

Disentangling Occupation- and Sector-specific Technological Change*

Zsófia L. Bárány[†]

Christian Siegel[‡]

July 4, 2017

Abstract

Occupational and sectoral labor market patterns display a significant overlap. This implies that economic models can explain these patterns to a large degree through either sector- or occupation-specific technological change, but stay silent about the level of specificity. We propose a model where technologies evolve at the sector-occupation level, allowing us to extract sector-only and occupation-only components and to quantify their importance. We find that most of productivity changes are occupation-specific, but that there is also a sizable sector component. We contrast the data and our baseline model against implications of models where technological change is restricted to be either at the sector or at the occupation level, or both. All three restricted models can replicate both sectoral and occupational outcomes very well, but occupation-specific changes are crucial for within-sector changes of occupational employment and income shares.

*We wish to thank Francesco Caselli, Georg Duernecker, Tim Lee, Guy Michaels, as well as seminar and conference participants at the 24th CEPR European Summer Symposium in International Macroeconomics in Helsinki, the VII Workshop on Institutions, Individual Behavior and Economic Outcomes in Alghero, Kent, the 2016 Workshop on Structural Transformation and Macroeconomic Dynamics, the Joint Vienna Macro workshop, and the Society for Economic Dynamics Annual Meeting 2017 for comments and suggestions. We gratefully acknowledge funding from a British Academy/Leverhulme Small Research Grant (SC-07334).

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

[‡]University of Kent. Email: c.siegel@kent.ac.uk

1 Introduction

Over recent decades the labor markets in most developed countries have experienced substantial changes. There has been structural change, a massive reallocation of labor across sectors, while at the occupational level labor markets experienced polarization; employment shifted out of middle-earning jobs to low- and high-earning jobs, which also saw higher wage growth than the middle-income occupations. Both of these patterns are typically explained through non-neutral productivity growth. Whereas the literature on structural change focuses on productivity growth differentials across sectors (e.g. Kongsamut, Rebelo, and Xie (2001), Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008)), the emphasis of the polarization literature is typically differences across occupations (e.g. Autor, Katz, and Kearney (2006), Goos and Manning (2007), Autor and Dorn (2013), Goos, Manning, and Salomons (2014), Michaels, Natraj, and Van Reenen (2014)).

In this paper we aim to disentangle to what degree technological change is neutral or specific to sectors or to occupations. Since in the data one does not observe the output produced by workers in individual occupations (but only of firms or sectors), one cannot directly measure productivities at the occupation level. We therefore build a flexible model that we use to extract productivities at the occupation-sector level from household survey data. We do not restrict the nature of technological change, but infer productivities for each occupation-sector cell, which we then decompose into neutral, sector, and occupation components. This allows us to quantify their contributions. In addition, we compare our results to a model that restricts productivity growth to be specific either to the sector or the occupation, or to both.

The phenomena of structural change and polarization across occupations have been connected in recent literature. In Bárány and Siegel (2017) 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. 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)

and Lee and Shin (2017) argue that unbalanced occupational productivity growth by itself provides dynamics consistent with structural change.

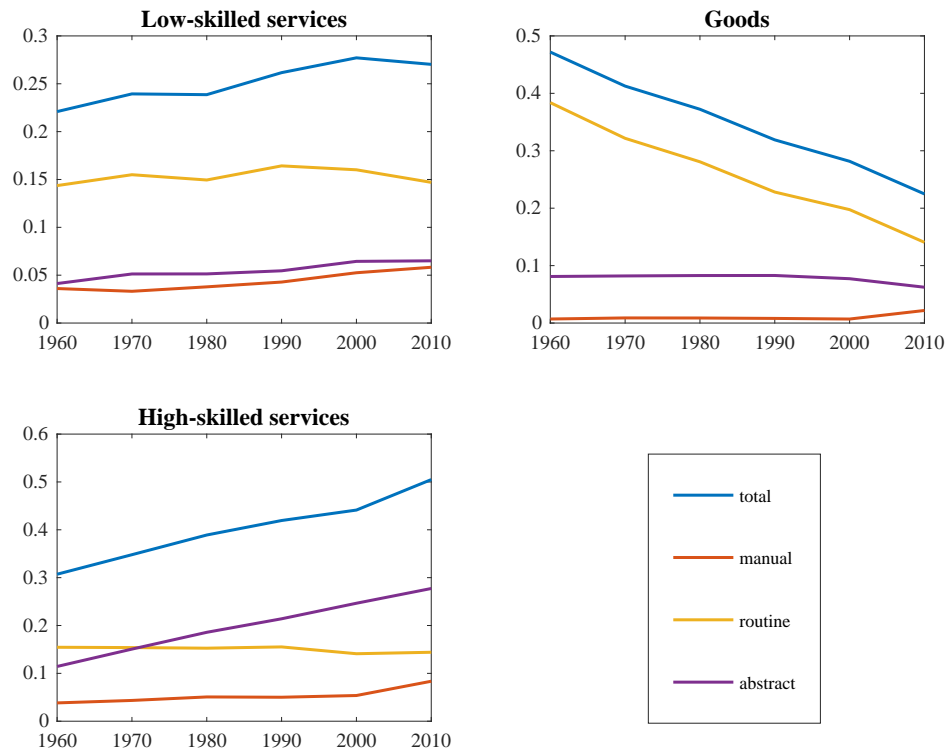


Figure 1: Sector-occupation employment shares 1960-2007

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 sector-occupation employment shares in the U.S. over 1960–2010. The black lines show the employment share of each sector, which within the sector panels is broken down by occupations. For the classification of occupations and industries see appendix A.1.

That the observed sectoral and occupational patterns in employment and wages can be explained by either productivity growth differentials between sectors or alternatively between occupations, is due to the large overlap between sectoral and occupational employment the data displays. For three broad sectors (low-skilled services, goods, high-skilled services) and three occupational categories (manual, routine, abstract), Figure 1 plots the evolution of sector-occupation employment shares in the U.S. between 1960–2010. The black lines show the employment share of each sector, which is then broken down into occupations. The structural change in the economy is apparent in the pronounced decline in goods sector employment and the rise in (particularly

high-skilled) service sector employment. Labor market polarization is manifested in the fall of the share of routine occupations (which traditionally are in the middle of the wage distribution). However, looking at occupations and sectors more carefully, two additional facts are apparent. First, the goods sector has the highest share of routine labor. Second, by far most of the decline in routine employment occurred in the goods sector, whereas in the two service sectors it declined only slightly. Similarly, almost all of the increase in the employment share of abstract occupations took place in the high-skilled service sector, and virtually all of the increase in manual employment up to 2000 was in low-skilled services. Because of this large overlap between the evolution of occupational and sectoral employment, models that allow for productivity growth differences only at the sectoral level or only at the occupational level can do a relatively decent job in generating patterns in line with the data. However, such restricted models load all technological differences on one type of factor, therefore not allowing to identify whether these technological differences are indeed at the sector or occupational level. Yet distinguishing the nature of productivity growth matters for policy implications. The goal of our paper is to disentangle and quantify different forms of technological change. Our setup is flexible enough to allow for productivity changes that are neutral (economy-wide), specific to firms in particular industries (producing particular products), specific to workers in certain occupations (linked to their task content), or specific to occupation-sector cells (a residual).

We build a parsimonious model featuring multiple sectors that employ various occupations, possibly at different intensities. In this model we do not impose any structure on the nature of technological change. Using only the production-side of the model, we extract from US IPUMS data the time series of sector-occupation specific technological change. Then, using a factor model we decompose the inferred productivities into a neutral component, sector-specific components, occupation specific components, and a residual. Our results suggest that technological change is biased both at the sector and at the occupation level. When considering the three broad sectors and three broad occupational categories used in Figure 1, we find that occupation factors account for 50-53 percent of productivity changes between 1960 and 2010, sector factors for for 20-34 percent, and that between 4 and 9 percent are due to neutral

technologies, whereas 24-34 percent are specific to occupation-sector cells.

