dashed blue are based on the sector-only components, and the dashed-dotted yellow lines are generated from the occupation-only components.
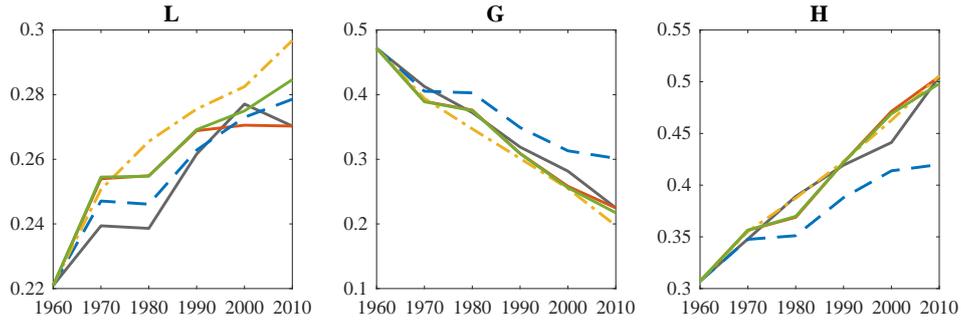
(a) Sectoral employment
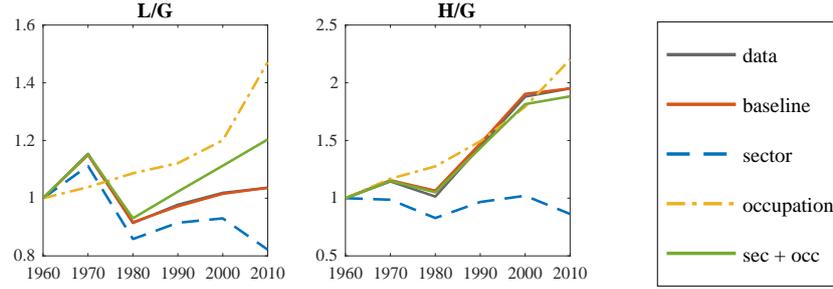

(b) Sectoral relative prices

Figure 5: The evolution of sectoral outcomes

cells. While all counterfactual cell productivities lead to outcomes that are similar to the data, the evolution of cell employment shares are furthest away in the case of sector-only cell productivities. For this value of η, the predictions of the occupation-only and the full factor model are very close to each other, but there is a gap between their predictions and the data for some cells. This implies that the idiosyncratic cell productivities play an important role in the evolution of cell employment shares, which is also reflected in Figure 3, where the distance measure for cell employment is bounded away from zero for all counterfactual cell productivities and for all values of η.

Figure 7 shows the predicted changes in labor income shares within sectors. This figure shows that the sector-only component predicts hardly any change in the θs. This can be understood from equations (22) and (23) in the Appendix: labor income shares change if relative occupational wages or relative cell productivities within a sector change. The sector-only model shuts down the second channel, and predicts quantitatively small changes in relative occupational wages, thus implying changes in the θs
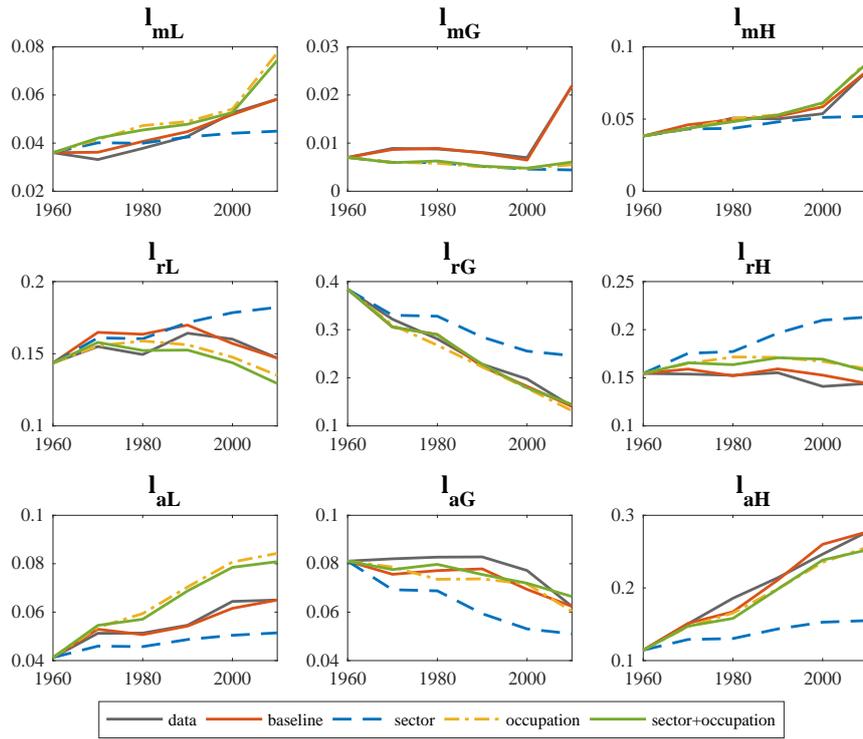
Figure 6: Cell employment shares

that are in line with the data, but which are quantitatively very small. Furthermore, in terms of the other counterfactual cell productivities the message is similar to Figure 6: the full factor model and the occupation-only component get the general patterns right, but there is a gap between the full factor and the baseline model/data pointing to the importance of the cell-specific productivity component. These predictions are reflected in Figure 3; while the distance measure for θ of the sector-only model is quite far from zero for all values of η, the full factor and occupation-only models are also bounded away from zero, implying that productivity changes idiosyncratic to the cell are important for changes in occupational labor income shares within sectors.

