Biased Technological Change and Employment Reallocation*

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

September 2020

Abstract

To study the drivers of the employment reallocation across sectors and occupations between 1960 and 2017 in the US we present a model where technology evolves at the sector-occupation cell level. Drawing on key equations of the production side we infer technologies directly from the data. We assess the magnitude of neutral, sector-, and occupation-specific components in technological change and study their consequences for labor market outcomes in general equilibrium where occupational choice and demands for sectoral outputs change endogenously with technology. Our findings indicate a major role for occupation-specific technological changes.

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

*Parts previously circulated as 'Disentangling Occupation- and Sector-specific Technological Change'. We wish to thank the editor and the referees, as well as Francesco Caselli, Georg Duernecker, Tim Lee, Miguel León-Ledesma, Guy Michaels, Mathan Satchi, Ákos Valentinyi and numerous seminar and conference participants for their comments and suggestions.

[†]Sciences Po, CEPR and ROA. Email: zsofia.barany@sciencespo.fr

[‡]University of Kent (School of Economics and MaGHiC), CEPR, GLO and ROA. Email: c.siegel@kent.ac.uk

1 Introduction

There have been substantial changes in the structure of employment over recent decades in most developed countries. Most economies are undergoing structural change, whereby labor reallocates across sectors, while at the occupational level labor markets have been polarizing, with employment shifting out of middle-earning routine jobs to low-earning manual and high-earning abstract jobs. There is a wide consensus in the literature that one of the main drivers of each of these patterns separately is biased technological change. Technological change biased across sectors is a key mechanism for structural change, technological change biased across production factors (occupations/tasks) is a prominent explanation for job polarization.[1] The benefits of such non-neutral technological progress are not equally distributed, with some workers displaced or their skills becoming less valuable in the labor market. These distributional impacts have generated an interest in policies to counteract the adverse effects of technological change. To devise and implement such policies, the biases of technological change and their implications for the labor market needs to be understood.

Figure 1 shows that the evolution of sectoral and of occupational employment shares are closely connected. Virtually all of the decline in the goods-producing sector's employment share is due to a decline in routine employment in this sector. Conversely, most of the increase in high-skilled service employment is a rise in employment in abstract occupations in that sector.[2] Because of this tight connection between employment reallocations across sectors and across occupations, models that do well in matching sectoral outcomes do also quite well in replicating certain aspects of occupational outcomes, and vice versa. However, this also poses a challenge for establishing what the true drivers of these phenomena are; is it technological change

[1]Baumol (1967) and Ngai and Pissarides (2007) argue that productivity growth differences lead to structural change. Other mechanisms have been proposed such as income effects in Kongsamut, Rebelo, and Xie (2001) and Boppart (2014), as well as differential factor intensities across sectors as in Caselli and Coleman (2001) and Acemoglu and Guerrieri (2008). Autor, Katz, and Kearney (2006), Goos and Manning (2007), Autor and Dorn (2013), Goos, Manning, and Salomons (2014), Michaels, Natraj, and Van Reenen (2014) argue that differences in technological progress across occupations (routinization) lead to polarization. Other mechanisms for polarization include offshoring (Grossman and Rossi-Hansberg (2008)) and consumption spill-overs (Manning (2004), Mazzolari and Ragusa (2013)).

[2]We show in Figure A2 in the appendix that even for a finer classification of occupations and sectors there is a clear, though somewhat less tight, link between occupational and sectoral employment changes.

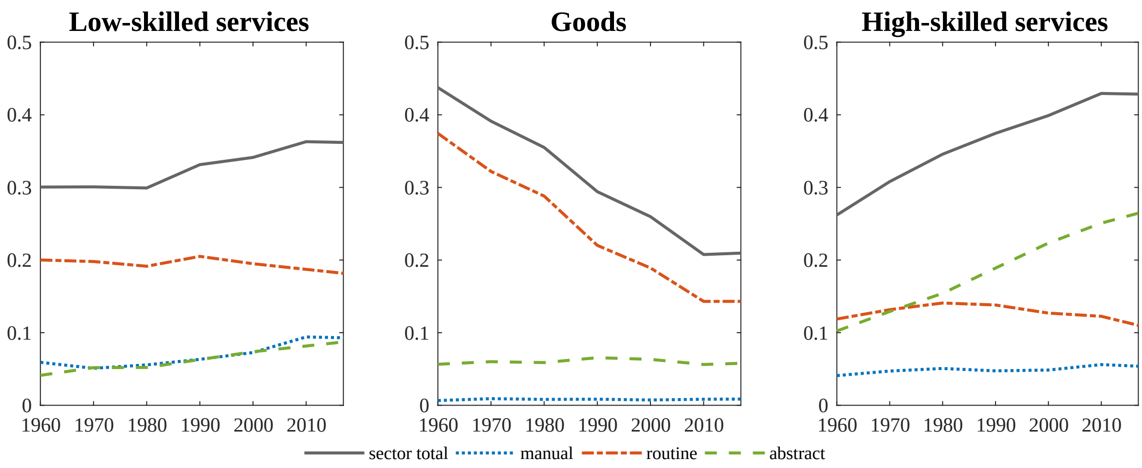

Figure 1: sector-occupation hours worked shares 1960-2017

Notes: The data is taken from IPUMS US Census data for 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) for 2010 and 2017. 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 and 2017. 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. See Appendix A.1 for the classification of occupations and industries.

biased across industrial sectors or across tasks/occupations? Sector-specific technological change can reflect, for example, product innovations, and occupation-specific technological change can capture innovations in processes, i.e. in how tasks are completed. We believe that the technology used by workers depends on both their sector and their occupation. Hence we allow for technologies to evolve at the sector-and-occupation level.[3] Next we identify patterns common to sectors and to occupations. With this at hand we can study which type of technological biases are responsible for the observed changes in market outcomes. We view our model as an important and useful first step towards evaluating policies responding to the rapid changes in the labor market.

In this paper we shed light on the types of technological change that drive the observed changes in (labor) market outcomes, without imposing any restrictions a priori on how technological change is biased. Specifically, we take a simplified version of

[3]For example consider an accountant in the healthcare sector. An improvement in general accounting software (process-innovation) is likely to translate into improvements in accounting software in healthcare. Equally, changes in the range of services offered in the healthcare sector (product-innovation) is likely to affect the accountant's productivity. But potentially there is an interaction between the two: the more wide-ranging or complex the set of services are, the larger the potential gain of a better software is. Such interactions are what make technology sector-and-occupation specific.

the model in our previous work, Bárány and Siegel (2020), and use the same methodology as there to identify technologies. As in that paper, we decompose the changes in these technologies into occupation- and sector-specific components using a factor model. The novelty of this paper is that we study the contribution of these components to labor market outcomes. To do this, we specify a general equilibrium model, where individuals' occupational choice and the demand for sectoral output responds endogenously to changing technologies. We use this general equilibrium model to quantify the role of the various components via counterfactual simulations in the reallocation of employment across sectors and occupations, as well as in the evolution of occupational wages and of sectoral prices.

We assume a CES production function in manual, routine, and abstract labor in each sector. This is a simplified version of the sectoral production function we use in Bárány and Siegel (2020). As in that paper, we draw on key equations of the production side of the model together with data from the US Census and from the U.S. Bureau of Economic Analysis between 1960 and 2017 to extract sector-and-occupation specific technologies. We decompose the changes in the inferred sector-occupation cell technologies using a factor model into neutral, sector- and occupation-specific components.[4] We find that these components jointly explain around 90 percent of the variation in cell technology growth. We then construct the following counterfactual cell technology series: a sector-only path where technology is allowed to be biased only across sectors, an occupation-only path allowing for a bias only across occupations, and a neutral technology path that shuts down all biases of technological change. This approach is more rigorous than calibrating sector- and occupation-specific growth rates, as those would depend on the normalizations implemented.[5]

The goal of this paper is to evaluate the role of neutral, sector- and occupation-specific components of technological change in labor market outcomes. In order to

[4]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 technology series based on our factor model estimates.

[5]Each separately cannot be calibrated, some sector-specific growth rates and/or some occupation-specific growth rates need to be normalized. It can be shown that different normalizations have different implications in terms of the role of sector- and occupation-specific technology growth.

conduct counterfactual exercises we specify a general equilibrium model. We assume 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 feed the counterfactual technology paths into the model to determine how important each component of technological change is in explaining various outcomes of interest. We find that qualitatively all counterfactual technology series generate employment and wage paths in line with the data. Quantitatively however, the occupation-specific component is needed to get close to the data. To explain the evolution of sectoral prices, all components are needed. For occupational income shares within sectors and employment shares at the cell level, the neutral and the sector-specific component have almost no effect, whereas the technology component idiosyncratic to the sector-occupation cell has a significant role.