Finally, we study models that a priori restrict the nature of technological progress. We explore variants where productivity growth is restricted to occur only at the sector level, only at the occupation level, or at both levels (but not their interaction), and contrast their predictions against our flexible model and the data. We find that these restricted models also do a reasonable well job in matching some aspects of the data. However, while all of these restricted models can capture the general trends in relative wages and employment, both at the sectoral as well as at the occupational level, we find that occupation-specific productivity growth is crucial to replicate the observed within-sector changes in occupational employment and income shares. We therefore conclude that, despite our evidence for both sizable occupation and sector factors in technologies, studies that only allow for one of two are not misleading with respect to their predictions for sectoral or occupational outcomes. However, if also the occupational employment or income shares within sectors are objects of interest, it is necessary to model occupation-specific technological change. We further find that while the value of the elasticity of substitution is quantitatively (and for sectoral prices even qualitatively) important for the occupation-only models, it is essentially irrelevant for the predictions of sector-only and sector-occupation models. As there is no consensus in the literature on the value of the elasticity of substitution between different occupations, our analysis suggests that models should allow for technologies to evolve at the sector- and at the occupation-level for robust results.

2 Model

We assume that there is a continuum of measure one of heterogeneous workers in the economy. Each worker optimally selects his occupation, and can freely choose which sector of the economy to supply his 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. Taking as given the supply of the different types of workers over time, we use this parsimonious static model to pin down how the valuation 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, goods, and high-skilled services. 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:

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

where l_{oJ} is occupation o labor used in sector J , $\alpha_{oJ} > 0$ is an *occupation-sector* specific labor augmenting technology term for occupation $o \in \{m, r, a\}$ in sector J , and $\eta \in [0, \infty]$ is the elasticity of substitution between the different types of labor.¹ 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 occupation-sector 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.² We use the model to calculate from the data the

¹We 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.

²Given that the data shows a large overlap between sectors and occupations (e.g. the share of routine workers is highest in the goods sector), 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 that are actually due to the other factor. Our approach circumvents this problem as we do not impose any a priori restrictions.

occupation-sector 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 down the optimal relative labor use as:

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

$$\frac{l_{aJ}}{l_{rJ}} = \left(\frac{w_r}{w_a} \right)^\eta \left(\frac{\alpha_{aJ}}{\alpha_{rJ}} \right)^{\eta-1}. \quad (2)$$

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:

$$p_J = \left[\alpha_{mJ}^{\eta-1} \frac{1}{w_m^{\eta-1}} + \alpha_{rJ}^{\eta-1} \frac{1}{w_r^{\eta-1}} + \alpha_{aJ}^{\eta-1} \frac{1}{w_a^{\eta-1}} \right]^{\frac{1}{1-\eta}}. \quad (3)$$

Finally using (1), (2) and (3) to express sector J output, optimal sectoral labor use can be expressed as:

$$l_{mJ} = \left[\frac{p_J \alpha_{mJ}}{w_m} \right]^\eta \frac{Y_J}{\alpha_{mJ}}, \quad (4)$$

$$l_{rJ} = \left[\frac{p_J \alpha_{rJ}}{w_r} \right]^\eta \frac{Y_J}{\alpha_{rJ}}, \quad (5)$$

$$l_{aJ} = \left[\frac{p_J \alpha_{aJ}}{w_a} \right]^\eta \frac{Y_J}{\alpha_{aJ}}. \quad (6)$$

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 (L), goods (G), or high-skilled services (H), implying that in equi-

librium, 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 provides them with the highest income. Thus an individual i chooses sector j if

$$w_j - \chi_j^i \geq w_k - \chi_k^i \quad \text{for any } k \neq j, \quad k, j \in \{m, r, a\},$$

where w_j is the unit wage in occupation j and χ_j^i is individual i 's cost of entering occupation j . The optimal occupational choice is summarized in Figure 2.

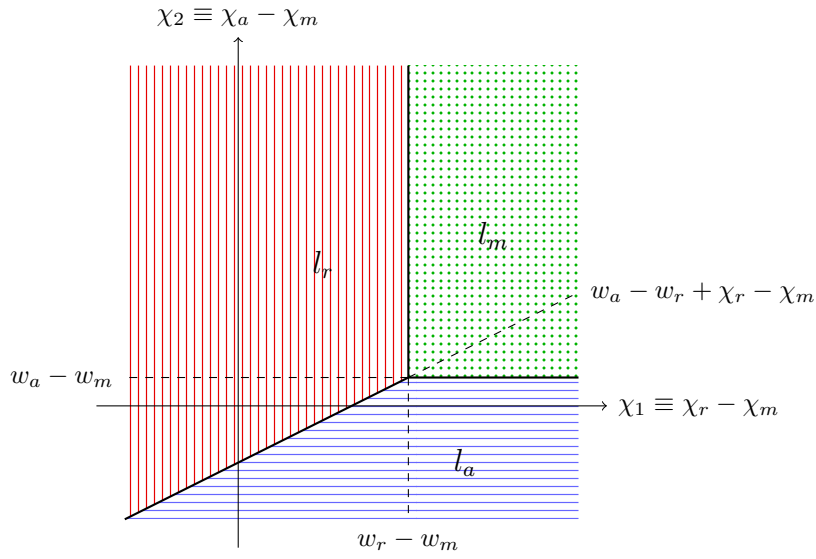


Figure 2: Optimal occupational choice

Notes: The graph shows the optimal selection of individuals into manual, routine and abstract occupations in terms of their idiosyncratic occupational cost differences ($\chi_1 \equiv \chi_r - \chi_m$ and $\chi_2 \equiv \chi_a - \chi_m$), 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:

$$l_m = \int_{w_r - w_m}^{\infty} \int_{w_a - w_m}^{\infty} f(\chi_1, \chi_2) d\chi_1 d\chi_2 \quad (7)$$

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

$$l_a = \int_0^{\infty} \int_{-\infty}^{\min\{w_a - w_r + \chi_1, w_a - w_m\}} f(\chi_1, \chi_2) d\chi_1 d\chi_2 \quad (9)$$

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:

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

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

where $a_L + a_G + a_H = 1$, $\varepsilon < 1$, i. e. the goods and services are complements in consumption, and c_L, c_H allow for non-homotheticity in consumption. 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 every type of good and service (i. e. an internal solution), optimality implies the following demand schedule:

$$C_L = \left(\frac{a_L}{p_L} \right)^{\varepsilon} \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^{\varepsilon} p_L^{1-\varepsilon} + a_G^{\varepsilon} p_G^{1-\varepsilon} + a_H^{\varepsilon} p_H^{1-\varepsilon}} - \bar{c}_L, \quad (10)$$

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

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

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: $w_r = 1$. 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 (3) and (10), (11), and (12). Then from (4), (5), or (6) 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:³

$$\begin{aligned} l_{mL}(w_m, w_a) + l_{mG}(w_m, w_a) + l_{mH}(w_m, w_a) &= l_m, \\ l_{rL}(w_m, w_a) + l_{rG}(w_m, w_a) + l_{rH}(w_m, w_a) &= l_r. \end{aligned}$$

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 micro data using only the production side. We therefore proceed in the following steps, similarly to Buera, Kaboski, and Rogerson (2015). First, we calibrate the objects of the sectoral production functions taking as given the occupational wage rates and employment shares, and 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 ACS data between 1960 and 2010 to calculate occupational wage rates and occupational labor income shares within sectors, from which we infer the cell-specific productivities. In addition, we compute each sector's share in labor

³Due to Walras' law the market for abstract labor automatically clears.

income⁴, which we use as targets for the calibration of the utility function. For these calculations, we categorize workers into our three occupations and our three sectors based on their occupational code (*occ1990*) and respectively on their industry code (*ind1990*).

We calculate manual, routine, and abstract wage rates as the average hourly wage of a narrowly defined group – 25 to 29 year old men – in the given occupation. We rely on this measure – rather than on the average hourly wage of all workers within an occupation – to limit the potential influence of composition changes, e.g. due to differential changes in the demographic composition of workers across occupations.⁵ The occupational wage rate targets are calculated as:

$$\begin{aligned}\omega_m &\equiv \frac{\text{average hourly wage of 25–29 year old men in manual jobs}}{\text{average hourly wage of 25–29 year old men in routine jobs}}, \\ \omega_r &\equiv \frac{\text{average hourly wage of 25–29 year old men in routine jobs}}{\text{average hourly wage of 25–29 year old men in routine jobs}} = 1, \\ \omega_a &\equiv \frac{\text{average hourly wage of 25–29 year old men in abstract jobs}}{\text{average hourly wage of 25–29 year old men in routine jobs}}.\end{aligned}$$

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 :

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

Finally, we calculate sectoral income shares as

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

Note that, we can express total occupational earnings in two ways. Either simply as total occupational labor supply times occupational wage, or as the sum of the given

⁴In our model the sectoral labor income shares are equal to value added shares as there are no other production factors.

