Engines of Sectoral Labor Productivity Growth*

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

July 2020

Abstract

We study the origins of labor productivity growth and its differences across sectors. In our model, sectors employ workers of different occupations and various forms of capital, none of which are perfect substitutes, and technology evolves at the sector-factor cell level. Using the model we infer technologies from US data over 1960-2017. We find that sectoral differences in labor productivity growth are largely due to sectoral differences in the growth rate of routine labor augmenting technologies. Neither capital accumulation nor the occupational employment structure within sectors explains much of the sectoral differences in labor productivity growth.

Keywords: biased technological change, structural transformation, labor productivity

JEL codes: O41, O33, J24

*We wish to thank Francesco Caselli, Georg Duernecker, Tim Lee, Miguel León-Ledesma, Guy Michaels, Rachel Ngai, Mathan Satchi, Ákos Valentinyi as well as seminar and conference participants. We gratefully acknowledge funding from a British Academy/Leverhulme Small Research Grant (SC-07334) and thank God'stime Eigbiremolen for research assistance.

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

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

1 Introduction

The fact that labor productivity growth is different across sectors is well known. Average annual labor productivity growth between 1960 and 2017 in the US, for instance, was 2.5% in the goods sector, much higher than the 1.5% in low-skilled and the 0.7% in high-skilled services. However, it is not clear what the origins of these differences are. Several explanations are possible, such as differences in technological progress across sectors or across production factors, or differential trends in capital deepening or in the use of other inputs. We study the drivers of sectoral labor productivity growth in a production-side framework. What sets our framework apart from the literature is that (i) we consider various types of occupational labor as distinct production factors, (ii) technological change is sector-and-factor specific, and (iii) we infer the evolution of the sector-and-factor specific technologies over time directly from the data. Essentially, we perform model-based sectoral growth accounting.

In our model we consider different occupations as distinct production factors for a variety of reasons.[1] First, given that occupations entail very different tasks, they are most likely not perfect substitutes. This implies that using the simple summation of hours worked within a sector might not capture labor's true contribution to a sector's output. The second reason is that occupations are likely to use different technologies, which might grow at different rates. Third, the effects of new technologies and of the accumulation of (different types of) capital on the various occupations might depend on the tasks performed by that occupation, in particular on their routine content and cognitive requirements. In our analysis we therefore differentiate between manual, routine and abstract labor. We also distinguish between Information and Communication Technology (ICT) capital and non-ICT capital, allowing for different substitutability with the various types of occupational labor. Altogether we thus have five factors of production: three types of occupational labor, and two types of capital.

The second key feature of our framework is that we allow technologies to evolve in a very flexible way, at the sector-factor level. It is common to think about sector-

[1]Our analysis is based on the firm side of the economy and on a production function where firms hire workers into distinct occupations. We therefore think of different occupational labor inputs as distinct production factors, irrespective of the fact that workers might be able to choose their occupation.

specific technological change and factor-augmenting technologies. However, on the one hand, technological change can occur differentially within sectors across factors of production, even within labor and within capital factors. For example in the automobile manufacturing industry various tasks performed by routine workers (on the production line) have been largely automated, without a direct impact on the productivity of manual workers (e.g. cleaners and janitors) or of abstract workers (e.g. accountants or managers). Thus technologies augmenting the different occupations in this industry did not evolve in the same way. On the other hand, technological change can occur differentially across sectors for a given factor. Production line workers, as we discussed above, have seen large changes in technologies, with many tasks automated. Workers in material moving occupations in the shipping industry used to lift heavy objects by hand or by hand-operated equipment. Now these occupations require operating cranes and forklifts to move containers around. Sales workers in real estate nowadays can communicate more information and more easily with clients, but still need to do viewings in person. These examples show how different the evolution of technology of routine occupations can be across sectors. To accommodate for all of these possible cases, we allow technological change to be sector-and-factor specific. Note as we infer from the data these sector-and-factor specific technologies, the data could suggest any pattern. For example technology could improve at the same rate either for all factors within a sector (i.e. sector-specific productivity improvements), or for a given factor across all sectors (i.e. factor-biased technological change).