6 Conclusion

In this paper we propose a novel approach to infer and study the nature of potentially non-neutral technological change. We set up a parsimonious yet flexible model which we use to extract sector-occupation cell productivities from observed changes in labor
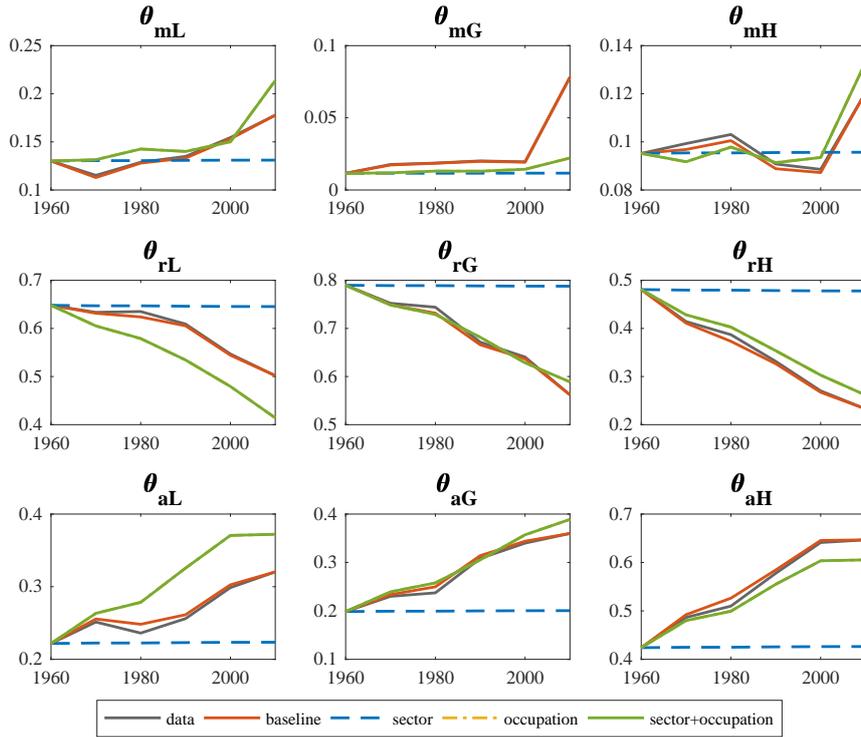
Figure 7: Income shares within sectors

income shares, sectoral prices and overall GDP growth over time.

We use a factor model on the extracted cell productivities to quantify to what degree technological change is biased across occupations and across sectors. For reasonable values of the production elasticity we find that 24-25 percent of the change in cell productivities is due to factors specific to the sector-occupation cell, i.e. is not driven by either sector- or occupation-specific factors. Within the about 75 percent of variation that the full factor model explains, the relative magnitude of the sector and the occupation component is quite sensitive to η (the lower η the more important the sector component), however, the occupation component's role is robustly larger.

To understand what each of these components imply for various outcomes of interest, we feed these as counterfactual productivity series into the model. While we see that sector and occupation productivity components by themselves generate occupational employment and wage, as well as sectoral employment paths qualitatively in line with the data, quantitatively the occupation component gets much closer. Moreover we find that the occupation component and the cell-specific elements are impor-

tant drivers of the occupational income shares within sectors and employment shares at the cell level. However both sector and occupation components are needed to explain the evolution of sectoral prices over time.

Overall, we infer that productivity growth is biased along both the occupational and the sectoral dimension. The occupation component goes a long way in explaining the reallocations across sectors and occupations. However, to understand the full picture of the evolution of prices, wages and employment not only across but also within sectors and occupations, models should allow for technologies to evolve at the sector-occupation cell level. While our model does not allow for any frictions and therefore does not warrant any policy interventions, the finding that virtually all of labor market outcomes are explained by the occupation component suggests that if policymakers wanted to respond to the observed reallocations, they should not focus on industrial policies but consider active labor market policies, including training programs.

References

Daron Acemoglu and David Autor. Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 4, Part B, pages 1043 – 1171. Elsevier, 2011.

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

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

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

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

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

William J. Baumol. Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis. *The American Economic Review*, 57(3):415–426, 1967.

Francisco J. Buera, Joseph P. Kaboski, and Richard Rogerson. Skill Biased Structural Change. Working Paper 21165, National Bureau of Economic Research, May 2015.

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

Georg Duernecker and Berthold Herrendorf. Structural Transformation of Occupation Employment. Working paper, February 2016.

Atish R. Ghosh and Holger C. Wolf. Geographical and Sectoral Shocks in the U.S. Business Cycle. Working Paper 6180, National Bureau of Economic Research, September 1997.

Maarten Goos and Alan Manning. Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics*, 89(1):118–133, February 2007.