⁵This 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.

occupation's earnings across different sectors. This identity allows us to calculate the occupational labor supplies, l_m , l_r and l_a , from the data on sectoral income shares, occupational labor income shares within sectors, and occupational wages.⁶

We use data from the Bureau of Economic Analysis (BEA) between 1960 and 2010 to get sectoral prices and the growth rate of GDP per worker between periods.⁷

3.2 Calibration of the production side

As mentioned before, given the structure of the model we can infer the productivity parameters directly from the micro 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 set this elasticity to $\eta = 0.6$, which is close to the value Duernecker and Herrendorf (2016) use for the elasticity between goods and service occupations. Since in our model η captures the elasticity between three occupational categories, this might not be the right value. We conduct sensitivity analysis with respect to the value of η . For our baseline model, its value hardly matters. For the factor model its value has a small impact on the quantitative importance of the various factors. Finally, its value turns out to be crucial for some of the outcomes of the restricted models, which we explore in a sensitivity analysis in section 5.3.

We calculate the nine cell-specific productivity parameters, the α_s , in each period. As discussed earlier, 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. The intuition of what pins down the different α_s is the following. The occupational income shares pin down the ratios of α_s within sectors in each period, given occupational wages (from (1) and (2)).

⁶Given the measurement of unit wages, these labor supply shares are not the same as the share of hours across occupations, but the trends are very similar. See the appendix for details of the calculations of occupational labor supplies.

⁷Table 3 in the appendix contains all the targets we use in the calibration.

The sectoral relative prices pin down the α s across sectors within each period, again given occupational wages (from (3)). Finally, the overall growth rate of the output per worker pins down the evolution of the α s over time, given the distribution of income across sectors and occupational labor supplies.⁸

3.3 Calibration of the cost distribution and of the consumption side

To close the model we need to parameterize the household side of the model. These choices matter only for model simulations but not for assessing the contributions of sector and occupation factors to productivity growth.

To calibrate the distribution of cost differences, we assume that $f(\chi_1, \chi_2)$ is distributed according to a bivariate normal, and we fix the correlation parameter to be 0.4. Given this correlation, we calibrate the two means and diagonal 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.

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 c_L , c_H , a_L , and a_G 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.

Table 1 contains all the time-invariant parameters of the model. These, together with the evolution of the α s as backed out from the data using the model fully specify the calibrated model.

It is important to note that the occupational wages and the sectoral income shares

⁸See the appendix for the details of these steps.

Table 1: Calibrated parameters

	Description	Value
η	elasticity of substitution in production	0.6
ε	elasticity of substitution in consumption	0.2
ρ	correlation of cost differences	0.4
\bar{c}_L	non-homotheticity term in L	0.0036
\bar{c}_H	non-homotheticity term in H	0.0058
a_L	weight on L	0.092
a_H	weight on H	0.908
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)

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 6b and 7.

4 Factor model decomposition

We set up a factor model to relate 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

$$\Delta \ln \alpha_{Jo,t} = \beta_t + \gamma_{J,t} + \delta_{o,t} + \varepsilon_{Jot}, \quad (13)$$

In the regression the residual ε_{Jot} captures productivity changes that are idiosyncratic to that sector-occupation cell in period t . The sector dummy that is omitted from the regression is the one for the goods sector, the omitted occupation dummy is the manual one, and the omitted time dummy is for 1970. 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 in 1970.

To gauge the quantitative importance of each factor we generate four different cell productivity series. First, we generate the predicted productivity series using all factors as:

$$\begin{aligned}\widehat{\ln \alpha_{oJ,0}} &= \ln \alpha_{oJ,0} \\ \widehat{\ln \alpha_{oJ,t}} &= \widehat{\ln \alpha_{oJ,t-1}} + \widehat{\beta}_t + \widehat{\gamma}_{J,t} + \widehat{\delta}_{o,t}.\end{aligned}\quad (14)$$

To construct a productivity series that shuts down occupation-specific changes we therefore generate the following 'sector' series as:

$$\begin{aligned}\widehat{\ln \alpha_{oJ,0}}^{sec} &= \ln \alpha_{oJ,0} \\ \widehat{\ln \alpha_{oJ,t}}^{sec} &= \widehat{\ln \alpha_{oJ,t-1}}^{sec} + \widehat{\beta}_t + \widehat{\gamma}_{J,t} + \frac{\widehat{\delta}_{m,t} + \widehat{\delta}_{r,t} + \widehat{\delta}_{a,t}}{3},\end{aligned}\quad (15)$$

where the last term assigns the average occupation effect of that year.

Similarly to shut down sector-specific changes we generate the following 'occupation' series as:

$$\begin{aligned}\widehat{\ln \alpha_{oJ,0}}^{occ} &= \ln \alpha_{oJ,0} \\ \widehat{\ln \alpha_{oJ,t}}^{occ} &= \widehat{\ln \alpha_{oJ,t-1}}^{occ} + \widehat{\beta}_t + \widehat{\delta}_{o,t} + \frac{\widehat{\gamma}_{L,t} + \widehat{\gamma}_{G,t} + \widehat{\gamma}_{H,t}}{3}\end{aligned}\quad (16)$$

which gives a prediction for productivities based on (time-varying) neutral and occupation-specific factors only.

Finally, we generate a time-only productivity change series as:

$$\begin{aligned}\widehat{\ln \alpha_{oJ,0}}^{time} &= \ln \alpha_{oJ,0} \\ \widehat{\ln \alpha_{oJ,t}}^{time} &= \widehat{\ln \alpha_{oJ,t-1}}^{time} + \widehat{\beta}_t + \frac{\widehat{\delta}_{m,t} + \widehat{\delta}_{r,t} + \widehat{\delta}_{a,t}}{3} + \frac{\widehat{\gamma}_{L,t} + \widehat{\gamma}_{G,t} + \widehat{\gamma}_{H,t}}{3}\end{aligned}\quad (17)$$

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 .⁹

⁹The α s themselves change as we back out cell productivities conditional on a value of the elasticity,

factors	time & sector & occupation	time	time & sector	time & occ
$\eta = 0.4$	0.6627	0.0457	0.2036	0.5048
$\eta = 0.6$	0.6673	0.0549	0.2012	0.5210
$\eta = 0.8$	0.6779	0.0639	0.2115	0.5304
$\eta = 1.2$	0.7125	0.0796	0.2623	0.5298
$\eta = 1.4$	0.7335	0.0858	0.2975	0.5218
$\eta = 1.6$	0.7552	0.0908	0.3357	0.5104

Table 2: R^2 of factor model decomposition

It shows the R^2 of the factor model regression (13) as well as when only the time factor, or time and sector or time and occupation are used in the predictions. Time-varying sector and occupation effects explain about 66 and 76% of all cell productivity changes. This implies that between a third and a fourth of the variation is due to effects idiosyncratic to the sector-occupation cell. The time fixed effect by itself, which captures neutral technological progress, explains only 4-9% of productivity changes, whereas time-varying sector effects roughly 20-34% and occupation effects between 50 and 53%. Our interpretation of this decomposition is that most of productivity changes are not neutral, but biased across sectors and occupations. The contribution of occupation-specific changes to the evolution of cell productivities is between 1.7 and 2.5 times the one of sector-specific changes. To study what the different kinds of specificity in technological progress imply for various economic outcomes, we conduct in the next section a series of counterfactual experiments.

5 Quantitative experiments

In this section we first calibrate three alternative models, one where technological progress is restricted to be sector-specific and one where it is restricted to be occupation-specific. We then show that our baseline model does very well in matching the data and compare its results against from the restricted models.

but it is important to bear in mind that this series is independent of any other part of the model. A different η will imply a different parametrization of the household side to match the 1960 and 2010 data, but that part of the model does not affect the analysis of cell productivities.