In order to infer factor-augmenting technologies we need to make assumptions about the structure of production.[2] Similarly to Katz and Murphy (1992) and Krusell, Ohanian, Ríos-Rull, and Violante (2000) we assume a (nested) CES production function.[3] Taking values for the substitution elasticities from the literature, we infer sector-specific factor-augmenting technologies in each period from firm optimality conditions. This approach is similar to Caselli (2005), Caselli and Coleman (2006), and

[2]Observing factor inputs and output allows the computation of a neutral productivity. This is how total factor productivity (for instance at the sectoral level) is extracted; note that changes in measured TFP might actually be driven by technological change augmenting only one individual factor of production.

[3]These papers impose a specific process for factor-biased technological change, allowing them to estimate the elasticity of substitution. In contrast, we do not impose any restrictions on technological change.

Buera, Kaboski, Rogerson, and Vizcaino (2018). In order to infer technologies we need data on the occupations of workers within each sector, which we take from the US Census and American Community Survey (ACS). Further, we combine data from the U.S. Bureau of Economic Analysis (BEA) and EU KLEMS 2017 to obtain value added, labor income shares, prices, employment and capital of different types by sector.

Our results show that allowing for technologies to be sector-and-factor specific is crucial. Technologies have evolved at very differential rates, both across factors within each sector and across sectors for a given factor (occupation or type of capital), echoing the general conclusions of Caselli (2016).[4] In particular, technologies augmenting routine occupations have been growing the fastest in all sectors, but at very different rates across sectors: at 5.6% per year in goods, at 2.9% in low-skilled services and at 1.3% in high-skilled services.

Through a series of counterfactual simulations, we study the role of technological change and of inputs in labor productivity growth. We find that the single most important driver of sectoral labor productivity growth differences is sector-specific routine labor augmenting technological change. Without this type of technological change, labor productivity growth would have been almost equalized across sectors. Specifically, sector-specific routine labor augmenting technological change explains at least 59 percent of labor productivity growth in low-skilled services, 74 percent in goods and 21 percent in high-skilled services.

As we found such an important role for sector-specific routine labor augmenting technological change, we want to assess to what extent this can be assigned to overall routine labor augmenting technology growth or to sectoral components in technological progress. Therefore, we conduct a factor model decomposition of the growth rates of all sector-specific occupational labor-augmenting technologies (or sector-occupation specific technologies for short). With this, we establish that these growth rates are well described as the sum of sector-specific and occupation-specific components, and both types of components are required.[5] Moreover, we show through counterfactuals that

[4]While Caselli investigates technological biases across labor and capital, and across workers of different education or experience, we consider biases across different factors of production (including occupations and different types of capital).

[5]While these components of technological change have important equilibrium implications for the labor market, the analysis of these effects is beyond the scope of this paper. In Bárány and Siegel (2019)

the occupation-specific components do not explain sectoral differences in labor productivity growth. This suggests that sector heterogeneity in routine labor augmenting technological change is the key driver of sectoral differences in labor productivity growth rates.

Further counterfactuals allow us to evaluate the role of various other channels proposed in the literature for sectoral productivity growth differences. As suggested by Acemoglu and Guerrieri (2008), differential capital intensities and capital accumulation could be driving the faster productivity growth in the goods sector. While we find that capital accumulation contributes to labor productivity growth (without it growth would have been 39 percent lower on average), it does not generate the sectoral differences in labor productivity growth observed in the data. Instead, we confirm the finding of Herrendorf, Herrington, and Valentinyi (2015), that differences in labor-augmenting technological progress across sectors are crucial. They do not differentiate labor by occupation, and as such cannot identify the sources of the sectoral differences in labor-augmenting technological change. In principle, these could be driven by differences in sectoral intensities of occupational employment and technological change specific to occupations, as suggested by Duernecker and Herrendorf (2016) and by Lee and Shin (2017). However, in our framework we find that differences in occupational employment structure across sectors do not contribute much to sectoral labor productivity growth differences. Instead our results indicate that the sectoral differences in the growth rate of routine labor augmenting technologies themselves are crucial.

We conduct our analysis at the sectoral level, and it also has implications for aggregate productivity growth, as emphasized in Duarte and Restuccia (2010), Duernecker, Herrendorf, and Valentinyi (2017), and Duarte and Restuccia (2019). We find that in our context, i.e. taking into account five factors of production and their technologies, technological change is much more important than input use for labor productivity growth also in the aggregate. A novel finding of our paper is that the contribution of routine labor augmenting technological change is large and increasing over time. In its absence aggregate growth would have been lower by about a third between 1960-1990, and there would have been hardly any growth over 1990-2017.