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

Jeremy Greenwood, Zvi Hercowitz, and Per Krusell. Long-Run Implications of Investment-Specific Technological Change. *American Economic Review*, 87(3):342–62, 1997.

Lawrence Katz and Kevin M. Murphy. Changes in Relative Wages, 1963–1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

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

Miklós Koren and Silvana Tenreyro. Volatility and Development. *The Quarterly Journal of Economics*, 122(1):243–287, 2007.

Per Krusell, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. Capital-skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, 68 (5):1029–1053, 2000.

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

Guy Michaels, Ashwini Natraj, and John Van Reenen. Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years. *The Review of Economics and Statistics*, 96(1):60 – 77, March 2014.

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

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

Alan C. Stockman. Sectoral and national aggregate disturbances to industrial output in seven European countries. *Journal of Monetary Economics*, 21(2):387 – 409, 1988.

U.S. Bureau of Economic Analysis. Gross domestic product by industry 1947-2014. *retrieved from U.S. Bureau of Economic Analysis*, 2015.

U.S. Bureau of Economic Analysis. Hours worked by full-time and part-time employees. *retrieved from FRED, Federal Reserve Bank of St. Louis*, 2017.

A   Appendix

A.1   Classification

Occupations are classified as:

1. <u>Manual: low-skilled non-routine</u>

   housekeeping, cleaning, protective service, food prep and service, building, grounds cleaning, maintenance, personal appearance, recreation and hospitality, child care workers, personal care, service, healthcare support

2. <u>Routine</u>

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

3. <u>Abstract: skilled non-routine</u>

   managers, management related, professional specialty, technicians and related support

 Industries are classified into sectors in the following way:

1. <u>Low-skilled services:</u> personal services, entertainment, low-skilled transport (bus service and urban transit, taxicab service, trucking service, warehousing and storage, services incidental to transportation), low-skilled business and repair services (automotive rental and leasing, automobile parking and carwashes, automotive repair and related services, electrical repair shops, miscellaneous repair services), retail trade, wholesale trade

2. <u>Goods:</u> agriculture, forestry and fishing, mining, construction, manufacturing

3. <u>High-skilled services:</u> professional and related services, finance, insurance and real estate, communications, high-skilled business services, communications, utilities, high-skilled transport , public administration

Figure 8 shows the share of hours worked in each of the three sectors and also in each sector-occupation cell between 1960 and 2010 in the US. This figure demonstrates the patterns we described in the introduction.
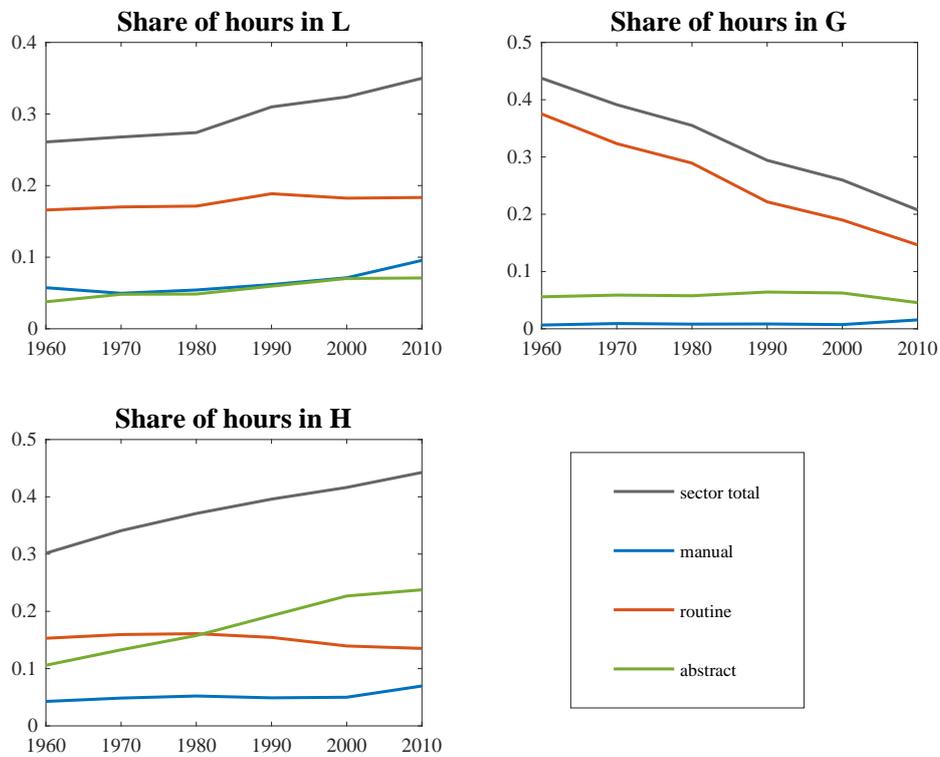

Figure 8: sector-occupation hours worked shares 1960-2007

Notes: The data is taken from IPUMS US Census data for 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) for 2010. For three broad sectors, low-skilled services (L), goods (G) and high-skilled services (H) and three occupational categories (manual, routine, abstract), this figure plots the evolution of the share of hours supplied in sector-occupation cells, as well as in sectors in the US between 1960–2010. The dark grey lines show the share of hours supplied in each sector, which are broken down into occupations within the sector in each panel.