5.1 Restricted models

Our model is flexible enough to allow productivity growth to be neutral or specific to sectors, occupations, or both. Most existing literature, however, relies on frameworks that a priori restrict the nature of technological progress to be specific to sectors or to occupations. To explore how these types of models compare to our richer and more flexible setup, in this section we restrict technological progress to be either sector, or occupation specific, or a combination of the two. In particular we assume that the productivity process is described by one of the following processes:

$$\tilde{\alpha}_{oJ,t}^{sec} = Z_{J,t}\alpha_{oJ,0}, \quad (18)$$

$$\tilde{\alpha}_{oJ,t}^{occ} = Z_{o,t}\alpha_{oJ,0}, \quad (19)$$

$$\tilde{\alpha}_{oJ,t} = Z_{J,t}Z_{o,t}\alpha_{oJ,0}, \quad (20)$$

Note that we assume that the growth rate of technology is either specific to the sector, $Z_{J,t}$ as in (18) or to the occupation, $Z_{o,t}$ as in (19), or there is both a sector and an occupation component as in (20), but the initial level, $\alpha_{oJ,0}$ is sector and occupation specific. To evaluate the model's predictions under these alternative assumptions we re-calibrate the productivity growth terms. We take from our baseline calibration all the time-invariant parameters: the elasticity of substitution in production, and the parameters of the cost distribution and of preferences. Given these, we calibrate the growth rates (either 3 sector, or 3 occupation specific, or 3 of each) in each period after 1960 to match the same nine targets (six occupational income shares within sector, two relative sectoral prices, and overall GDP growth) as in the calibration of the production side of our baseline model in general equilibrium. Numerically we solve in each period for the three (six) productivity growth rates that minimize the (equally-weighted) sum of percentage deviations between the model-implied statistics and these targets. Thus in general the restricted model does not perfectly match the data except in the initial period. Figure 3 shows the path of the baseline model, where the productivities are extracted from the data (red solid line), as well as the various restricted models: sector-only (blue dashed), occupation-only (yellow dashed-dotted), and sector- and

occupation-specific (green solid line).

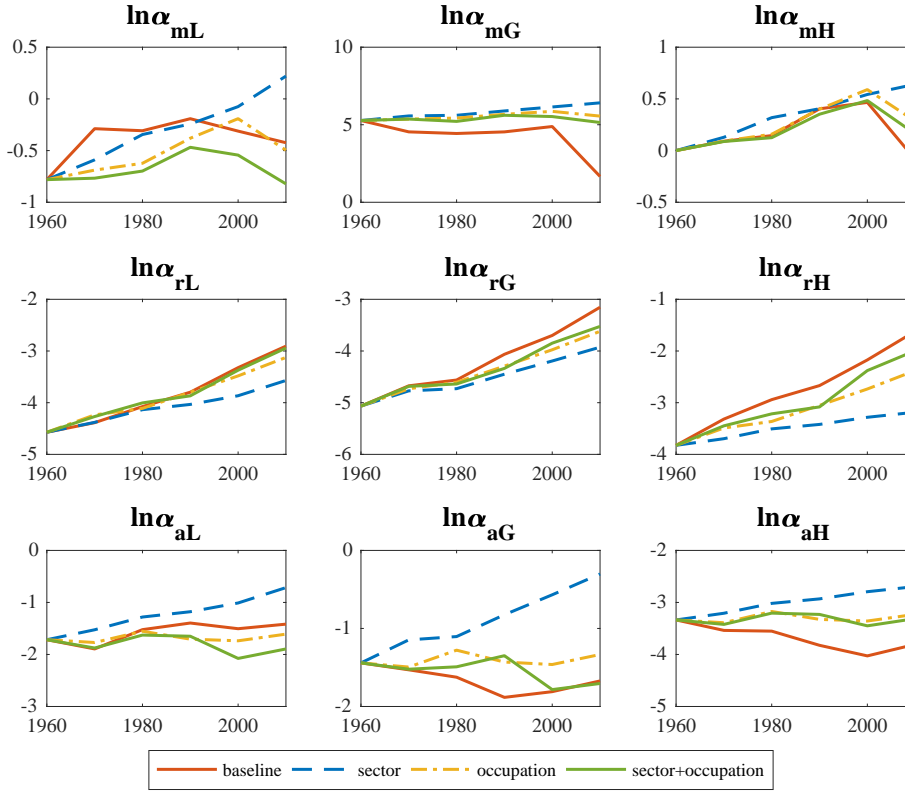


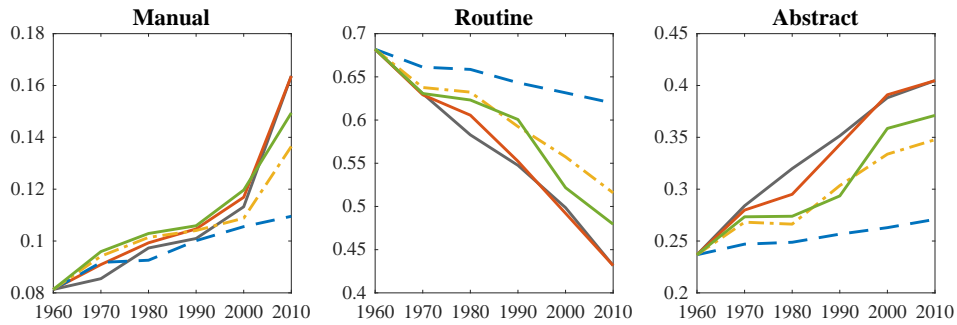
Figure 3: Calibrated productivity growth: baseline and restricted models

This figure shows in the solid red line the natural logarithm of the cell productivities as calculated from the data using the economic model. The columns refer to the three sectors (L , G , H) and the rows to the three occupations (m , r , a). The other lines show the calibrated productivity paths, when restricting productivity growth to be sector-specific (as shown in (18), blue dashed line), occupation-specific (as in (19), yellow dashed-dotted), and sector- and occupation-specific (as in (20), green solid line).

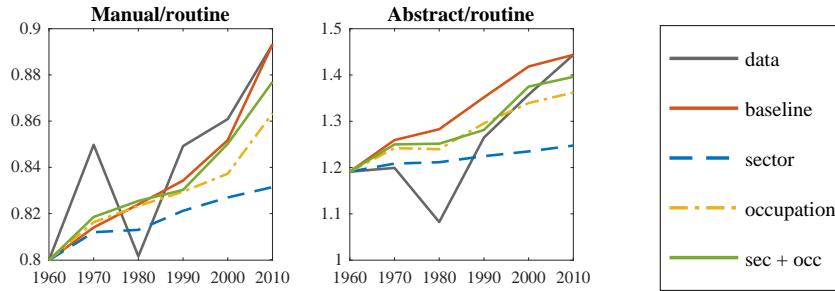
5.2 Model vs data: baseline and restricted models

Figures 4–7 show the baseline model economy’s evolution given the calibrated cell productivities (solid red line) contrasted with the data (solid grey line). The other lines show the predictions feeding in the cell productivities as calibrated in the restricted models. In particular, the dashed blue lines are generated using the sector-only growth ($\tilde{\alpha}^{sec}$), the dashed-dotted yellow lines the occupation-only growth ($\tilde{\alpha}^{occ}$), whereas the solid green lines are based on both sector and occupation specific growth rates ($\tilde{\alpha}$).

The first thing to note in Figures 4, 5 and 6 is that our baseline model does very well in matching sectoral and occupational employment shares and relative wages, as well as sectoral income shares and relative prices. This is partially by construction:



(a) Occupational employment



(b) Occupational relative wages

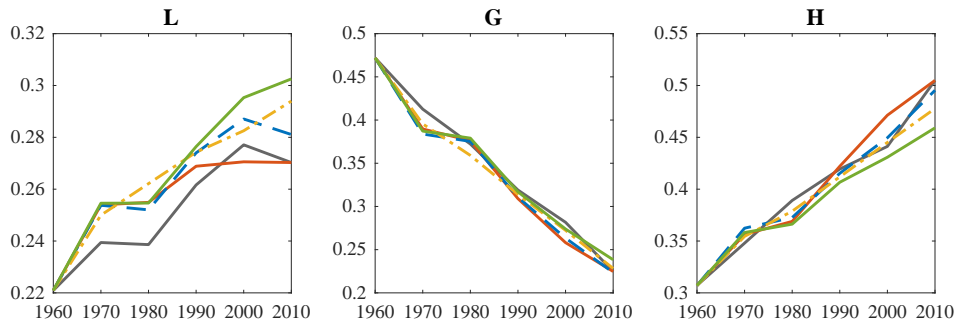
Figure 4: The evolution of occupational labor outcomes

This figure plots the path of the occupational employment shares (on the top) and relative wages (on the bottom) in the data, and for various paths of cell productivities: from the baseline and the restricted models.