There are many papers on the impact of technological change biased across tasks (occupations) or across sectors, but there is no consensus in the literature about the true nature of technological biases, potentially due to the trends shown in Figure 1. In Bárány and Siegel (2018) we show that 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 because of differences in occupation intensities across sectors. Goos et al. (2014) argue conversely that differences in productivity growth across occupations together with differences in occupation intensity across sectors can lead to employment reallocation across sectors. In a similar vein, Duernecker and Herrendorf (2016), Lee and Shin (2017) and Aum, Lee, and Shin (2018) show that differences in occupational productivity growth can generate both structural transformation and changes in occupational employment consistent with the data.[6] In contrast, we take a fully flexible approach to technological change, where it is sector-and-occupation specific. This is what allows us to evaluate the role of sector- and occupation-biases in technological change on an equal footing. An overarching conclusion of our analysis is that occupation-biases in technological change are quantitatively more important. One implication of this finding is that if one wants

[6]Lee and Shin (2017) and Aum et al. (2018) also consider sector-bias in technology, but not sector-occupation bias, and differently from us, calibrate all technological processes.

to introduce biases in technological change into models, the most important bias to introduce is across occupations. Another implication is that policies targeting workers' occupational choice might be better at improving labor market outcomes than industrial policies.[7]

The paper proceeds as follows: section 2 introduces the model. Section 3 presents the data and the model parameterization. In section 4 we first identify neutral, sector- and occupation-specific components of technological change, and then analyze the role that each of these components play in our general equilibrium model. 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 sector-occupation-specific technologies change over time, which we then decompose into common factors, as described in section 4.1. This is similar to the approach we take in Bárány and Siegel (2020). Note, here we do not model non-labor inputs but the technology

[7] In our analysis we focus on the technology side of the economy. There are no frictions or externalities in our model which would justify policy interventions. Given that we find a large role for the occupation-bias in technological change, models with occupation-bias in technology and with a frictional labor market or job specific human capital accumulation would be well suited to study policies.

parameters subsume effects stemming from changes in inputs other than labor. Not modeling capital as a distinct input to labor-augmenting technologies is consistent with the view that technological change is embodied into capital. While the focus of our work in Bárány and Siegel (2020) is to account for sources of sectoral productivity growth, here we study the implications of various forms of technological change for labor market outcomes. To do this we specify a general equilibrium model allowing us to conduct counterfactual exercises where occupational choice and demands for goods and services respond endogenously to changing technologies.

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 a *sector-occupation* 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.[8] In the initial year α_{oJ} reflects the initial level of technology as well as the intensity at which sector J uses occupation o, whereas any subsequent change over time reflects sector-occupation specific technological change. We do not make any assumptions about whether technological change occurs at the occupation or the sector level, but instead allow for α_{oJ} to evolve freely over time without imposing any restrictions.[9]

[8]We assume the same elasticity of substitution in all sectors as sector-specific estimates are not available.

[9]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 on the evolution of technologies.

All firms take prices and wages as given, and profit maximization implies that the optimal relative labor demand within a sector has to satisfy:

$$\frac{l_{oJ}^d}{l_{rJ}^d} = \left(\frac{w_r}{w_o}\right)^{\eta} \left(\frac{\alpha_{oJ}}{\alpha_{rJ}}\right)^{\eta-1} \quad \text{for } o \in \{m, a\}. \tag{2}$$

This equation shows that it is optimal to use less of one occupational labor input compared to routine labor if that occupation's wage relative to routine is higher. In addition, the occupational labor inputs' technologies within a sector matter. The larger the term $\left(\frac{\alpha_{oJ}}{\alpha_{rJ}}\right)^{\eta-1}$ is in a sector, the more occupation o compared to routine labor the sector employs optimally. 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, equal to the marginal cost, of sector J output in terms of wage rates:

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

Finally using (2) and (3) to express sector J output, optimal occupation o labor use in sector J can be expressed as:[10]

$$l_{oJ}^d = \left[\frac{p_J \alpha_{oJ}}{w_o}\right]^{\eta} \frac{Y_J}{\alpha_{oJ}}. \tag{4}$$

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

[10]The full derivations can be found in appendix A.2.

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

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

where w_o is the unit wage in occupation o and $\chi_o^i \geq 0$ is a net-of-cost multiplier for individual i when entering occupation o. It is convenient to define $\tilde{\chi}_o^i = \ln\{\chi_o^i\}$. Then an equivalent formulation to the above is

$$\ln\left(\frac{w_o}{w_k}\right) \geq \tilde{\chi}_k^i - \tilde{\chi}_o^i \quad \text{for any} \quad k \neq o, \;\; k, o \in \{m, r, a\}.$$

Note, since only $\tilde{\chi}_k^i - \tilde{\chi}_o^i$ matter, we define occupational cost differences as $\tilde{\chi}_1^i \equiv \tilde{\chi}_a^i - \tilde{\chi}_r^i$ and $\tilde{\chi}_2^i \equiv \tilde{\chi}_a^i - \tilde{\chi}_m^i$. The optimal occupational choice is summarized in Figure 2.

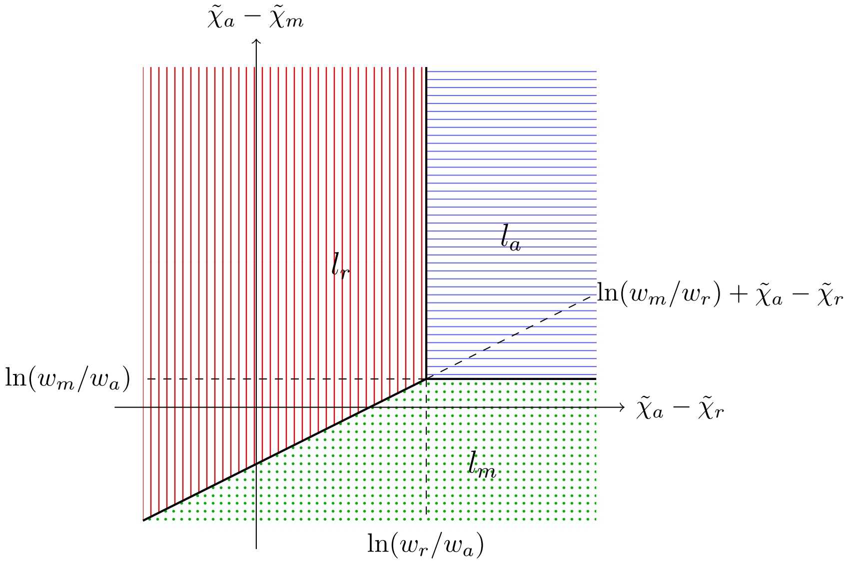

Figure 2: Optimal occupational choice

Notes: The graph shows the optimal selection of individuals into manual, routine and abstract occupations in terms of their idiosyncratic occupational cost differences 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:

$$l_m^s = \int_{-\infty}^{\infty} \int_{-\infty}^{\min\{\ln(w_m/w_r)+\tilde{\chi}_1, \ln(w_r/w_a)\}} f(\tilde{\chi}_1, \tilde{\chi}_2) d\tilde{\chi}_1 d\tilde{\chi}_2, \tag{5}$$

$$l_r^s = \int_{-\infty}^{\ln(w_r/w_a)} \int_{\ln(w_m/w_r)+\tilde{\chi}_1}^{\infty} f(\tilde{\chi}_1, \tilde{\chi}_2) d\tilde{\chi}_1 d\tilde{\chi}_2, \tag{6}$$

$$l_a^s = \int_{\ln(w_r/w_a)}^{\infty} \int_{\ln(w_m/w_a)}^{\infty} f(\tilde{\chi}_1, \tilde{\chi}_2) d\tilde{\chi}_1 d\tilde{\chi}_2, \tag{7}$$

where $f(\tilde{\chi}_1, \tilde{\chi}_2)$ is the joint probability density function of 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 by p_G, and 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:

$$C_S = \left(\frac{a_S}{p_S} \right)^{\varepsilon} \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^{\varepsilon} p_L^{1-\varepsilon} + a_G^{\varepsilon} p_G^{1-\varepsilon} + a_H^{\varepsilon} p_H^{1-\varepsilon}} - \bar{c}_S \quad \text{for } S \in \{L, H\}, \tag{8}$$

$$C_G = \left(\frac{a_G}{p_G} \right)^{\varepsilon} \frac{f_m w_m + f_r w_r + f_a w_a + p_L \bar{c}_L + p_H \bar{c}_H}{a_L^{\varepsilon} p_L^{1-\varepsilon} + a_G^{\varepsilon} p_G^{1-\varepsilon} + a_H^{\varepsilon} p_H^{1-\varepsilon}}. \tag{9}$$

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 depend on the endogenous occupational wage rates, w_m and w_a, through (3), and hence can be written as $p_J = p_J(w_m, w_a)$. Using this in (8) and (9), sectoral demands also only depend on occupational wage rates, $C_J = C_J(w_m, w_a)$. Then (4) shows that optimal occupation o labor use in sector J can be expressed as a function of manual and abstract wage rates. The equilibrium then boils down to finding wage rates w_m and w_a such that the labor markets clear.

3 Extracting technologies and calibrating the worker side

To evaluate how sector-occupation-specific technology evolved over time and to study their implications for labor markets, we parameterize the model. In our model setup, there is a dichotomy that allows to back out the sector-occupation cell technologies from the data using only the production side. We therefore proceed in the following steps, similarly to Buera, Kaboski, Rogerson, and Vizcaino (2018). First, we compute cell technologies 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 (similarly to Bárány and Siegel (2020)). 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 Data targets

We use US Census and American Community Survey (ACS) data between 1960 and 2017 from IPUMS, provided by Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek (2010), to calculate occupational wage rates and occupational labor income shares within sectors, as well as each sector's share in labor income.[11] 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 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:

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

To get measures for the occupational wage rates we employ the following Mincer wage regression to control for workers' observable skills,

$$\log w_{iot} = \delta_{ot} + \beta' X_{it} + \varepsilon_{iot}, \tag{10}$$

where δ_{ot} are occupation-time effects and X_{it} is a vector of worker characteristics. From this regression we back out for each year t a wage for occupation o that is not confounded by changes in composition of worker characteristics, X_{it}. In particular, we run this regression on the Census/ACS data where the vector of worker i characteristics X_{it} is comprised of a third-order polynomial in potential experience (defined as age minus years of schooling minus 6), interacted with a gender dummy, as well as a dummy for foreign-born and non-white race. From the estimates of this regression we construct for each year the manual wage rate as $w_{mt} = exp\{\delta_{mt} - \delta_{rt}\}$ and abstract wage rate as $w_{at} = exp\{\delta_{at} - \delta_{rt}\}$, maintaining the normalization of $w_r = 1$ in all years. We measure wages in this way – rather than as 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 or in educational attainment of workers across occupations. This is similar to Buera et al. (2018), and it implies that all differences within an occupational group in hourly wages are due to

[11] In our model the share of each sector in labor income and in value added is the same as there are no other factors of production. As our main focus in this paper is on labor market outcomes, we calibrate the model to match the each sector's share in labor income, rather than in value added.

differences in the endowment of efficiency units of labor.