---

we use a simplified production function to study these effects in general equilibrium.

The paper proceeds as follows: section 2 shows the facts about sectoral production on which we base our analysis. Section 3 introduces the production-side framework used to infer technologies and explains its implementation. In section 4 we analyze the role of inputs and technologies in labor productivity growth through counterfactuals. In section 5 we verify that our results are very robust to alternative substitution elasticities, to a different nesting of the production function, to controlling for workers' human capital, and to different treatments of the data. The final section concludes.

2 Factor use and factor income shares by sector

In this section we describe the data used to inform our specification of the sectoral production functions. From this data we will also infer – using our model's optimality conditions – the evolution of the sector-specific factor augmenting technologies. We combine data from the U.S. Bureau of Economic Analysis (BEA) on sectoral value added and its components, on sectoral prices, on sectoral employment, and on fixed assets, with data on the allocation of capital across sectors from EU KLEMS 2017. To get more detailed information on the occupations of workers within each sector, 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). Since we draw on various data sources which are based on different industry classification systems, we map the fine industries of each system into our broad sector categories, as explained in detail in Table A1 in the appendix.

We use annual data on nominal value added, real value added and prices by industry from the BEA.[6] We group all non-service industries into the goods sector, and similarly to much of the recent literature on structural transformation, we break services into two, based on the skill or education level of workers in the industry.[7] It is common to split services, as already in 1947 the service industries as a whole constituted

[6]The industry categories in this dataset are based on the North American Industry Classification System (NAICS)).

[7]Services are split based on whether they are high- or low-skilled in Buera and Kaboski (2012), whether they are low- or high-productivity growth in Duernecker et al. (2017), or whether they are traditional/modern services in Duarte and Restuccia (2019). While these splits are based on different criteria, in practice the overlap between such classifications is substantial. In appendix E.5 we also report results when dividing the goods sector into agriculture and industry.

around 60 percent of total value added. We aggregate real value added and price data on fine industry categories into our three broad sectors – low-skilled services, goods, and high-skilled services – using the cyclical expansion procedure, as for example in Herrendorf, Rogerson, and Valentinyi (2013). The left panel of Figure 1 shows the

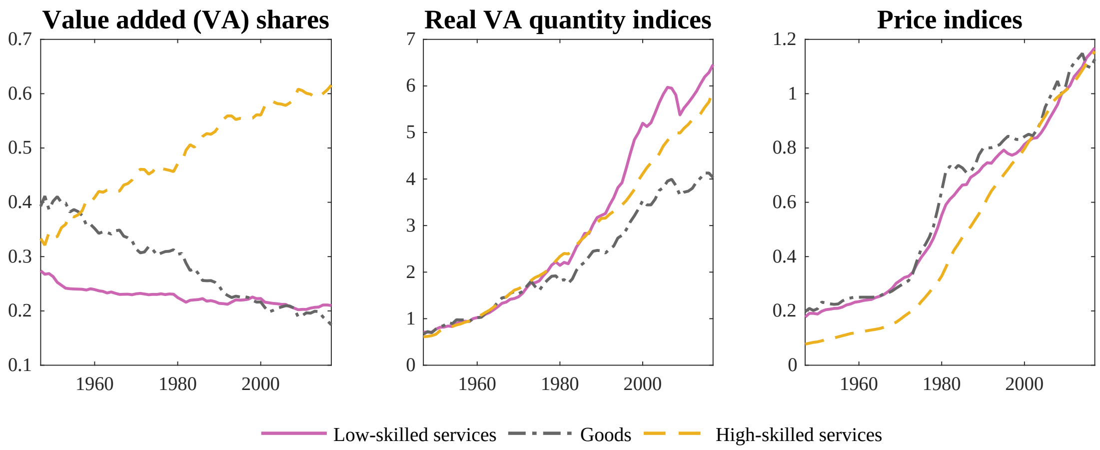

Figure 1: Nominal value added, real value added and prices

Notes: Authors' own calculations based on Value Added by Industry data from the BEA for the years 1947-2017. The left panel shows the evolution of nominal value added shares across sectors (summing to 1 in each period). The middle panel shows the evolution of real value added quantity indices by sector (these are normalized to be 1 in 1960 on the graph). The right panel shows the evolution of sectoral price indices (normalized to 1 in 2009).