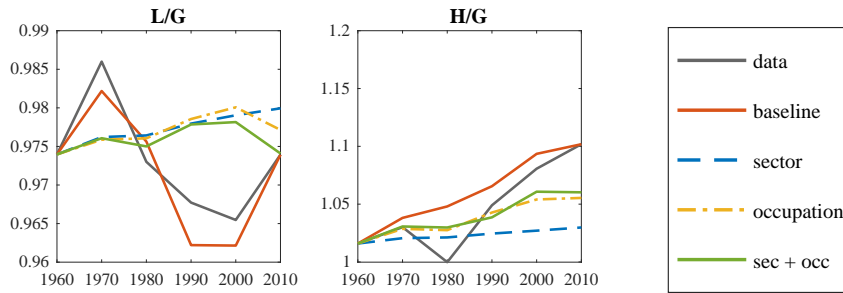
we calibrated the model such that in the initial and final period this data is exactly matched. For the periods in between this is not the case, as in the calibration of the utility function we did not target these periods. The gray (of the baseline model) and red (data) lines meet in 1960 and 2010 by construction, but they are also quite close in the periods in between. The baseline model also accounts very well for the evolution occupational income shares within sectors (Figure 7), which were used in identifying the cell-specific productivities.¹⁰

Second, these figures establish that models that restrict the nature of technological progress replicate by large most features of the data. In particular, the variant with productivity growing at the sector and at the occupation level (green solid line) gets for these outcomes close to our flexible model (red line) and the data (gray line). From this we infer that sector-specific and occupation-specific productivity changes are suf-

¹⁰The model matches the θ s perfectly in 1960 and in 2010 when it also matches occupational wage rates perfectly. In the interim periods, due to occupational wages in the model slightly deviating from the data, there is a tiny discrepancy between the model-implied and the observed occupational income shares within sectors.



(a) Sectoral employment



(b) Sectoral relative wages

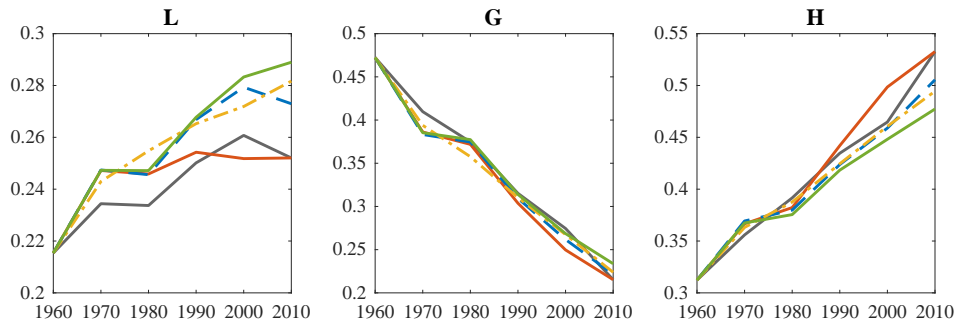
Figure 5: The evolution of sectoral labor outcomes (i)

This figure plots the path of the sectoral employment shares (on the top) and relative wages (on the bottom) in the data, and for various paths of cell productivities: from the baseline and the restricted models.

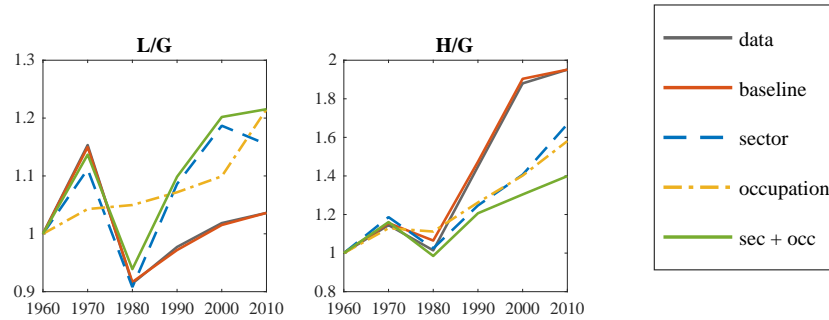
efficient to give a good description of the labor market facts we focus on. The models that allow only for sector-level (blue dashed line) or only occupation-level (yellow dashed-dotted line) technological change have predictions that are qualitatively consistent with the data along these dimensions.

Figure 4 shows that the decline in routine employment and relative wage compared to manual and abstract workers is well accounted for in our baseline model, but at least qualitatively also in the restricted models. While in the version with occupation-only technological change, for example routine workers are directly adversely affected by a change in productivities, the sector-only model also generates this due to a decline in the goods sector, which is the sector that uses routine workers the most intensely. Quantitatively, at least under our calibration, this latter effect by itself is not strong enough to generate changes as large as seen in the data.

The conclusions for sectoral employment and relative wages in Figure 5 and sectoral income shares and relative prices in Figure 6 are very similar. It is important



(a) Sectoral income shares



(b) Sectoral relative prices

Figure 6: The evolution of sectoral outcomes (ii)

This figure plots the path of the sectoral income shares (on the top) and relative sectoral prices (on the bottom) in the data, and for various paths of cell productivities, as discussed in the text.

to note that for relative sectoral prices this crucially depends on the elasticity of substitution being smaller than 1, we return to this issue in Section 5.3. The predictions of models with sector-only, occupation-only, or both sector- and occupation-level productivity changes are qualitatively in line and quantitatively close to the data.

To summarize, in terms of employment and relative wages at the sectoral and at the occupational level, as well as sectoral income shares and relative sectoral prices all restricted models imply changes that are qualitatively in line with the data (Figures 4, 5 and 6).

However, Figure 7 shows that the sector-only and the occupation-only technological change models have some distinct implications, in particular for within-sector occupational income shares. As already mentioned, the baseline model does very well in generating patterns in line with the data, as these are used as calibration targets for the productivities and when –due to the parameterization of the utility function– occupational wage rates are perfectly matched in 1960 and 2010, then also are these

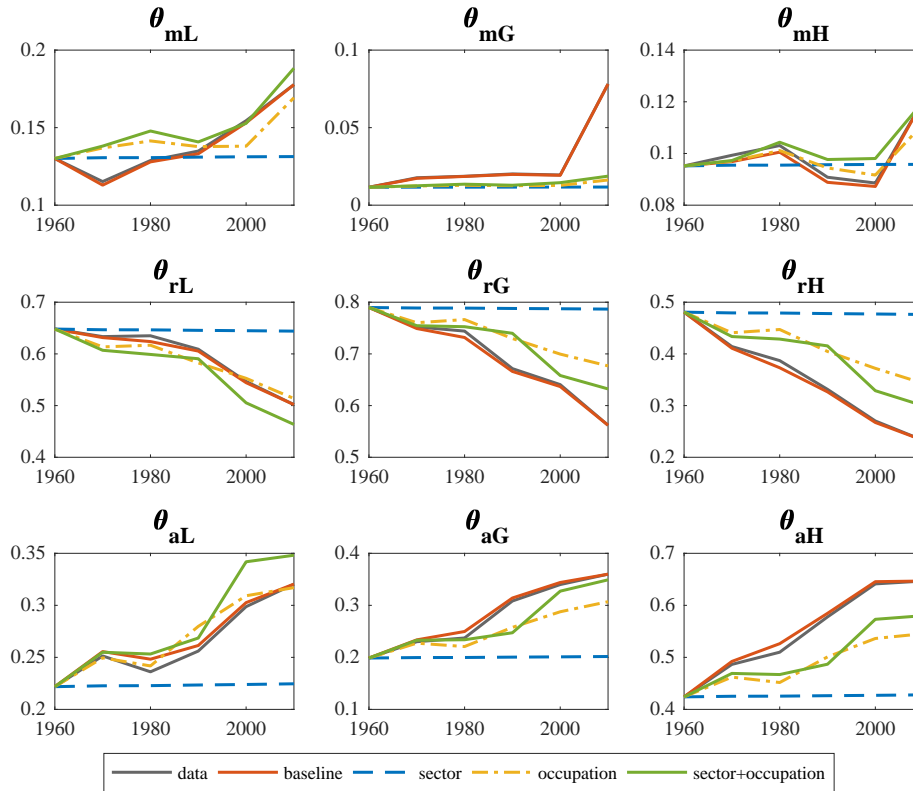


Figure 7: Income shares within sectors

This figure plots the path of the the share of occupational incomes within a sector in the data, and for various paths of cell productivities, as discussed in the text.