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 is normalized to one.[12]

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 2017 to get sectoral prices and the growth rate of GDP per full-time equivalent worker between periods.[13] Table A2 in the appendix contains all the calibration targets, and these are also plotted along with the model outcomes in section 4.

3.2 Extracting sector-occupation cell technologies

As mentioned before, given the structure of the model we can infer the technology parameters using key equations from the model's production side 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, following similar steps as in Bárány and Siegel (2020).

We set the elasticity of substitution in production to 0.6 in our baseline and conduct robustness checks around this value. While there is no consensus on the value of this in the literature, there seems to be wide agreement that occupations are complementary, implying a value of η less than one. To our knowledge the only estimate of this

[12]Note that these are not exactly the same as the share of hours worked by occupation in the data. Figure A1 in the appendix shows the evolution of the model implied sector-occupation employment shares over time in the same format as Figure 1. Comparing these two figures reveals that the trends in actual hours and in implied employment shares are very similar.

[13]The industry classification system changed from SIC to NAICS in the middle of our sample, both systems are different from the classification used in the IPUMS Census/ACS. Table A1 in the appendix shows the mapping of fine industries of each system into our broad sector categories.

elasticity is in Goos et al. (2014), who estimate this for 21 occupation 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 for 2 occupations, while Lee and Shin (2017) calibrate a value of 0.70, and Aum et al. (2018) a value of 0.81 for this same parameter for 10 occupations. With fewer, coarser occupations this elasticity is likely to be smaller, and we hence, as in Bárány and Siegel (2020), set $\eta = 0.6$ for our 3 occupations.

We calculate the nine cell-specific technologies, 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 also take for now as given occupational wage rates, occupational labor supplies, and the sectoral distribution of income. However, the calibration of the household side of our model guarantees that these are matched in general equilibrium in the initial and final period (1960 and 2017). Our model allows us to express cell-specific technologies 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 condition (2):

$$\frac{\alpha_{oJ}}{\alpha_{rJ}} = \left(\frac{\theta_{oJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_o}{w_r} \quad \text{for } o \in \{m, a\}. \tag{11}$$

The relevant wages in the above equation are wages per effective unit of labor. As we explained earlier we measure occupational wage rates from the Mincer wage regression (10). This procedure accounts for workers' observable characteristics, and assuming that there are no other unobservable skills, this gives us the relevant wages per efficiency units of labor. However, if workers are self-selecting on unobservable skills, this would confound our measured wage rates. Then, if the selection changes over time, our framework would assign it to changing technologies. In our analysis we abstract from such self-selection and hence we can utilize wage data to infer technologies. Note, our modeling of occupational choice satisfies the requirement of no selection on unobservable skills.

Using the expression for relative αs within sectors (11) in the expression for sectoral prices (3) allows us to express relative αs across sectors within each period as:

$$\frac{\alpha_{mJ}}{\alpha_{mK}} = \frac{p_K}{p_J}\left(\frac{\theta_{mJ}}{\theta_{mK}}\right)^{\frac{1}{\eta-1}}. \tag{12}$$

This equation gives the relative αs across sectors such that relative marginal costs, given optimal labor use, are equal to relative prices in the data.

Finally, the overall growth rate of output per (full-time-equivalent) 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.[14]

Thus, we have shown how to extract sector-occupation specific technologies from the data conditional on the elasticity of substitution across occupations.

3.3 Calibration of the cost distribution and of the consumption side

To close the model we need to parameterize the household side. In calibrating the distribution of the occupational cost differences, we assume that $f(\tilde{\chi}_1, \tilde{\chi}_2)$ is a time-invariant bivariate normal distribution,[15] and in our baseline we assume that the two cost differences are uncorrelated, i.e. $\rho = 0$. Given this ρ, we calibrate the two means (μ_1, μ_2) and the diagonal elements of the variance-covariance matrix (σ_1^2, σ_2^2) such that in the initial and final period for given unit wages the occupational cost difference distribution is able to match the employment shares. While all these parameters are calibrated jointly, we can infer which moments of the data are more informative about which parameter. The intuition can be gained by inspecting Figure 2. The initial employment shares given the wages are informative about the means of these cost differ-

[14]Note that we only rely on relative wages and relative prices within a period, not changes over time. This is important, as in our data prices are subject to a different normalization than wages, implying that changes in prices are not directly comparable with changes in wages.

[15]For simplicity we assume that the distribution is time invariant. Allowing for a changing distribution of cost differences (for example as in Caselli and Coleman (2001)) would require more parameters to be calibrated, and it would not affect the sector-occupation cell technologies, neither their decomposition into various components, nor the results from the baseline model.

ences, while changes in employment given the change in wages are informative about the variance of these cost differences.[16]

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 sectoral income shares in the initial and final year, i.e. in 1960 and 2017. This also guarantees that the relative occupational wages in 1960 and 2017 are met in equilibrium. These four parameters are calibrated jointly in the general equilibrium model, but the change in expenditures given the change in prices between 1960 and 2017 informs the magnitude of the non-homotheticity terms, $\bar{c}_J$, and the level of expenditures in 1960 are informative about the weights, a_J. One way to understand the magnitude of the non-homotheticity terms is to calculate their share in the total consumption of the good, i.e. $\bar{c}_J/(\bar{c}_J + C_{J,t})$. We show these values for 1960 and 2017 in Table 1. In 1960 over 80 percent of the consumption value of low- and high-skilled services comes from the endowment values, and while the share of consumption value coming from the endowment falls for both types of services, it still represents a large fraction in 2017. These high shares imply that there are sizable and important non-homotheticities. This in line with the results of Boppart (2014) and Comin, Lashkari, and Mestieri (2020) who find a substantial role for income effects as drivers of structural change. Table 1 contains the calibrated parameters of the model which, together with the evolution of the αs as backed out from the data, fully specify the calibrated model.

We conduct robustness checks on the importance of the correlation parameter, ρ, and the elasticity of substitution in consumption, ε, and find neither to matter much; see appendix A.5. This is partly due to the calibration procedure, as the initial and final period outcomes are guaranteed to be the same in the baseline across all calibrations.

[16]This procedure can be implemented for any value of the correlation parameter, ρ. This means that the moments that we target are not informative about the value of this correlation. However, this value turns out to have hardly any impact on any (counterfactual) model outcome, see appendix A.5.

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 $\tilde{\chi}$ distribution	0
μ_1, μ_2	mean of $\tilde{\chi}$ distribution	(-0.3364, 1.0003)
σ_1^2, σ_2^2	variance $\tilde{\chi}$ distribution	(0.0649, 2.0789)
$\overline{c}_L$	non-homotheticity term in L	0.0034
	implying $\overline{c}_L/(\overline{c}_L + C_{L,t})$ in 1960 and 2017	0.82 & 0.59
$\overline{c}_H$	non-homotheticity term in H	0.0120
	implying $\overline{c}_H/(\overline{c}_H + C_{H,t})$ in 1960 and 2017	0.88 & 0.73
a_L	weight on L	0.00550
a_H	weight on H	0.99448

Notes: The table shows the calibrated parameters of the model. The first panel shows the parameters set outside the model. The next panel shows the parameters of bivariate normal distribution of $(\tilde{\chi}_1, \tilde{\chi}_2)$, calibrated separately based on occupational relative wages and employment shares in 1960 and 2017, conditional on ρ. The final panel shows the parameters of the utility function calibrated in general equilibrium conditional on all other parameters, including the sector-occupation specific technologies, $\alpha_{oJ,t}$s. We also show the implied value of purchased consumption relative to the total consumption value, $\overline{c}_S/(\overline{c}_S + C_{S,t})$, in the two service sectors, $S = L, H$ and for $t = 1960$ and $t = 2017$.

4 The role of technological biases

In this section we first decompose the change in the extracted cell level technologies into neutral, sector- and occupation-specific components, similarly to Bárány and Siegel (2020). We then use our general equilibrium model to quantify the role of each component of technological change for various outcomes of interest. It is important to note that in order to correctly assess the impact of various technological biases, we need to consider the endogenous response of sectoral demands and of occupational labor supplies to technological change.

4.1 Decomposition of technological change

We decompose the change in the extracted series of sector-occupation specific technologies into neutral, sector- and occupation-specific components using a factor model. In particular, we run the following regression on the log difference of the cell technologies,

$$\Delta \ln \alpha_{oJ,t} \equiv \ln \alpha_{oJ,t} - \ln \alpha_{oJ,t-1} = \beta_t + \gamma_{J,t} + \delta_{o,t} + \varepsilon_{oJ,t}. \quad (13)$$

In the regression we use each cell's average labor income share between period $t-1$ and t as weights $\omega_{oJ,t}$ to reflect the relative importance of the sector-occupation cell.[17] We restrict the average sector effect ($\sum_o \sum_J \omega_{oJ,t}\gamma_{J,t}$) and the average occupation effect ($\sum_o \sum_J \omega_{oJ,t}\delta_{o,t}$) across all cells to be zero. These restrictions imply that β_t, the time effect, is the average technological change between period $t-1$ and t in all cells. The time-varying sector effects, $\gamma_{J,t}$, capture sector-wide innovations that affect the technology of all workers in that sector equally regardless of their occupation. Technological change that is common to workers of a given occupation, but are independent from the sector, are assigned to the time-varying occupation effects, $\delta_{o,t}$. The residual $\varepsilon_{oJ,t}$ reflects technological changes idiosyncratic to workers in a sector-occupation cell. The R^2 of this regression is 89.4% meaning that neutral, sector- and occupation-specific components jointly describe the evolution of cell technologies very well, and that only around 10% of the variation is idiosyncratic to the sector-occupation cell.