evolution of (nominal) value added shares, which displays structural transformation: the share of value added produced in high-skilled services increased steadily from the 1940s, the share produced in goods steeply declined, and in low-skilled services it also declined albeit at a lower rate. The evolution of real value added by sector (depicted in the middle panel) together with the evolution of sectoral employment gives us sectoral labor productivity growth. The steady increase in the nominal value added share in high-skilled services can be reconciled with its lower growth in real terms vis-a-vis low-skilled services by the steep increase in the relative price of high-skilled services, as shown in the right panel.

We next investigate the use of various factor inputs and their income shares in each sector. As a first step, we calculate the share of sectoral income paid to capital (Θ_J) and to labor ($1 - \Theta_J$), using data on the Components of Value Added by Industry from the

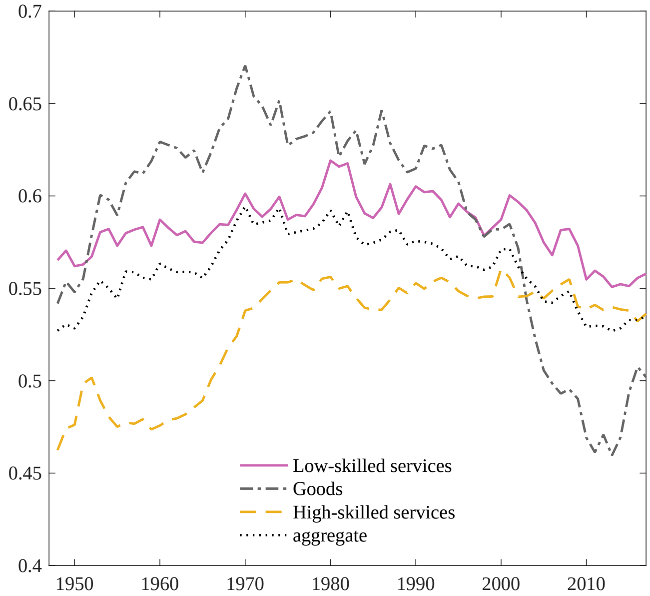

Figure 2: Labor income share by sector

Notes: Compensation of employees relative to gross value added in a sector calculated from Components of Value Added by Industry data provided by the BEA for 1947-2017.

BEA. We calculate the labor income share as:[8]

$$1 - \Theta_J = \frac{\text{Compensation of employees in sector } J}{\text{Gross value added in sector } J}.$$

The difficulty is that for the period before 1987 this data is only available based on the Standard Industrial Classification (SIC), whereas for the period post 1997 it is only available based on the NAICS classification of industries. Therefore we have to combine these two data sources based on different industry classification systems. While the individual industries are not the same in these two classifications, when we aggregate them up to our three broad sectors, the two give similar results for the period of the overlap. As the NAICS data was introduced in 1997, we use the (native) SIC data until 1997, and the NAICS data from that point onwards.[9] Figure 2 plots the evolution of the labor income share by sector as well as for the aggregate economy. The

[8]This definition of the labor income share excludes proprietors' income. We choose to do this for two reasons. First, Elsby, Hobijn, and Şahin (2013) call this the *unambiguous part* of the labor income share. Second, we take data on workers from the Census and the ACS, and there we only include employees, which makes this definition of labor income share consistent with our approach there. However, in an extension we include the self-employed in our analysis. While there are some differences in the sectoral labor income shares, our main results still carry through, see appendix E.4 for details.

[9]Herrendorf et al. (2015) also combine data on the labor income and employment shares across different industries based on the SIC and the NAICS classification.

labor income share in the economy as a whole increased until the early 1970s, which was followed by a virtually equal reduction thereafter.[10] There are two important observations. First, there are substantial sectoral differences in the labor income share. For most of the period between 1947 and 2017 the goods sector had the highest labor income share, while high-skilled services had the lowest labor income share. The second thing to note is that these labor income shares are far from constant: following a common increasing trend until the 1970s, the labor income share declined steeply in the goods sector, declined slightly in low-skilled services, whereas it stayed roughly constant in high-skilled services. Thus to be able to replicate these patterns, we need sectoral production functions which allow the labor income share to change over time, e.g. not of the Cobb-Douglas form.