shares. The restricted model based on occupation-only technologies and on sector- and occupation-technology both do well in replicating these. However, the restricted model based on sector-only technological change has practically flat predictions; for example it fails to generate a drop in the routine labor income share in all sectors. As we show in Figure 11 in the Appendix, this version of the restricted model also fails to predict the observed fall in routine employment relative to other occupations within the low- and the high-skilled service sectors. The reason for this lies in the fact that routine labor is most intensively used in the goods sector. Forces that lead to structural change, i. e. employment movement out of the goods sector, based on any form of technological progress, have to increase wages in abstract and manual occupations relative to routine, which is in line with the data (see Figure 4b). Unless however the productivity of routine workers changed relative to other occupations, firms in each sector would optimally employ relatively more routine workers, as indicated by equation (5), which is at odds with the data. To explain the within-sector changes in oc-

cupational employment and income shares occupation-specific technological changes are therefore of crucial importance.

However, the restricted model based on sector-only technology has counterfactual predictions; for example it fails to generate a drop in the routine labor income share in all sectors. As we show in Figure 11 in the Appendix, this version of the restricted model also fails to predict the observed fall in routine employment relative to other occupations within each sector. The reason for this lies in the fact that routine labor is most intensively used in the goods sector. Forces that lead to structural change, i. e. employment movement out of the goods sector, based on any form of technological progress, have to increase wages in abstract and manual occupations relative to routine, which is in line with the data (see Figure 4b).

Overall this suggest that to understand the full picture of sector-occupation employment and wages as well as of sectoral prices, one needs to allow for technologies to evolve at least at both the sector- and the occupation-level. However, if one is only interested in wage or employment outcomes at the sector or the occupation level, restricting technologies to only change at the sector or only at the occupation-level is sufficient. Yet none of the restricted models does as well as our parsimonious yet flexible baseline model which does not impose any restrictions on the form of technological progress. One drawback of our method is that we rely on a parameterization of the model's production side to back out cell productivities. In particular, the elasticity of substitution between different occupations might be a crucial parameter for the restricted models.¹¹ To see how our results depend on parameter values we conduct a sensitivity analysis in the next section.

5.3 Sensitivity Analysis

In this section we study how the results of the restricted models depend on the value of the elasticity of substitution between different occupations, η . In sum, we find that the predictions stemming from sector-only or sector-and-occupation-only technological change models hardly depend on η , while the occupation-only model's predictions

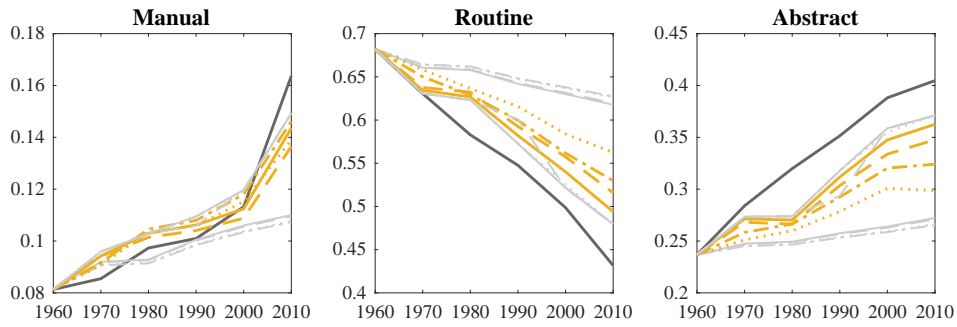
¹¹It is important to bear in mind that we back out cell productivities conditional on this elasticity. When changing η , the implied series of α s changes as well.

vary quantitatively much more with this elasticity, and for relative prices even qualitatively. Figure 8 shows the restricted models' predictions for occupational labor market outcomes for different values of η . As already mentioned, the models restricted to sector-only or to sector-and-occupation-only technology have predictions for occupational employment shares and relative wages that are very robust to different elasticities. As such, we show these predictions only in the light gray lines, which for the various values are very close to each other and to the baseline of 0.6 for η , i.e. the plots in Figure 4 which we replicate here in the dashed lines. The predictions of models with occupation-only technological change, however, depend much more on this elasticity. These predictions are shown in the yellow lines, where the solid line corresponds to a value of 0.4 for η , dashed to 0.6 (our baseline), dash-dotted to 1.4, whereas the dotted line is based on $\eta = 1.6$. While quantitatively the predictions based on occupation-only productivity growth depend on η , they all are qualitatively aligned with the data (dark gray line).

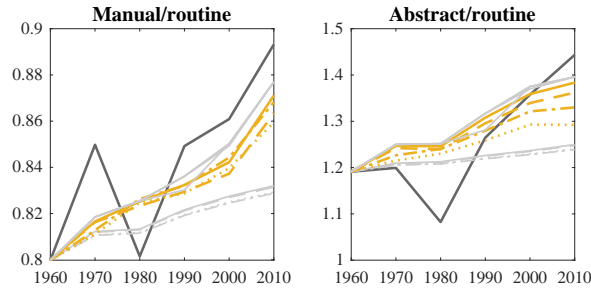
Similarly, the elasticity of substitution in production hardly matters for predictions from models featuring sector-level technologies, but it matters quantitatively for models restricted to occupation-only technologies, as can be seen in Appendix Figure 12. Yet qualitatively, irrespective of the parameterization of this elasticity, all restricted models have predictions for the evolution of employment shares and relative wages, at the sector and at the occupation level, as well as sectoral income shares, that are consistent with the data.

A general pattern emerges from these figures: the closer is the elasticity of substitution to 1, i. e. as the production function gets closer to the Cobb-Douglas case, the predictions of the occupation-only models get closer to those of the sector-only models. This is very intuitive, as for Cobb-Douglas production functions factor-augmenting and neutral technologies are indistinguishable, implying that the sector-only and the occupation-only models are able to capture exactly the same processes. For the elasticities of substitution further from 1 (smaller for complements, larger for substitutes), the prediction of the occupation-only models get closer to those of the sector-occupation models.

However, a different picture emerges for the predictions for relative sectoral prices.



(a) Occupational employment



(b) Occupational relative wages

Figure 8: The evolution of occupational labor outcomes

This figure plots occupational employment shares and relative wages in the data (dark gray) and in the restricted models for different elasticities of substitution in production. The light gray lines show predictions from the model restricted to sector-only or occupation-and-sector-only technological change, and the yellow lines of the occupation-only model. The pattern of the lines refer to the following values of η : 0.4 in the solid, 0.6 (the baseline value) in the dashed, 1.4 in the dotted, and 1.6 in the dash-dotted lines.

Figure 9 shows the models' predictions for these prices under alternative values for the elasticity of substitution between occupations. While it is still true that the predictions of the sector-only and the sector-occupation technology models are essentially independent of this elasticity, the implications of occupation-only technological change models hinge crucially on the value of η . When $\eta > 1$ this model implies movements in relative prices that are at odds with the data, but with $\eta < 1$ its predictions are consistent.