To asses how much of the evolution of cell technologies is explained by the neutral, the sector-specific and the occupation-specific components respectively, we build counterfactual cell technology series which in turn shut down various components. All series are constructed starting from the extracted initial cell technology $\ln \alpha_{Jo,t}$. We then add to this in turn counterfactual series for $\Delta \ln \alpha_{Jo,t}$ constructed from the neutral component only ($\widehat{\beta}_t$, 'neutral'), from neutral and sector-specific components only ($\widehat{\beta}_t + \widehat{\gamma}_{J,t}$, 'sector-only' as there are biases only across sectors), from neutral and occupation-specific components only ($\widehat{\beta}_t + \widehat{\delta}_{o,t}$, 'occupation-only'), and from neutral, sector- and occupation-specific components ($\widehat{\beta}_t + \widehat{\gamma}_{J,t} + \widehat{\delta}_{o,t}$, 'all factors').

When evaluating the explanatory power of the sector-only and the occupation-only predictions, it is important to bear in mind that these series 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. Those series would pick up differential technological change across sectors (or occupations) that originates from the sectors using occupations at different intensities (or the occupations being used at different intensities across sectors).

[17] 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 A2. The results are very robust to alternatives, such as using employment shares, or using year $t-1$ or year t shares, rather than averages.

To assess how well the various counterfactual series describe the extracted technologies, we calculate the following distance measure:

$$D = \frac{\sum_{o,J,t} \omega_{o,J,t}(\widehat{\Delta \ln \alpha_{Jo}}_{,t} - \ln \alpha_{Jo,t})^2}{\sum_{o,J,t} \omega_{o,J,t}(\Delta \ln \alpha_{Jo,t} - \overline{\ln \alpha})^2}. \quad (14)$$

The smaller this number, the closer the prediction is to the data; a value of zero implies a perfect fit. Note that $1 - D = R^2$ for the predictions generated from all factors and from time-specific components. Table 2 shows this distance measures for our baseline classification with 3 sectors and 3 occupational groupings in the top row. The occupation-only prediction with a distance of 0.168 is close to the prediction based on all factors (0.106). These small numbers imply these predictions are quite close to the data. In contrast, the distance of the sector-only prediction is much larger and quite close to that of the neutral prediction. Nonetheless, the difference in the distance measure between the all factors and the occupation-only predictions indicates that there is a role for the sector-bias in technological change as well.

Table 2: Distance measures

	neutral	sector-only	occ-only	all factor
3 sectors – 3 occupations	0.944	0.906	0.168	0.106
12 industries – 3 occupations	0.964	0.782	0.555	0.302
12 industries – 10 occupations	0.980	0.914	0.616	0.445

Notes: The table shows the distance measure defined in (14) for the four counterfactual predictions in the columns. Each row represents an alternative classification of industries and occupations, with our baseline in the first row. See appendix A.4 for the details of the classifications.

The table also shows distance measures for finer classifications with 12 industries, as well as with 10 occupations. For a finer classification of industries and occupations the fit of the factor model gets worse (the distance measure of all factors increases from 0.106 to 0.445), indicating that the relevance of technological change specific to the industry-occupation pair increases. This suggests that there is a more significant role for technologies that affect more narrowly defined occupations within a given industry than in our more aggregated analysis with the 3 industrial sectors and 3 occupations.[18] However, whatever the slicing of the data, the extracted technologies are very far from neutral (with the distance measure close to 1, the R^2 of neutral technologies is close to 0). Moreover, also in the partition with 12 industries and 3 or 10 occupations, occupational shocks still drive most of the variation in cell technologies relative to sectoral shocks.[19]

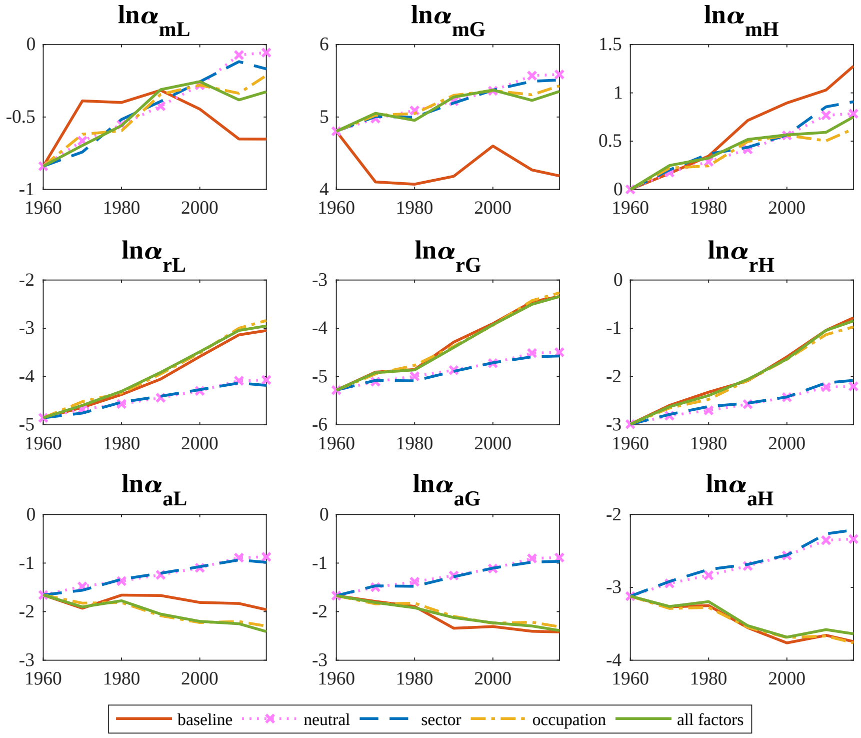

Figure 3: Baseline and counterfactual cell technologies

In Figure 3 we show the path of cell technologies as extracted from the data, as well as the different predicted technology paths as explained above. The red solid lines show the baseline cell technologies, the green solid lines show the prediction based on all factors, the pink dotted lines with the x-marker the neutral, the blue dashed lines the sector-only, while the yellow dashed-dotted lines show the occupation-only

[18]Across the different specifications there are more cells (as there are more industries and more occupations), and hence the number of observations (changes in log cell technologies) increases. While the number of cells is the product of the number of occupations and industries, the number of regressors is linear in the sum of the number of occupations and industries. Therefore, the reduction in the R^2 across these regressions is partly to be expected: the number of observations increases much more than the number of regressors. It is important to note, however, that the extracted cell technologies itself change, and this by itself could affect the explanatory power of neutral, occupation and industry components.

[19]This echoes the findings in Lee and Shin (2017) who consider 2 industries and 11 occupations and find a more important role for occupation-specific technological change.

counterfactuals. The figure confirms the findings of Table 2. The all factors and the occupation-only predictions are quite close to each other and to the baseline, whereas the neutral and the sector-only predictions are close to each other, but further away from the baseline.

Figure 3 reveals that in some cells the extracted technologies show a regress, i.e. a decline in the sector-occupation specific technology.[20] While this might seem surprising, it is not uncommon to see some technological regress when considering sector-specific factor-augmenting technologies.[21] Could this finding be an artifact of omitting capital from the sectoral production functions? In Bárány and Siegel (2020) we consider two types of capital, and still find technological regress in some sector-occupation cells, so it cannot be entirely due to omitting capital.[22] There are three different explanations for the technological regress revealed in Figure 3. First, this might reflect a compositional change in the tasks performed within a sector-occupation cell, perhaps towards more time consuming ones. Second, it is possible that technological change does not only take a factor-augmenting form, but also directly alters the relative importance of occupational labor inputs (see for instance Acemoglu and Autor (2011) or Acemoglu and Restrepo (2019)). In particular, one could imagine that in the sectoral CES production functions there are also weights on the effective input of different occupations, and that these weights also change over time.[23] In our framework we

[20]Note, we assume that wages differ only across occupations. As a robustness check we have also extracted technologies with sector-occupation specific wages and the resulting technology series are very similar to our baseline results.

[21]For instance, Herrendorf, Herrington, and Valentinyi (2015) find negative capital-augmenting technological change in the manufacturing and service sectors, and Antràs (2004) finds in aggregate U.S. data decreasing capital-augmenting technology as well.

[22]It is not hard to show that with the simplest modeling of capital, where it is combined with the labor aggregate in a CES fashion, equation (11) would be unchanged, implying the same relative change in technologies within a sector. This also implies that including capital in such a way would change all technologies within a sector proportionally. In fact, in Bárány and Siegel (2020) we find very similar technology growth patterns within sector across occupations, but the differences across sectors are much more pronounced. This is due to differential paths in capital accumulation and in capital income shares across sectors.