A.2   Derivations

In this subsection we show how the αs can be a expressed as a function of observables. The labor income shares of different occupations within a sector pin down the αs

within a sector. To see this multiply (2) with w_m/w_r and (3) with w_a/w_r to get:

$$\frac{\theta_{mJ}}{\theta_{rJ}} = \left(\frac{w_r}{w_m}\right)^{\eta-1}\left(\frac{\alpha_{mJ}}{\alpha_{rJ}}\right)^{\eta-1}, \tag{22}$$

$$\frac{\theta_{aJ}}{\theta_{rJ}} = \left(\frac{w_r}{w_a}\right)^{\eta-1}\left(\frac{\alpha_{aJ}}{\alpha_{rJ}}\right)^{\eta-1}. \tag{23}$$

Re-arrange to get:

$$\frac{\alpha_{mJ}}{\alpha_{rJ}} = \left(\frac{\theta_{mJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}}\frac{w_m}{w_r},$$

$$\frac{\alpha_{aJ}}{\alpha_{rJ}} = \left(\frac{\theta_{aJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}}\frac{w_a}{w_r}.$$

The relative prices across sectors pin down the relative αs across sectors. To see this, use (4) for sectors J and K to get:

$$\frac{p_J}{p_K} = \frac{\alpha_{mK}}{\alpha_{mJ}}\left[\frac{\frac{1}{w_m^{\eta-1}} + \left(\frac{\alpha_{rJ}}{\alpha_{mJ}}\right)^{\eta-1}\frac{1}{w_r^{\eta-1}} + \left(\frac{\alpha_{aJ}}{\alpha_{mJ}}\right)^{\eta-1}\frac{1}{w_a^{\eta-1}}}{\frac{1}{w_m^{\eta-1}} + \left(\frac{\alpha_{rK}}{\alpha_{mK}}\right)^{\eta-1}\frac{1}{w_r^{\eta-1}} + \left(\frac{\alpha_{aK}}{\alpha_{mK}}\right)^{\eta-1}\frac{1}{w_a^{\eta-1}}}\right]^{\frac{1}{1-\eta}}.$$

Using the above expressions on the relative αs within sector and re-arranging we get:

$$\frac{\alpha_{mJ}}{\alpha_{mK}} = \frac{p_K}{p_J}\left(\frac{\theta_{mK}}{\theta_{mJ}}\right)^{\frac{1}{1-\eta}}.$$

The growth rate of the economy pins down the evolution of the αs over time. First, note that we express the evolution of cell productivities over time conditional on the sectoral income shares. The sectoral income shares, using equations (5), (6) and (7), can be expressed as:

$$\frac{\Psi_G}{\Psi_H} = \frac{p_G Y_G}{p_H Y_H} = \frac{l_{mG}p_G^{1-\eta}w_m^\eta\alpha_{mG}^{1-\eta}}{l_{mH}p_H^{1-\eta}w_m^\eta\alpha_{mH}^{1-\eta}},$$

$$\frac{\Psi_L}{\Psi_H} = \frac{p_L Y_L}{p_H Y_H} = \frac{l_{mL}p_L^{1-\eta}w_m^\eta\alpha_{mL}^{1-\eta}}{l_{mH}p_H^{1-\eta}w_m^\eta\alpha_{mH}^{1-\eta}}.$$

Re-arranging and using the above expressions to substitute out α_{mJ}/α_{mK}:

$$\frac{l_{mG}}{l_{mH}} = \frac{\Psi_G}{\Psi_H}\left(\frac{p_H}{p_G}\right)^{1-\eta}\left(\frac{\alpha_{mH}}{\alpha_{mG}}\right)^{1-\eta} = \frac{\Psi_G}{\Psi_H}\frac{\theta_{mG}}{\theta_{mH}},$$

$$\frac{l_{mL}}{l_{mH}} = \frac{\Psi_L}{\Psi_H}\left(\frac{p_H}{p_L}\right)^{1-\eta}\left(\frac{\alpha_{mH}}{\alpha_{mL}}\right)^{1-\eta} = \frac{\Psi_L}{\Psi_H}\frac{\theta_{mL}}{\theta_{mH}}.$$

Using that $l_{mL} + l_{mG} + l_{mH} = l_m$, we can express

$$l_{mH} = \frac{l_m}{\frac{\Psi_L}{\Psi_H}\frac{\theta_{mL}}{\theta_{mH}} + \frac{\Psi_G}{\Psi_H}\frac{\theta_{mG}}{\theta_{mH}} + 1}.$$

We can express sector-H price as a function of observables by plugging (22) and (23) into (4), and using that the θs sum to 1 within sector:

$$p_H = \left[\left(\frac{\alpha_{mH}}{w_m}\right)^{\eta-1} + \left(\frac{\alpha_{rH}}{w_r}\right)^{\eta-1} + \left(\frac{\alpha_{aH}}{w_a}\right)^{\eta-1}\right]^{\frac{1}{1-\eta}} = \frac{w_m}{\alpha_{mH}}\left(\frac{1}{\theta_{mH}}\right)^{\frac{1}{1-\eta}}.$$