Since there is no consensus in the literature on the value of the elasticity of substitution between different occupations (yet) we draw two conclusions from our sensitivity analysis. First, our analysis highlights that even small changes in η matter for models that allow for productivities to evolve only at the occupation level, this suggests that to infer robust results from economic models these should allow for technological progress to have a sector-specific component. Second, occupation-only technological

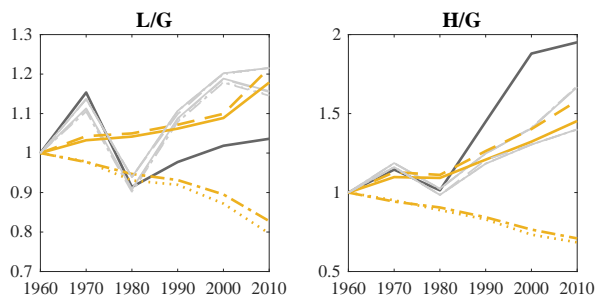


Figure 9: Sectoral relative prices for different η

This figure plots the path of the relative sectoral prices in the data (dark gray) and in the restricted models for different elasticities of substitution in production. The light gray lines show predictions from the model restricted to sector-only or occupation-and-sector-only technological change, and the yellow lines of the occupation-only model. The pattern of the lines refer to the following values of η : 0.4 in the solid, 0.6 (the baseline value) in the dashed, 1.4 in the dotted, and 1.6 in the dash-dotted lines.

change gives rise to changes in relative prices, and changes in all other outcomes, in line with the data when $\eta < 1$. This is in line with the work by Duernecker and Herrendorf (2016) who show in a two-sector two-occupation model that a production elasticity of substitution between occupations less than one is needed to explain structural change based on occupation-specific productivity growth alone.

6 Conclusion

To understand the nature of technological changes that drive structural change and labor market polarization, we set up a parsimonious yet flexible model, since it is not possible to directly compute productivity growth at the occupation level. After specifying functional forms for sectoral production functions, we extract sector-occupation cell productivities over time from observed occupational wages and labor income shares. We then use a factor model to decompose these into (time-varying) neutral, sector, and occupation factors as well as a residual component. We find that occupation-specific effects explain most of the changes in cells' productivity, but also a sizable contribution from sector effects and components idiosyncratic to sector-occupation cells.

Given the large overlap between sectoral and occupational employment, we find that models that restrict the nature of technological change to only one factor, can also

perform very well in generating patterns in employment and wages as seen in sector or occupation data. However, our analysis of restricted models highlights that only those that feature occupation-specific productivity changes can replicate reallocations within sectors, i.e. the relative fall of routine employment and income within goods and both types of services. Yet, the conclusions from a model with occupation-only technological change are sensitive to assumed values of the elasticity of substitution between labor in different occupations, whereas models with sector-specific technologies are very robust to this elasticity. This suggests that there is a good case for writing multi-sector multi-occupation models in a way that allows for technologies to evolve not only in one dimension, but at least at the sector- and the occupation-level or in a flexible way as in the model we propose.

References

- Daron Acemoglu and Veronica Guerrieri. Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy*, 116(3):pp. 467–498, 2008.
- David H. Autor and David Dorn. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97, 2013.
- David H. Autor, Lawrence F. Katz, and Melissa S. Kearney. The Polarization of the U.S. Labor Market. *The American Economic Review*, 96(2):pp. 189–194, 2006.
- Zsófia L. Bárány and Christian Siegel. Job Polarization and Structural Change. *American Economic Journal: Macroeconomics*, forthcoming, 2017.
- Francisco J. Buera, Joseph P. Kaboski, and Richard Rogerson. Skill biased structural change. Working Paper 21165, National Bureau of Economic Research, May 2015.
- Georg Duernecker and Berthold Herrendorf. Structural transformation of occupation employment. Working paper, February 2016.
- 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.
- Piyabha Kongsamut, Sergio Rebelo, and Danyang Xie. Beyond Balanced Growth. *The Review of Economic Studies*, 68(4):pp. 869–882, 2001.
- 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):pp. 429–443, 2007.

A Appendix

A.1 Classification

Occupations are classified as:

1. Manual: low-skilled non-routine

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

2. Routine

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. Abstract: skilled non-routine

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

Industries are classified into sectors in the following way:

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

2. Goods: agriculture, forestry and fishing, mining, construction, manufacturing

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

A.2 Derivations

Expressing the α s as a function of observables. Multiply (1) with w_m/w_r and (2) with w_a/w_r to get:

$$\begin{aligned}\frac{\theta_{mJ}}{\theta_{rJ}} &= \left(\frac{w_r}{w_m}\right)^{\eta-1} \left(\frac{\alpha_{mJ}}{\alpha_{rJ}}\right)^{\eta-1}, \\ \frac{\theta_{aJ}}{\theta_{rJ}} &= \left(\frac{w_r}{w_a}\right)^{\eta-1} \left(\frac{\alpha_{aJ}}{\alpha_{rJ}}\right)^{\eta-1}.\end{aligned}$$

Re-arrange to get:

$$\begin{aligned}\frac{\alpha_{mJ}}{\alpha_{rJ}} &= \left(\frac{\theta_{mJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_m}{w_r}, \\ \frac{\alpha_{aJ}}{\alpha_{rJ}} &= \left(\frac{\theta_{aJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_a}{w_r}.\end{aligned}$$

A.3 Calibration

Using that the total earnings of any occupation can be expressed as either their labor supply times their wage rate, or as the sum of their earnings across all sectors, the following two equations must be satisfied by the data:

$$\begin{aligned}\frac{f_a}{1-f_a-f_m} \frac{\omega_a}{\omega_r} &= \frac{\Psi_L \theta_{aL} + \Psi_G \theta_{aG} + \Psi_H \theta_{aH}}{\Psi_L \theta_{rL} + \Psi_G \theta_{rG} + \Psi_H \theta_{rH}}, \\ \frac{f_m}{1-f_a-f_m} \frac{\omega_m}{\omega_r} &= \frac{\Psi_L \theta_{mL} + \Psi_G \theta_{mG} + \Psi_H \theta_{mH}}{\Psi_L \theta_{rL} + \Psi_G \theta_{rG} + \Psi_H \theta_{rH}}.\end{aligned}$$

Given that we calculated relative occupational wage rates, $\frac{\omega_m}{\omega_r}$ and $\frac{\omega_a}{\omega_r}$, occupational labor income shares within sectors, θ_{oJ} for $o \in \{m, r, a\}$ and $J \in \{L, G, H\}$, as well as the labor income share of each sector, Ψ_L, Ψ_G, Ψ_H , we can calculate the implied values for f_m, f_r , and f_a .

It is worth to note that these equations always hold in the model. Therefore when we match the labor income share of each sector, i. e. the Ψ_J s, in the calibration of the consumption side of the model, these equations guarantee that we also match the occupational wage rates.

Table 3 contains the targets used in the calibration.

Table 3: 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	1.329	1.637	2.050	2.538	2.794
ω_m/ω_r	0.800	0.850	0.801	0.849	0.861	0.893
ω_a/ω_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 Factor Model

Figure 10 shows the counterfactual evolution of cell productivities using the predictions based on the factor model and also when shutting down some factors.

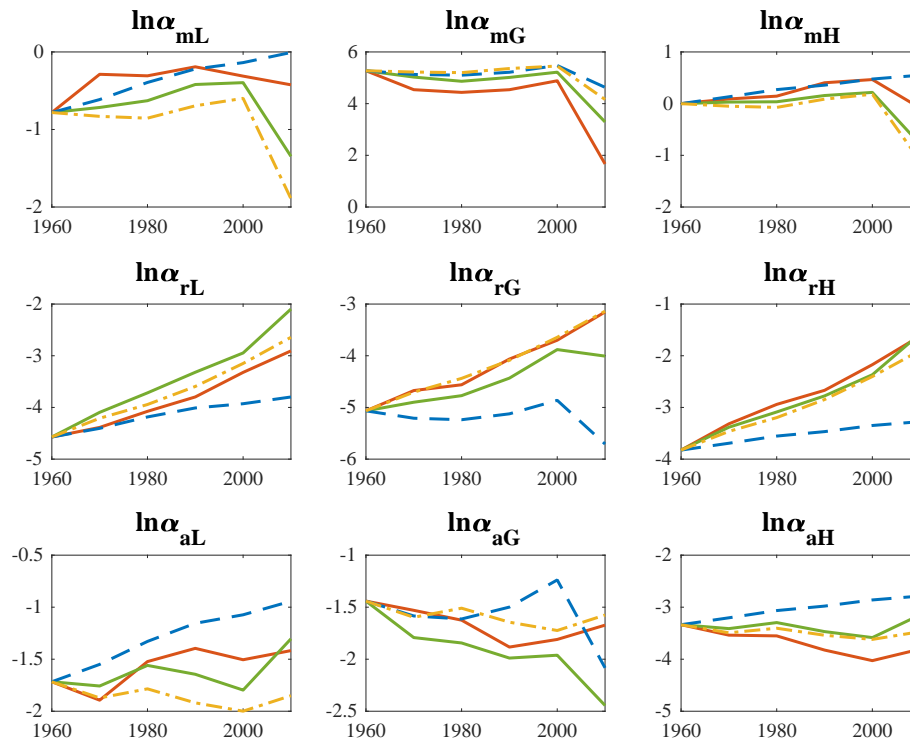


Figure 10: The evolution of cell productivities: baseline and factors

This figure shows in the solid red line the cell productivities as calculated from the data using the economic model. The columns refer to the three sectors (L, G, H) and the rows to the three occupations (m, r, a). The dash-dotted line shows the predictions based on both sector and the occupation factors (as specified in (13)). The dashed line shows the contribution of the time-varying sector factor only (as constructed in (15)) and the dotted line of the occupation factor only ('time & occ' as constructed in (16)).