[23]For example the following variant of equation (1) could describe sectoral production:

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

In this formulation the β_{oJ}s are share parameters or weights, and the $\tilde{\alpha}_{oJ}$s are factor-augmenting technologies. In all first order conditions these two show up together as $\beta_{oJ}\tilde{\alpha}_{oJ}^{(\eta-1)/\eta}$ which is equivalent to $\alpha_{oJ}^{(\eta-1)/\eta}$ in our original formulation.

implicitly assume that such weights are constant. Thus what we infer as a decline in a factor-augmenting technology could in fact be an increase in the corresponding weight. There is nothing that would allow us to distinguish between changes in the weights and changes in the factor-augmenting technologies. However, the two types of changes have the same implications in our model. Third, selection into occupations based on unobservable skills potentially could explain part of the technological regress that we found. If there is selection on unobservable skills, then the average efficiency typically declines in expanding occupations and increases in shrinking occupations. As we discussed in relation to (11), such changes in average worker efficiency would be picked up in the changes of our extracted technologies. Indeed Figure 3 shows that it is mainly in the expanding cells, where worker efficiency might have fallen due to selection, that we find technological regress. However, it is difficult to assess the magnitude of such a selection effect. In Bárány and Siegel (2018) we found that selection across sectors accounts for only about 10% of observed sectoral labor productivity growth. In light of this, it seems unlikely that selection could explain a large part of the technological regress.

4.2 The role of technologies in equilibrium outcomes

We now study the role of the different components of technological change in the evolution of various outcomes in our general equilibrium model. Figure 4 shows occupational employment and wages, Figure 5 sectoral employment and prices, and Figure 6 occupational income shares within sectors. In all figures we show the evolution of the data in solid grey, contrasted with the model's predictions for the various counterfactual technology paths color coded as before: the baseline in solid red, the all factors in solid green, the neutral in pink dotted lines marked with x, the sector-only in dashed blue, and occupation-only in dashed-dotted yellow lines.

It is worth to note in Figures 4 to 6 is that our baseline model does very well in matching the data. It is important to recall that our baseline model matches the data exactly in the initial and final period by construction, but not in the interim periods. Nonetheless, even in the interim periods the differences with the data turn out to be small for almost all outcomes of interest, implying that in all periods the baseline model's predictions are extremely close to the data except for occupational relative wages, where they do not pick up the (short-lived) drop in the 1980 values.[24]

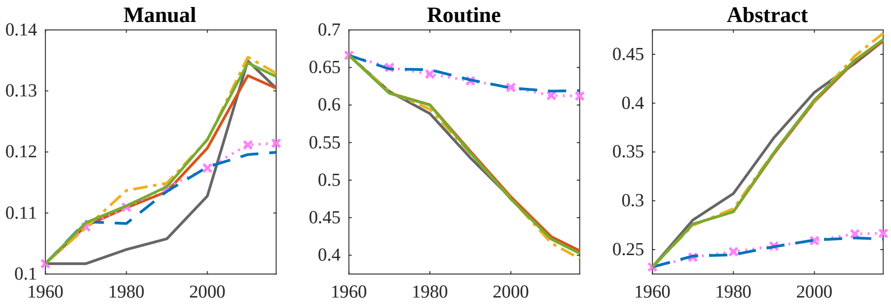

(a) Occupational employment

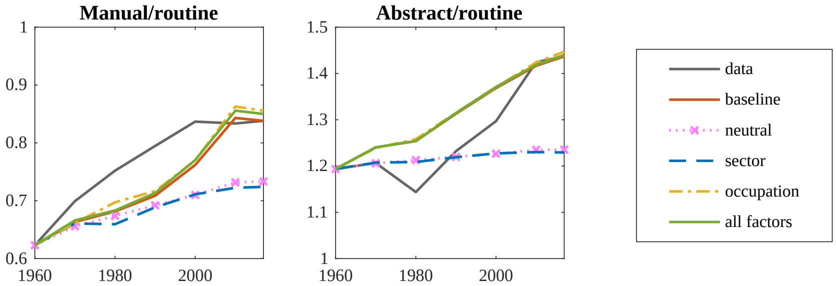

(b) Occupational relative wages

Figure 4: The evolution of occupational outcomes

We quantify for each counterfactual technology series how well the model predictions match the data in terms of the 1960-2017 change by averaging the following:

$$\left| \frac{\Delta x_k^m}{\Delta x_k^d} - 1 \right|, \tag{15}$$

where Δx_k^j is the change between 1960 and 2017 in our outcome of interest, in the data for $j = d$ and in the model for $j = m$. In other words, we report the average of the absolute value of the percent deviation of the change in the model from the change in the data. We average over $k \in \{m, r, a\}$ for occupational employment shares, over $k \in \{m, a\}$ for occupational wages relative to r, over $k \in \{L, G, H\}$ for sectoral

[24]Our model's failure to match these paths can be understood by looking at Figure 4b, where the dark grey solid line shows the data and the red solid line the values predicted by our baseline model. 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.

employment shares, over $k \in \{L, H\}$ for sectoral prices relative to G, and over $k \in \{m, r, a\} \times \{L, G, H\}$ for occupational income shares within sectors.

Table 3: Average percent deviation from 1960–2017 changes

	Outcomes				
	occupational		sectoral		cell
	empl.	rel. wages	empl.	rel. prices	θ
all factors	0.03	0.03	0.05	0.65	0.28
occ.	0.05	0.06	0.16	0.67	0.28
sector	0.69	0.69	0.42	0.83	1.00
neutral	0.65	0.66	0.25	0.91	1.00

Notes: The table shows the average of the absolute value of the percent deviation of the model from the 1960-2017 change in the data, as in (15), for 5 outcomes of interest in the columns: occupational employment shares and relative wages, sectoral employment shares and relative wages, and occupational income shares within sectors. Each row shows this model outcome for a different counterfactual technological change (based on all factors, occupation-only (occ.), sector-only (sector), and neutral).

Table 3 contains the resulting numbers, which show the percent by which on average the model is off compared to the data. As our baseline model perfectly matches the data in the initial and the final year, it gives a deviation of 0 (for each outcome and each k, i.e. on average as well), and hence we do not show the deviation for the baseline model in the table.

The first row shows that the model based on the all factors technology series does almost as well as the baseline model for occupational employment and wages and for sectoral employment, as the average deviation is at most 5 percent from the data. The only difference between this and the baseline model is that the former does not contain the change in technology that is idiosyncratic to the sector-occupation cell. The fact that these two models perform equally well for occupational employment and wages and for sectoral employment implies that the component of technological change idiosyncratic to the cell is not the key driver of these outcomes. However, for sectoral prices and for occupational income shares within sectors (θ_{oJ}), the discrepancy between the all factors model prediction and the data (and thus the baseline model) is on average 65 and 28 percent respectively. This highlights that technological change idiosyncratic to the sector-occupation cell plays a more important role in these outcomes.

Before analyzing in detail the predictions based on the different components separately, it is worth to point out what channels operate in our general equilibrium model.

The first channel is that relative technologies within a sector impact optimal relative occupational labor demand within the sector, as shown in (2). The second channel is that all technologies within a sector impact the sectoral prices, as in (3). The third channel is that changing technologies affect the stand-in household's income. Both changes in prices and in income impact the sectoral consumption demands through (8) and (9). In turn, this impacts how much employment needs to re-allocate across sectors. The price effect and the income effect together with the first channel determine by how much occupational labor demands and thus market-clearing occupational wages change. The occupation-only counterfactual technological change works through all of these channels. The sector-only technology path, which scales all α_{oJ} within sector J by the same factor, exerts a price and an income effect, but does not alter the relative labor demands within a sector. Neutral technological change works only through the income effect, as it does not directly alter relative sectoral prices or relative labor demands within a sector.

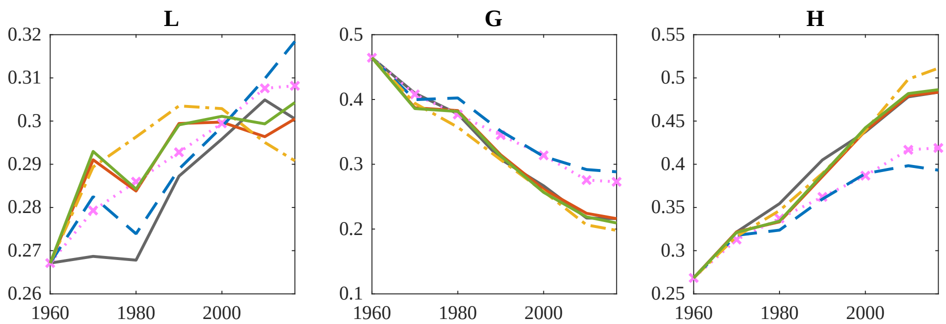

(a) Sectoral employment

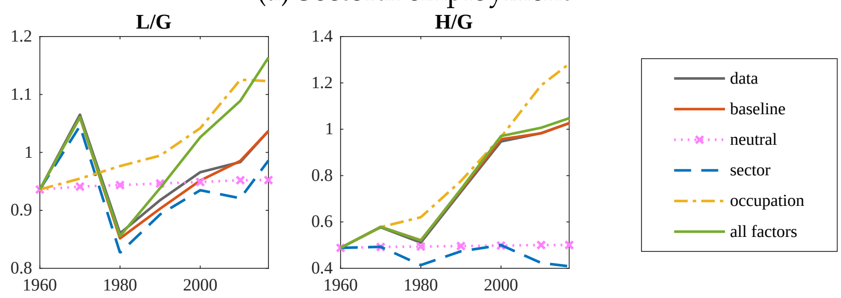

(b) Sectoral relative prices

Figure 5: The evolution of sectoral outcomes

Figure 4 shows (i) that the predictions based on occupation-only technological change are very close to the data for all occupational outcomes (with average deviations around 5 percent, see Table 3), and (ii) that while the predictions of sector-only and of neutral technological change are qualitatively in line with the data, quantitatively they fall short (with average deviations above 60 percent). Figure 5 shows (i) that for sectoral employment the predictions of occupation-only, of sector-only and of neutral technological change are close to the data (with average deviations of 16, 42 and 25 percent respectively), and (ii) that for sectoral prices neither the occupation-only nor the sector-only prediction does very well, and the neutral technology predicts virtually no changes (with average deviations of 67, 83 and 91 percent respectively).[25]