Similarly using (2), (3) and relative αs within sectors as expressed above, as well as that within sectors the θs sum to 1, sectoral output can be expressed as:

$$Y_L = \left[(\alpha_{mL}l_{mL})^{\frac{\eta-1}{\eta}} + (\alpha_{rL}l_{rL})^{\frac{\eta-1}{\eta}} + (\alpha_{aL}l_{aL})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}$$

$$= \alpha_{mL}l_{mL}\left[1 + \left(\frac{\alpha_{rL}l_{rL}}{\alpha_{mL}l_{mL}}\right)^{\frac{\eta-1}{\eta}} + \left(\frac{\alpha_{aL}l_{aL}}{\alpha_{mL}l_{mL}}\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} = \alpha_{mL}l_{mL}\left(\frac{1}{\theta_{mL}}\right)^{\frac{\eta}{\eta-1}}$$

$$= \alpha_{mH}\frac{p_H}{p_L}\left(\frac{\theta_{mH}}{\theta_{mL}}\right)^{\frac{1}{1-\eta}}\left(\frac{1}{\theta_{mL}}\right)^{\frac{\eta}{\eta-1}}l_{mH}\frac{\Psi_L}{\Psi_H}\frac{\theta_{mL}}{\theta_{mH}} = \alpha_{mH}l_{mH}\theta_{mH}^{\frac{\eta}{1-\eta}}\frac{p_H}{p_L}\frac{\Psi_L}{\Psi_H},$$

$$Y_G = \alpha_{mH}l_{mH}\theta_{mH}^{\frac{\eta}{1-\eta}}\frac{p_H}{p_G}\frac{\Psi_G}{\Psi_H},$$

$$Y_H = \alpha_{mH}l_{mH}\theta_{mH}^{\frac{\eta}{1-\eta}}.$$

Using the above and the expressions for p_H and l_{mH} we can then write the value of output at current prices as:

$$p_L Y_L + p_G Y_G + p_H Y_H = w_m l_m \frac{\frac{\Psi_L}{\Psi_H} + \frac{\Psi_G}{\Psi_H} + 1}{\frac{\Psi_L}{\Psi_H}\theta_{mL} + \frac{\Psi_G}{\Psi_H}\theta_{mG} + \theta_{mH}}.$$

We can express the value of output at initial prices, where we denote by 0 the initial

period and we omit the subscript t in all other periods for brevity, as:

$$p_{L,0}Y_L + p_{G,0}Y_G + p_{H,0}Y_H$$

$$= \frac{\alpha_{mH}}{\alpha_{mH,0}} w_{m,0} \left(\frac{\theta_{mH}}{\theta_{mH,0}} \right)^{\frac{1}{1-\eta}} \frac{l_{mL}}{\frac{\Psi_L}{\Psi_H}\theta_{mL} + \frac{\Psi_G}{\Psi_H}\theta_{mG} + \theta_{mH}} \left(\frac{p_{L,0}}{p_{H,0}} \frac{\Psi_L}{\Psi_H} \frac{p_H}{p_L} + \frac{p_{G,0}}{p_{H,0}} \frac{\Psi_G}{\Psi_H} \frac{p_H}{p_G} + 1 \right).$$

The equivalent of output growth in our model is:

$$1 + \gamma = \frac{p_{L,0}Y_L + p_{G,0}Y_G + p_{H,0}Y_H}{p_{L,0}Y_{L,0} + p_{G,0}Y_{G,0} + p_{H,0}Y_{H,0}}.$$

The evolution of α_{mH} over time is therefore pinned down by:

$$\frac{\alpha_{mH}}{\alpha_{mH,0}} = \frac{(1+\gamma)}{\left(\frac{\theta_{mH}}{\theta_{mH,0}} \right)^{\frac{1}{1-\eta}} \frac{l_m}{l_{m,0}} \frac{\frac{\Psi_{L,0}}{\Psi_{H,0}}\theta_{mL,0} + \frac{\Psi_{G,0}}{\Psi_{H,0}}\theta_{mG,0} + \theta_{mH,0}}{\frac{\Psi_L}{\Psi_H}\theta_{mL} + \frac{\Psi_G}{\Psi_H}\theta_{mG} + \theta_{mH}} \frac{\frac{p_{L,0}}{p_{H,0}}\frac{\Psi_L}{\Psi_H}\frac{p_H}{p_L} + \frac{p_{G,0}}{p_{H,0}}\frac{\Psi_G}{\Psi_H}\frac{p_H}{p_G} + 1}{\frac{\Psi_{L,0}}{\Psi_{H,0}} + \frac{\Psi_{G,0}}{\Psi_{H,0}} + 1}}.$$

A.3 Calibration

Table 4 contains the targets used in the calibration.