A.5 Restricted Models

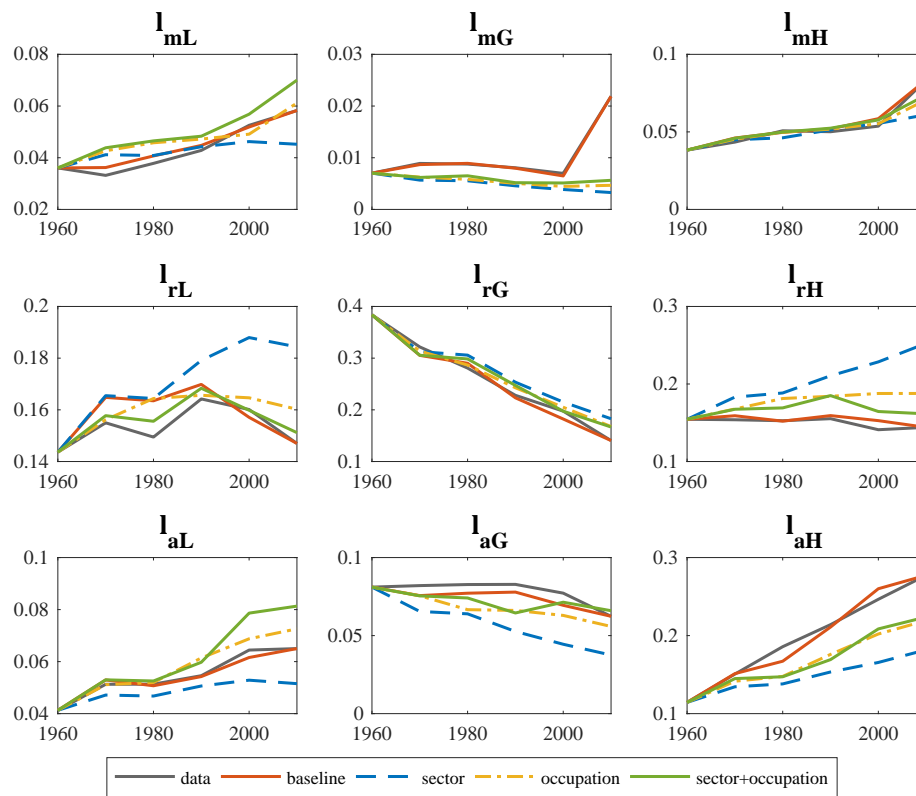
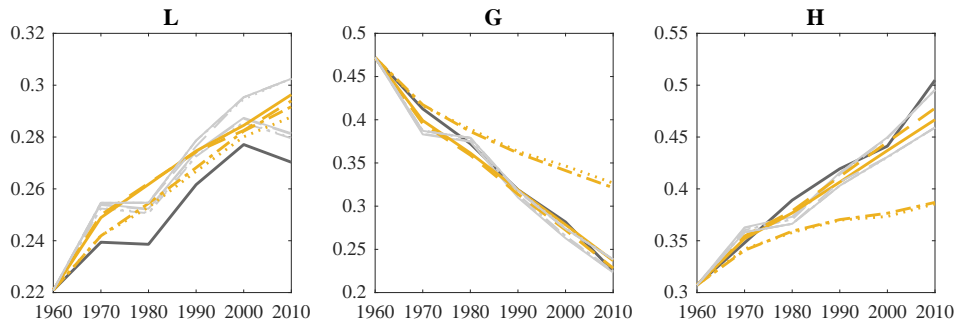
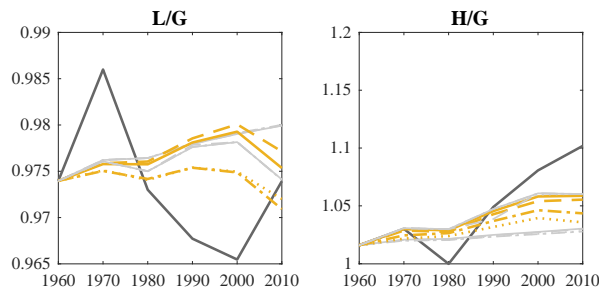


Figure 11: Cell employment shares

This figure plots the path of the the cell employment share in the data and for various paths of cell productivities, as discussed in the text.



(a) Sectoral employment



(b) Sectoral relative wages

Figure 12: The evolution of sectoral labor outcomes for different η

This figure plots the path of the sectoral employment shares and relative wages in the data (dark gray) and in the restricted models for different elasticities of substitution in production. The light gray lines show predictions from the model restricted to sector-only or occupation-and-sector-only technological change, and the yellow lines of the occupation-only model. The pattern of the lines refer to the following values of η : 0.4 in the solid, 0.6 in the dashed (the baseline value), 1.4 in the dotted, and 1.6 in the dash-dotted lines.

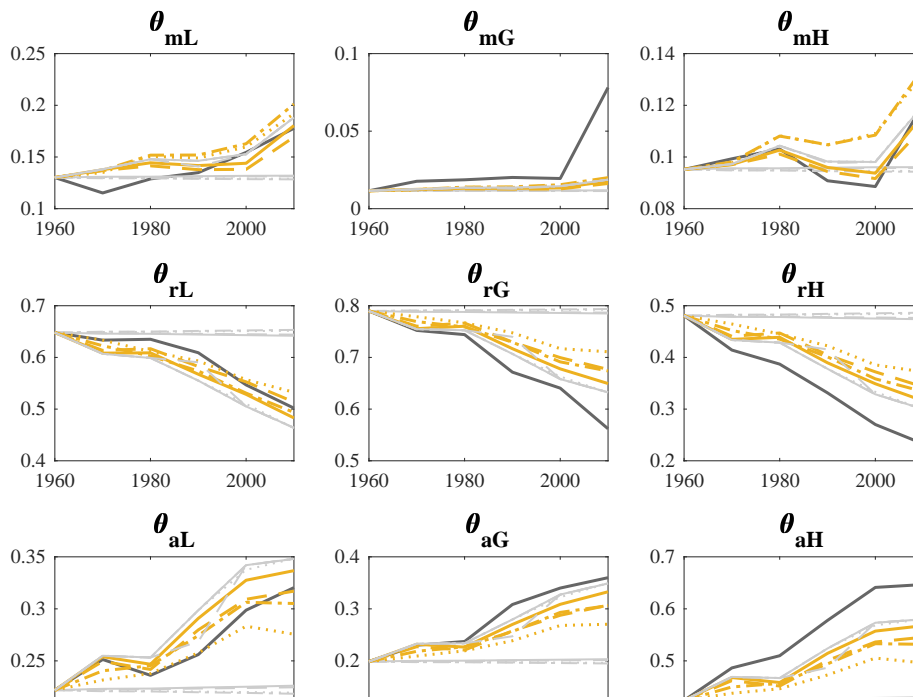


Figure 13: Income shares within sectors for different η

This figure plots the path of income shares within sectors in the data (dark gray) and in the restricted models for different elasticities of substitution in production. The light gray lines show predictions from the model restricted to sector-only or occupation-and-sector-only technological change, and the yellow lines of the occupation-only model. The pattern of the lines refer to the following values of η : 0.4 in the solid, 0.6 in the dashed (the baseline value), 1.4 in the dotted, and 1.6 in the dash-dotted lines.