Figure 4 and 5a demonstrate that occupation-only, sector-only, and even neutral technological change by itself, generate occupational employment and wage, as well as sectoral employment paths qualitatively in line with the data. However, it is evident that the occupation-component plays a much larger role for labor market outcomes. That neutral technological change by itself moves employment and wage outcomes in the direction of the changes seen in the data, highlights the importance of the income effect. The non-homothetic terms in the utility function are large, as can be seen in Table 1, which explains why even neutral technological change leads to large reallocations across sectors, and in turn across occupations as well.[26]

The fact that neutral and sector-only technological change have almost identical implications in general equilibrium should not be surprising given that sectoral shocks are almost identical across sectors, as can be seen from Figure 3. Note, that this does not mean that the goods sector did not experience more rapid technological progress. Our sector components capture common trends within a sector once having accounted for the occupation components. In our data productivity growth was faster in the goods sector, but according to our decomposition this was mostly because it used

[25]Note that qualitatively the path for sectoral income shares in the data and in the various counterfactuals are very similar to those of sectoral employment shares.

[26]To quantify the role of non-homotheticities, we calibrated a homothetic version of the utility function (imposing $\bar{c}_S = 0$ for $S \in \{L, H\}$) to match sectoral employment shares in 1960. Comparing it to the baseline model we conclude that non-homotheticities explain in terms of sectoral employment about two thirds and in terms of occupational employment about a sixth of the changes in the data. Specifically, the homothetic model generates a 7.5 percentage point decrease in employment (out of 24.8 in the data) in the goods sector, and a 9.9 percentage point increase in the high-skilled service sector (out of 21.5 in the data), consequently predicting a small decrease in the low-skilled service sector (as opposed to a small increase in the data). In terms of occupational employment shares it predicts a 1.4 percentage point increase in manual employment (2.9 in the data), a 21.7 decrease in routine (26 in the data) and a 20.2 increase in abstract (23.2 in the data).

routine occupations more intensively.[27]

That occupation-only technological change implies effects very similar to the all factors prediction is not surprising, given that Figure 3 already established that occupation-only technological change mimics most of the evolution of sector-occupation cell technologies. As mentioned above, neutral and sector-only technological change generates dynamics consistent with the data as it induces shifts in consumption demands and thus in sectoral employment towards the service sectors, i.e. structural transformation, in line with the data. Since the high-skilled service sector is the most intensive in abstract and the low-skilled service sector in manual occupations, these sectoral shifts also lead to a decline in the relative demand for routine occupations, leading to the polarization of occupational employment and wages. However, as relative technologies across occupations within a sector have not changed, neither neutral nor sector-only technological change induces a decline of routine employment within a sector.[28] Therefore they both understate the overall changes in occupational outcomes. Moreover, they also somewhat understate the sectoral reallocations, reflecting the fact that some of the technological improvements within a sector are occupation-specific (as established by (13)).

Figure 6 shows the predicted changes in labor income shares within sectors. This figure shows that neutral or sector-only technological change predicts hardly any change in the θs. This can be understood from equation (16) in the appendix: labor income shares change if relative occupational wages or relative cell technologies within a sector change. The neutral and the sector-only predictions shut down the second channel, and predict quantitatively small changes in relative occupational wages, thus implying changes in the θs that are in line with the data, but which are quantitatively very small. The figure also reveals that in general the implications of the occupation-only and of the all factors technological change are virtually identical and are quite close to the data (with an average deviation of 28 percent). For manual occupations and in low-skilled services there are larger discrepancies which highlight that the component

[27] As we discussed earlier, here we do not account for capital accumulation or for capital-augmenting technological change. In Bárány and Siegel (2020) we account for both of these and find more pronounced sector components which resemble the patterns of sectoral labor productivity growth.

[28] In fact, as w_m/w_r and w_a/w_r increase, it leads to a small (and counterfactual) rise in the routine employment share of each sector.

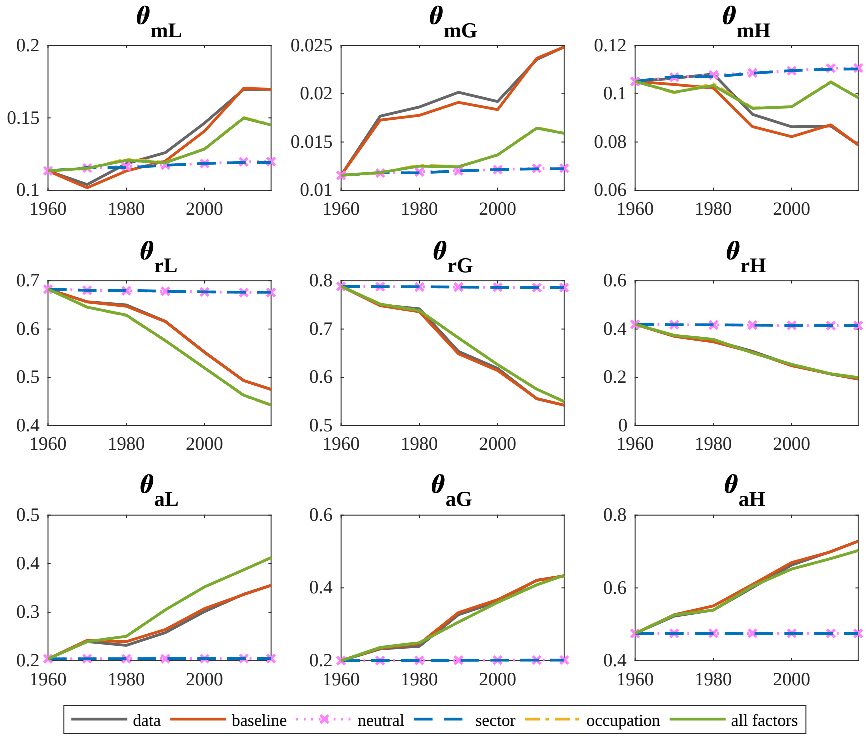

Figure 6: Income shares within sectors

of technological change idiosyncratic to the cell is important for the evolution of θ_{oJ}.[29]

5 Conclusion

In this paper we use a general equilibrium model to infer and study the consequences of biased technological change in labor market reallocations. Drawing on key firm side equations we infer sector-occupation-specific technologies from the data, which we then decompose into neutral, sector- and occupation-specific components, as well as sector-occupation cell specific residuals. This decomposition shows that neutral, sector- and occupation-specific components jointly explain about 90 percent of the variation in cell level technological change. We evaluate the role of these components by feeding these as counterfactual technological paths into the model. We find that qualitatively any counterfactual technological change series would generate structural transformation, the observed sectoral reallocation of employment, as well as polar-

[29]Note, the behavior of the sector-occupation cell employment shares in response to the evolution of the various components of technological change is very similar to Figure 6, and not shown for brevity.

ization of occupational employment and wages. This suggests an important role for income effects working through non-homotheticities in preferences. However, quantitatively we find a major role for occupation-bias in technological change as the driver of these reallocations. Moreover we find that occupation components and cell-specific elements are important drivers of occupational income shares within sectors. To explain the evolution of sectoral prices over time both sector and occupation components are needed.

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 that help workers switch occupations.

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.

Daron Acemoglu and Pascual Restrepo. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30, Spring 2019.

Pol Antràs. Is the U.S. aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution. *Contributions to Macroeconomics*, 4(1), 2004.

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

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):189–194, 2006.

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

Zsófia L. Bárány and Christian Siegel. Engines of Sectoral Labor Productivity Growth. *Review of Economic Dynamics*, forthcoming, 2020.

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

Timo Boppart. Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences. *Econometrica*, 82(6):2167–2196, 2014.

Francisco J. Buera, Joseph P. Kaboski, Richard Rogerson, and Juan I. Vizcaino. Skill Biased Structural Change. Working paper, 2018.

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

Diego A. Comin, Danial Lashkari, and Martí Mestieri. Structural Change with Long-run Income and Price Effects. *Econometrica*, forthcoming, 2020.

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.

Gene M. Grossman and Esteban Rossi-Hansberg. Trading Tasks: A Simple Theory of Offshoring. *American Economic Review*, 98(5):1978–97, 2008.

Berthold Herrendorf, Christopher Herrington, and Ákos Valentinyi. Sectoral Technology and Structural Transformation. *American Economic Journal: Macroeconomics*, 7(4): 104–33, October 2015.

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.

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.

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

Francesca Mazzolari and Giuseppe Ragusa. Spillovers from High-skill Consumption to Low-skill Labor Markets. *The Review of Economics and Statistics*, 95(1):74 – 86, March 2013.

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

The classification of workers into occupational categories and into industrial sectors is identical to the assignments in Bárány and Siegel (2020). For ease of reference, these mappings are listed below.

We combine three different industry classification systems, the NAICS, the SIC, and the ind1990. Table A1 summarizes our categorization in terms of each system.