Table 4: Calibration targets

	1960	1970	1980	1990	2000	2010
p_L/p_G	1	1.153	0.914	0.977	1.019	1.036
p_H/p_G	1	1.145	1.014	1.449	1.880	1.951
Ψ_L	0.215	0.234	0.234	0.250	0.261	0.252
Ψ_G	0.472	0.410	0.375	0.315	0.275	0.215
Ψ_H	0.312	0.356	0.392	0.435	0.465	0.533
growth	1	1.229	1.420	1.657	1.973	2.389
w_m/w_r	0.800	0.850	0.801	0.849	0.861	0.893
w_a/w_r	1.191	1.199	1.082	1.265	1.358	1.444
θ_{mL}	0.130	0.115	0.129	0.135	0.154	0.178
θ_{rL}	0.648	0.633	0.635	0.609	0.547	0.502
θ_{aL}	0.222	0.251	0.236	0.256	0.299	0.320
θ_{mG}	0.012	0.018	0.019	0.020	0.019	0.078
θ_{rG}	0.790	0.752	0.744	0.672	0.641	0.562
θ_{aG}	0.199	0.230	0.237	0.308	0.340	0.360
θ_{mH}	0.095	0.099	0.103	0.091	0.089	0.120
θ_{rH}	0.481	0.414	0.387	0.331	0.270	0.233
θ_{aH}	0.424	0.486	0.510	0.578	0.641	0.647

A.4 Robustness checks under alternative parameters

Here we explore the role of alternative values of ρ and ε. Note that these parameters do not change the productivity series we infer from the data, as the αs are fully pinned down by the production side of the model. However, these parameters impact the general equilibrium outcomes of the model. In each of these robustness checks we recalibrate the model based on the alternative values of the fixed parameters we consider.

We show in Figure 9 a sensitivity analysis with respect to ρ, the correlation in the cost differences of entering an occupation. Comparing the sub figures, for no correlation and $\rho = 0.6$ respectively, and Figure 3 in the main text, where $\rho = 0.4$, reveals that the model predictions are extremely robust with respect to this parameter; there is hardly any difference between the figures.

(a) Distance measure for $\rho = 0$ (b) Distance measure for $\rho = 0.6$
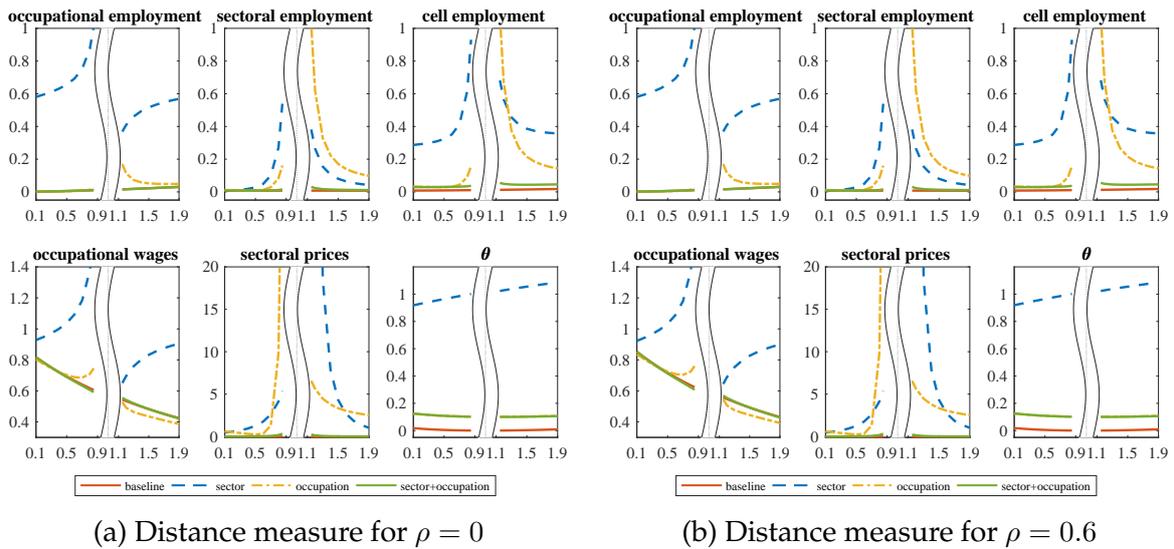
Figure 9: Sensitivity Analysis with respect to ρ

This figure plots the distance measure for the various outcomes of interest for alternative parameterizations of the correlation in the occupational choice costs. The baseline value in the main text is $\rho = 0.4$ for which these distance measures are shown in Figure 3.

Similarly Figure 10 shows the distance measures of model outcomes for alternative values of the elasticity of substitution in consumption. Comparing the various sub-panels where ε is set to 0.1, 0.3, 0.4 to each other and to Figure 3 which shows the baseline parametrization based on $\varepsilon = 0.2$ shows that the model predictions are very robust to this preference parameter.
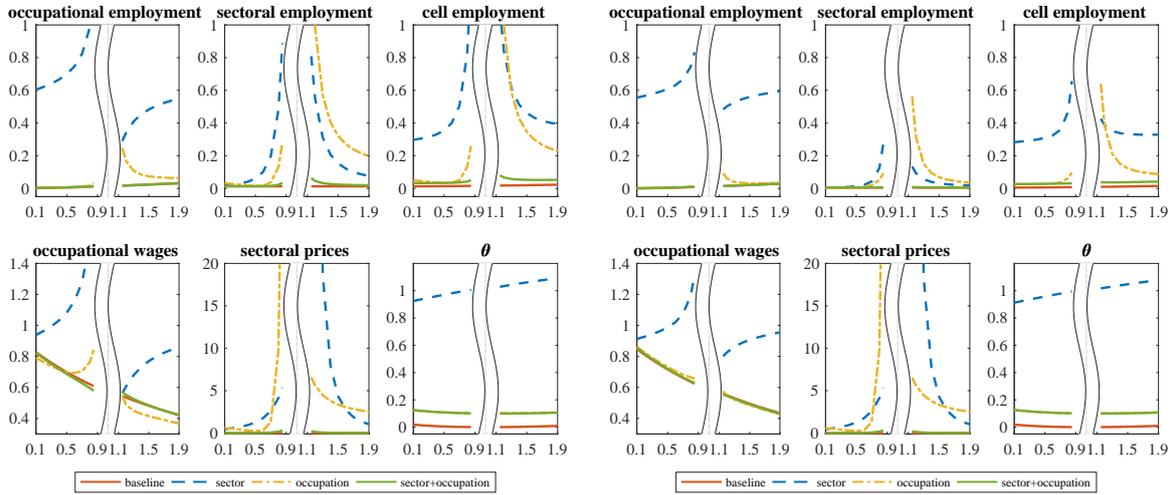
(a) Distance measure for $\varepsilon = 0.1$

(b) Distance measure for $\varepsilon = 0.3$
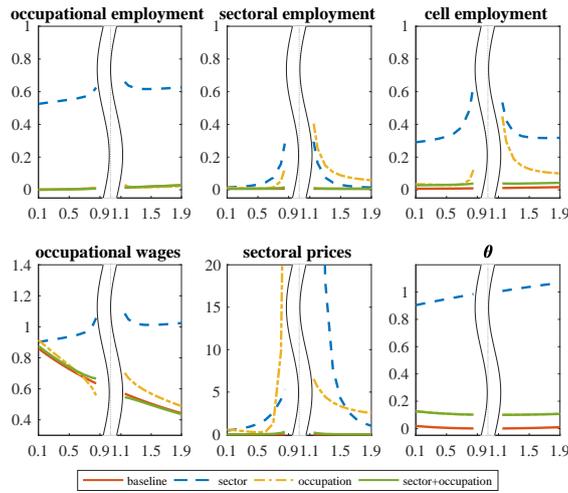
(c) Distance measure for $\varepsilon = 0.4$

Figure 10: Sensitivity Analysis with respect to ε

This figure plots the distance measure for the various outcomes of interest for alternative parametrizations of the consumption elasticity of substitution. The baseline value in the main text is $\varepsilon = 0.2$ for which these distance measures are shown in Figure 3.

A.5 Results based on the naive decomposition

We generate alternative cell productivity paths which are based on partial predictions. In particular the sector-only productivity series is generated as:

$$\widetilde{\ln \alpha}_{oJ,0}^{sec} = \ln \alpha_{oJ,0},$$

$$\widetilde{\ln \alpha}_{oJ,t}^{sec} = \widetilde{\ln \alpha}_{oJ,t-1}^{sec} + \widehat{\beta}_t + \widehat{\gamma_{J,t}} + \frac{\omega_{rJ,t}\widehat{\delta_{r,t}} + \omega_{aJ,t}\widehat{\delta_{a,t}}}{\omega_{mJ,t} + \omega_{rJ,t} + \omega_{aJ,t}},$$

and the occupation-only productivity series is generated as:

$$\widetilde{\ln \alpha}_{oJ,0}^{occ} = \ln \alpha_{oJ,0},$$

$$\widetilde{\ln \alpha}_{oJ,t}^{occ} = \widetilde{\ln \alpha}_{oJ,t-1}^{occ} + \widehat{\beta}_t + \widehat{\delta_{o,t}} + \frac{\omega_{oG,t}\widehat{\gamma_{G,t}} + \omega_{oH,t}\widehat{\gamma_{H,t}}}{\omega_{oL,t} + \omega_{oG,t} + \omega_{oH,t}}.$$

The difference to the main text is that we use the actual cell weights when assigning the average occupation(sector) effect to the sector-(occupation-)only series.

factors	full factor	sector	occupation	neutral
$\eta = 0.1$	0.806	0.396	0.567	0.047
$\eta = 0.2$	0.791	0.337	0.599	0.040
$\eta = 0.3$	0.776	0.276	0.629	0.033
$\eta = 0.4$	0.763	0.215	0.657	0.025
$\eta = 0.5$	0.752	0.156	0.681	0.018
$\eta = 0.6$	0.744	0.103	0.699	0.012
$\eta = 0.7$	0.739	0.059	0.709	0.007
$\eta = 0.8$	0.738	0.026	0.712	0.003
$\eta = 0.9$	0.741	0.007	0.705	0.001
$\eta = 1.1$	0.759	0.006	0.670	0.001
$\eta = 1.2$	0.771	0.023	0.644	0.003
$\eta = 1.3$	0.785	0.049	0.614	0.006
$\eta = 1.4$	0.800	0.080	0.583	0.010
$\eta = 1.5$	0.815	0.116	0.552	0.015
$\eta = 1.6$	0.829	0.154	0.521	0.020
$\eta = 1.7$	0.843	0.193	0.491	0.024
$\eta = 1.8$	0.856	0.231	0.464	0.029
$\eta = 1.9$	0.869	0.268	0.438	0.034