Table A1: Classification of industries into sectors

	NAICS	SIC	ind1990
Low-skilled services	• Wholesale trade • Retail trade • Transportation & warehousing • Arts, entertainment, recreation, accommodation & food serv. • Other serv., except government	• Wholesale trade • Retail trade • Transportation • Amusement & recreation serv. • Motion pictures • Hotels & other lodging places • Personal serv. • Auto repair, serv. & parking • Miscellaneous repair serv. • Private households	• Wholesale & retail trade • Transportation • Entertainment • Low-skilled business serv. • Personal serv.
Goods	• Agriculture, forestry, fishing & hunting • Mining • Construction • Manufacturing	• Agriculture, forestry, & fishing • Mining • Construction • Manufacturing	• Agriculture, forestry & fishing • Mining • Construction • Manufacturing
High-skilled services	• Utilities • Information • Finance, insurance, real estate, rental & leasing • Professional & business serv. • Educational serv., health care & social assistance • Government	• Electric, gas, & sanitary serv. • Communications • Finance, insurance, & real estate • Legal serv. • Business serv. • Miscellaneous professional serv. • Membership organizations • Educational serv. • Health serv. • Social serv. • Government	• Utilities • Communications • Finance, insurance & real estate • Professional serv. High-skilled business serv. • Public administration

We classify occupations based on their routine task content and cognitive requirements, similarly to Acemoglu and Autor (2011), into the following three categories: **Manual** (low-skilled non-routine): housekeeping, cleaning, protective service, food preparation and service, building, grounds cleaning, maintenance, personal appearance, recreation and hospitality, child care workers, personal care, service, healthcare support;

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

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

A.2 Derivations

In this subsection we show how the αs can be 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_o/w_r to get:

$$\frac{\theta_{oJ}}{\theta_{rJ}} = \left(\frac{w_r}{w_o}\right)^{\eta-1} \left(\frac{\alpha_{oJ}}{\alpha_{rJ}}\right)^{\eta-1} \quad \text{for } o \in \{m, a\}. \tag{16}$$

Re-arrange to get:

$$\frac{\alpha_{oJ}}{\alpha_{rJ}} = \left(\frac{\theta_{oJ}}{\theta_{rJ}}\right)^{\frac{1}{\eta-1}} \frac{w_o}{w_r} \quad \text{for } o \in \{m, a\}.$$

The relative prices across sectors pin down the relative αs across sectors. To see this, use (3) 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 expression 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 technologies over time conditional on the sectoral income shares. The sectoral income shares, using equation (4), can be expressed as:

$$\frac{\Psi_J}{\Psi_K} = \frac{p_J Y_J}{p_K Y_K} = \frac{l_{mJ} p_J^{1-\eta} w_m^{\eta} \alpha_{mJ}^{1-\eta}}{l_{mK} p_K^{1-\eta} w_m^{\eta} \alpha_{mK}^{1-\eta}}.$$

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

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

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

$$\begin{aligned}
Y_L &= \left[(\alpha_{mL} l_{mL})^{\frac{\eta-1}{\eta}} + (\alpha_{rL} l_{rL})^{\frac{\eta-1}{\eta}} + (\alpha_{aL} l_{aL})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \\
&= \alpha_{mL} l_{mL} \left[1 + \left(\frac{\alpha_{rL} l_{rL}}{\alpha_{mL} l_{mL}}\right)^{\frac{\eta-1}{\eta}} + \left(\frac{\alpha_{aL} l_{aL}}{\alpha_{mL} l_{mL}}\right)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} = \alpha_{mL} l_{mL} \left(\frac{1}{\theta_{mL}}\right)^{\frac{\eta}{\eta-1}} \\
&= \alpha_{mH} \frac{p_H}{p_L} \left(\frac{\theta_{mH}}{\theta_{mL}}\right)^{\frac{1}{1-\eta}} \left(\frac{1}{\theta_{mL}}\right)^{\frac{\eta}{\eta-1}} l_{mH} \frac{\Psi_L}{\Psi_H} \frac{\theta_{mL}}{\theta_{mH}} = \alpha_{mH} l_{mH} \theta_{mH}^{\frac{\eta}{1-\eta}} \frac{p_H}{p_L} \frac{\Psi_L}{\Psi_H}, \\
Y_G &= \alpha_{mH} l_{mH} \theta_{mH}^{\frac{\eta}{1-\eta}} \frac{p_H}{p_G} \frac{\Psi_G}{\Psi_H}, \\
Y_H &= \alpha_{mH} l_{mH} \theta_{mH}^{\frac{\eta}{1-\eta}}.
\end{aligned}$$

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:

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

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 targets

Table A2: Calibration targets

	1960	1970	1980	1990	2000	2010	2017
p_L/p_G	0.9360	1.0651	0.8604	0.9177	0.9658	0.9834	1.0369
p_H/p_G	0.4882	0.5769	0.5110	0.7301	0.9472	0.9831	1.0261
Ψ_L	0.2563	0.2606	0.2604	0.2747	0.2794	0.2805	0.2751
Ψ_G	0.4734	0.4121	0.3800	0.3068	0.2624	0.2120	0.2095
Ψ_H	0.2703	0.3274	0.3596	0.4186	0.4582	0.5075	0.5155
growth	1	1.2009	1.3410	1.5482	1.8230	2.1971	2.2489
w_m/w_r	0.6231	0.6993	0.7520	0.7947	0.8368	0.8335	0.8382
w_a/w_r	1.1936	1.2066	1.1438	1.2320	1.2969	1.4243	1.4370
θ_{mL}	0.1134	0.104	0.1183	0.126	0.1465	0.1697	0.1698
θ_{rL}	0.6827	0.6565	0.65	0.616	0.5522	0.4931	0.4748
θ_{aL}	0.2039	0.2395	0.2317	0.258	0.3013	0.3371	0.3554
θ_{mG}	0.0116	0.0177	0.0186	0.0202	0.0192	0.0236	0.0249
θ_{rG}	0.7887	0.7501	0.7419	0.6535	0.6184	0.5554	0.5421
θ_{aG}	0.1998	0.2323	0.2395	0.3264	0.3624	0.4211	0.4330
θ_{mH}	0.1051	0.1064	0.1082	0.0915	0.0863	0.0866	0.0788
θ_{rH}	0.4197	0.3709	0.3526	0.3075	0.2508	0.2134	0.1927
θ_{aH}	0.4752	0.5226	0.5392	0.6010	0.6629	0.7000	0.7284

Table A2 contains the targets used in the calibration. Based on these, occupational employment shares are inferred from income shares and wages per efficiency unit as described in section 3, and sector-occupation employment shares follow. Figure A1 shows these, and it is the analogue to Figure 1, which plots hours worked shares from the data.

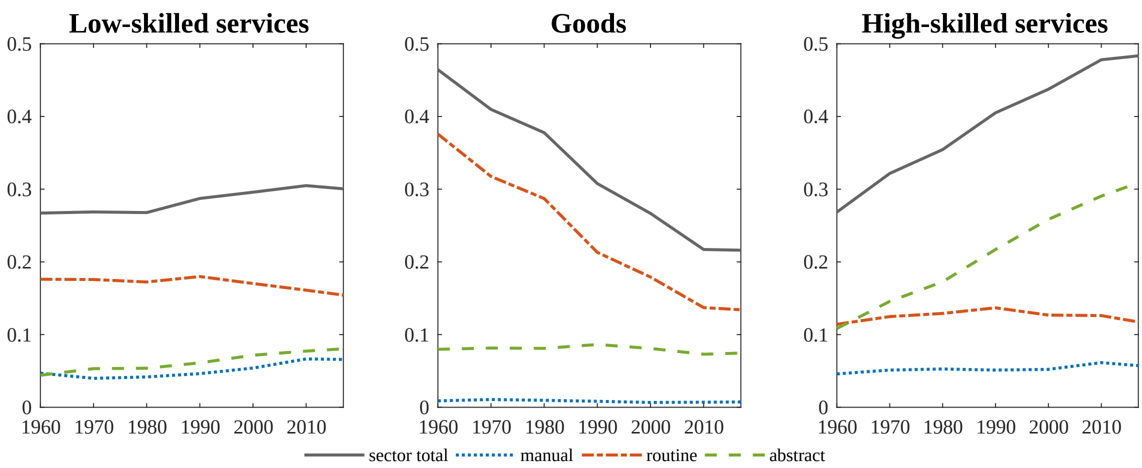

Figure A1: Implied employment shares 1960-2017

A.4 Finer industry and occupation classifications

We separate each of our three broad sectors into four industries each: We split low-skilled services into wholesale trade, retail trade, transport, and the remaining low-skilled service industries; the goods sector into agriculture, mining, construction, and manufacturing; and high-skilled services into utilities, FIRE, government services, and the remaining high-skilled services. Similarly, we break up the three broad occupational groups into a total of ten occupation categories. We divide manual occupations into personal care, food and cleaning services, and protective services; routine occupations into operators and laborers (incl. agricultural), production, office and admin., and sales; and abstract occupations into technicians, professionals, and managers.

Here, similar to Figure 1, we explore the connection between sectoral and occupational employment patterns for these finer classifications of industries and occupations. Figure A2 shows the allocation of hours worked across cells defined by the 12 industries and the 3 or 10 occupations. The top row shows the low-skilled service industries, the middle row the industries in the goods sector, and the third row shows the high-skilled service industries. The color coding of the occupations in the top panel is the same as in Figure 1, blue is manual, red is routine and green is abstract. In the bottom panel we maintain a similar color coding: the blueish lines show occupations within the manual group, the reddish are occupations in the routine group, and the greenish lines are occupations within the abstract group. First, we see that not all industries follow in terms of employment the same path as the broader sector that they are in: some low-skilled service industries (wholesale and transport) shrank, construction in the goods sector was largely stable, and utilities (within the high-skilled services) shrank. We further see that the contraction in all the goods sector industries happened through routine occupations (albeit not the same occupation within the routine group), the expansion in high-skilled service industries happened through abstract occupations, and in low-skilled service industries it happened mainly through manual occupations. While the connection between industrial and occupational employment in this finer classification is not as tight, the link between reallocations across these two dimensions is still clear.

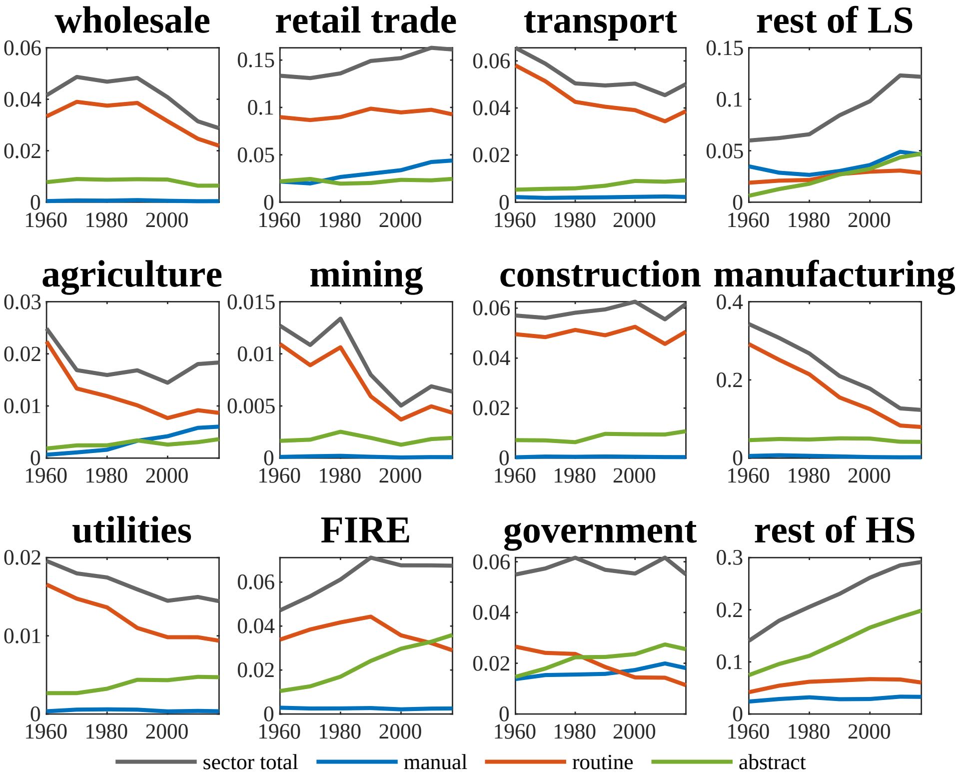

(a) 3 occupations

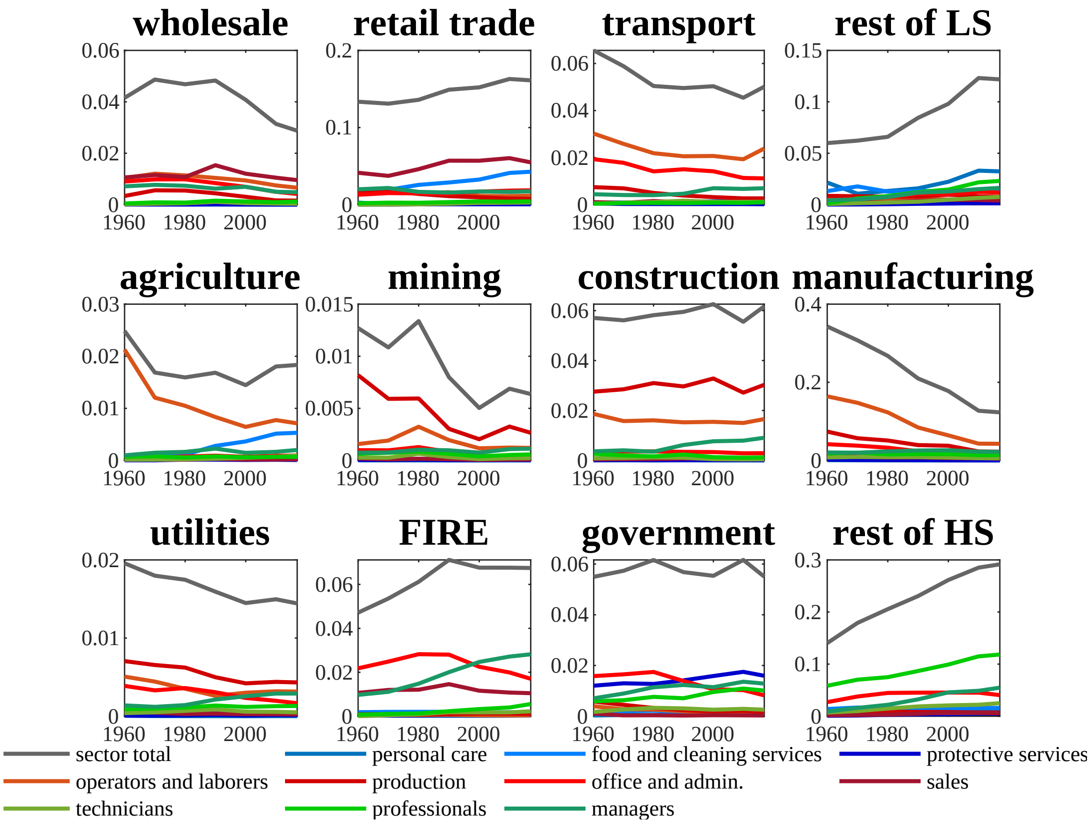

(b) 10 occupations

Figure A2: Allocation of hours worked across 12 industries and 3/10 occupations

A.5 Robustness checks

Here we explore the robustness of our results to alternative values of the elasticity of substitution in production, η, in consumption, ε, and to alternative correlation parameters in the distribution of in the occupational cost differences, ρ. In each of these robustness checks we recalibrate the model based on the alternative values of the fixed parameters we consider. Note that from these three parameters, only the value of η matters for the sector-occupation technologies that we infer from the data as these are fully pinned down by the production side of the model. However, all three parameters impact the general equilibrium outcomes of the model.

Table A3 summarizes the impact of changing each of these three parameters, one by one, on our results, using the same as in Table 3. The numbers show the percent by which on average the change in the model is off compared to the data. Within each panel each row represents a different counterfactual technological change (based on all factors, occupation-only (occ.), sector-only (sector), and neutral). Each column shows for an outcome of interest the average deviation of the model, as shown in Figures 4 – 6. The top panel repeats Table 3 by showing the statistics based on the baseline parameters. The subsequent panels show the same for alternative parameterizations. This table shows that all our conclusions are maintained qualitatively, and are quite robust quantitatively across the different parametrizations.

Table A3: Average deviation from 1960–2017 changes under alternative parameters

		Outcomes				
		occupational		sectoral		cell
		empl.	rel. wages	empl.	rel. prices	θ
Baseline:						
$\eta = 0.6$,	**all factors**	**0.03**	**0.03**	**0.05**	**0.65**	**0.28**
$\varepsilon = 0.2$,	**occ.**	**0.05**	**0.06**	**0.16**	**0.67**	**0.28**
$\rho = 0$	**sector**	**0.69**	**0.69**	**0.42**	**0.83**	**1.00**
	neutral	**0.65**	**0.66**	**0.25**	**0.91**	**1.00**
Alternative η:						
$\eta = 0.5$	all factors	0.03	0.03	0.03	0.51	0.28
	occ.	0.03	0.03	0.06	0.25	0.28
	sector	0.64	0.64	0.31	0.67	0.99
	neutral	0.64	0.64	0.25	0.90	0.99
$\eta = 0.7$	all factors	0.03	0.03	0.09	0.88	0.28
	occ.	0.09	0.10	0.33	1.47	0.28
	sector	0.76	0.76	0.57	1.05	1.00
	neutral	0.67	0.67	0.25	0.91	1.00
Alternative ε:						
$\varepsilon = 0.1$	all factors	0.03	0.03	0.10	0.65	0.28
	occ.	0.06	0.07	0.27	0.67	0.28
	sector	0.74	0.74	0.96	0.84	1.00
	neutral	0.69	0.70	0.65	0.92	1.00
$\varepsilon = 0.3$	all factors	0.03	0.03	0.05	0.65	0.28
	occ.	0.04	0.05	0.05	0.66	0.28
	sector	0.61	0.62	0.53	0.81	0.99
	neutral	0.60	0.60	0.43	0.90	0.99
Alternative ρ:						
$\rho = 0.3$	all factors	0.03	0.03	0.05	0.65	0.28
	occ.	0.05	0.06	0.16	0.67	0.28
	sector	0.69	0.68	0.42	0.82	1.00
	neutral	0.66	0.64	0.25	0.90	1.00
$\rho = -0.3$	all factors	0.03	0.03	0.05	0.65	0.28
	occ.	0.05	0.06	0.16	0.66	0.28
	sector	0.68	0.70	0.42	0.83	1.00
	neutral	0.64	0.67	0.25	0.91	0.99

Notes: The table shows the average of the absolute value of the percent deviation of the model from the 1960-2017 change in the data, as in (15), for 5 outcomes of interest in the columns, and for the baseline and 6 other parametrizations in the panels. In each panel each row shows this model outcome for a different counterfactual technological change (based on all factors, occupation-only (occ.), sector-only (sector), and neutral). The 5 outcomes are: occupational employment shares and relative wages, sectoral employment shares and relative wages, and occupational income shares within sectors.