Table 5: R^2 of the alternative decomposition

Table 5 contains the R^2 for the alternative predicted sector-only and occupation-only productivity paths. These alternative productivity components are naive in the

sense that they resemble the series one would have inferred when a priori restricting the nature of productivity growth to be specific to sectors or to occupations only. For comparison column 1 contains the R^2 of the full factor model, and the last column contains the R^2 of the neutral productivity path, these are the same as the respective columns of Table 2. In terms of the individual components' relative importance, the ranking is the same as in the main text: the R^2 of the occupation-only component is always substantially larger than the R^2 of the sector-only component. However, the exact magnitudes are somewhat different.
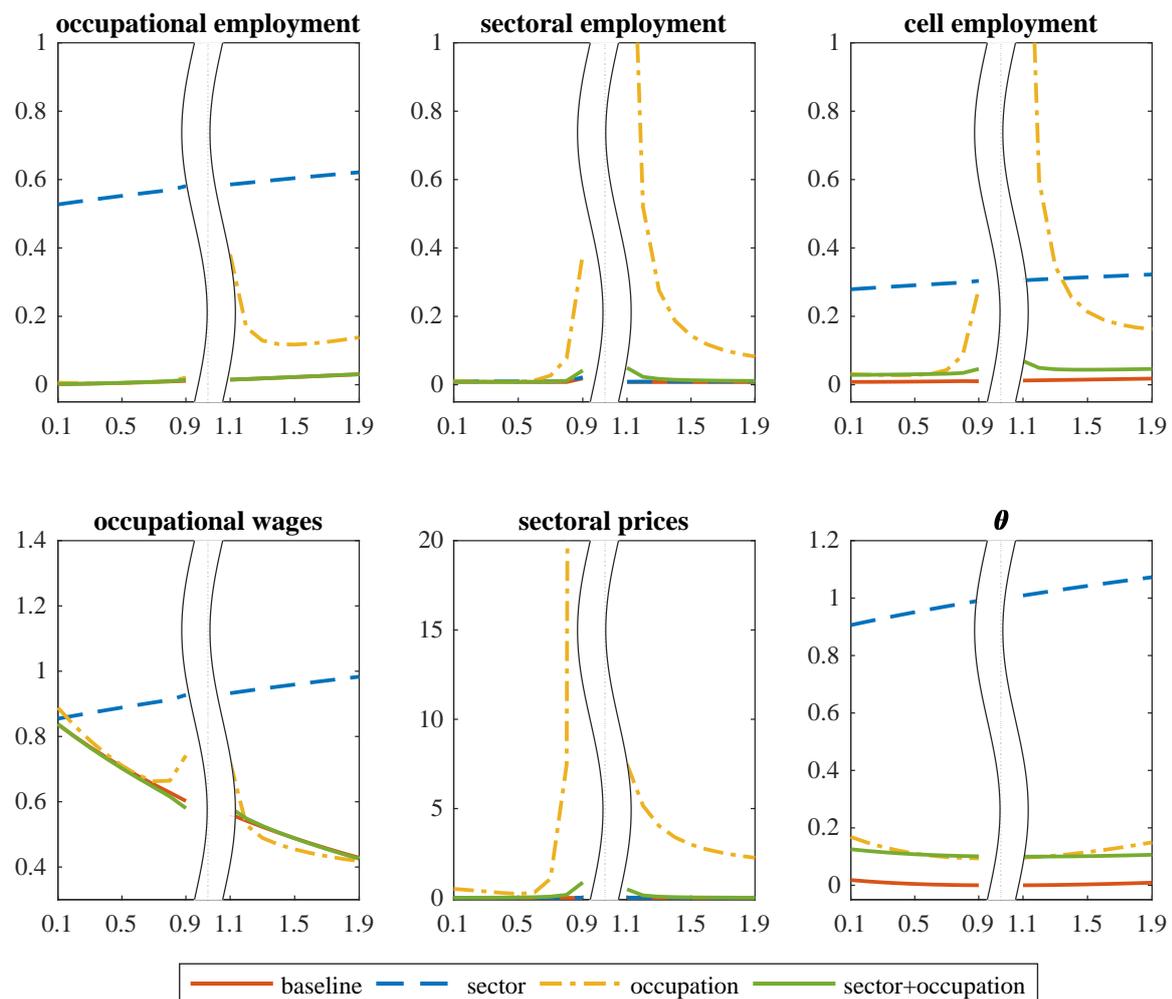

Figure 11: Alternative productivity components' ability to replicate the data

Figure 11 shows the distance measure between the data and the model using the above defined alternative productivity components. There are two things to note regarding the difference to Figure 3. The first is that here the sector-only model almost

perfectly replicates the data in terms of sectoral employment and relative sectoral prices. This is because in the naive model we load some of the occupational productivity growth differences on the sector-only component, as occupations are being used at different intensities in each sector. The second thing to note is that the performance of the occupation-only component is hardly affected by using this alternative prediction method.