We next analyze the use of capital. In our analysis we distinguish between two types of capital, ICT and non-ICT, as discussed in the introduction. The BEA Fixed Asset Accounts contains annual data on the nominal stock and on chain-type quantity indices of various types of capital for the entire period of our analysis. When constructing ICT capital from the BEA we include *Information processing equipment* and *Software*, while non-ICT capital comprises of all other non-residential capital.[11] Starting from data on these finer categories of capital we calculate quantity and price indices for our two aggregates using the cyclical expansion procedure. Figure 3 in the left and middle panel shows the evolution of the real quantity and price of ICT and non-ICT capital in the US economy between 1960 and 2017. The left panel shows that ICT capital grew much faster over this period than non-ICT capital. The huge improvement in ICT technology is reflected in the steep fall of ICT prices from the 1980s and the steep increase in ICT capital from the 1990s onwards.

In order to measure the allocation of ICT capital across sectors we use data from EU KLEMS 2017. The EU KLEMS 2017 release contains annual data on various types

[10]When comparing this series with the widely noted decline in the labor income share (Elsby et al. (2013) and Karabarbounis, Loukas and Neiman, Brent (2014) for example), it is important to bear in mind that we exclude proprietors' income from labor income. Since proprietors' income has been falling throughout this period, and especially until the 1970s, it roughly offsets the increase in the aggregate labor income share until the 1970s, and makes the subsequent decline slightly more pronounced.

[11]Non-ICT capital consist of *Industrial equipment*, *Transportation equipment*, *Other equipment*, *Nonresidential structures*, *Research and development* and *Entertainment, literary, and artistic originals*, as well as all non-residential government fixed assets except for *Software*, which is included in ICT capital.

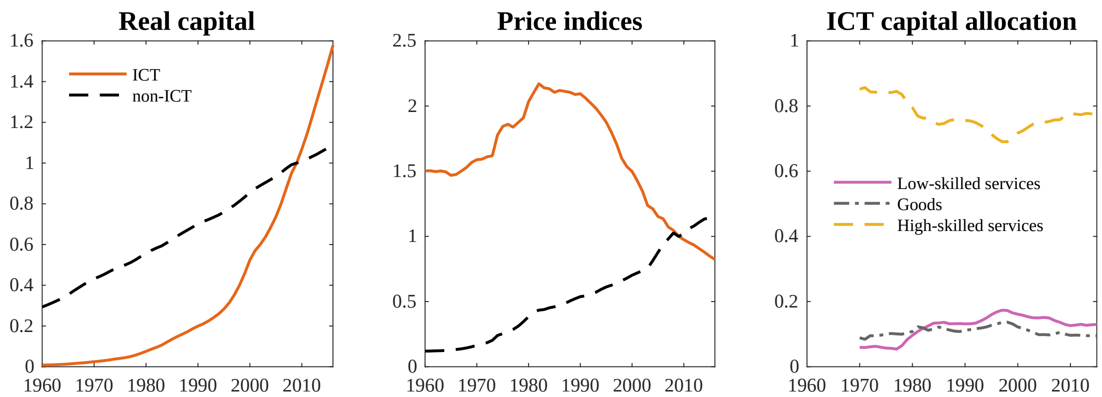

Figure 3: Real quantity and price of ICT and non-ICT capital, and allocation of ICT capital across sectors

Notes: The left panel shows real quantity indices, while the middle panel shows the price indices of ICT and non-ICT capital (all indices are normalized to 1 in 2009), both are computed based on data from the BEA Fixed Asset Accounts. The right panel shows the evolution of the share of ICT capital across sectors (summing to 1), calculated from EU KLEMS data.

of capital by industry (based on the International Standard Industrial Classification of All Economic Activities (ISIC)) from 1970 onwards. When constructing the allocation of ICT capital across sectors from the EU KLEMS data we include the following categories: *Computing equipment*, *Communications equipment*, and *Computer software and databases*. The right panel in Figure 3 shows the fraction of nominal ICT capital stock in each sector, and shows that most of the ICT capital stock is in the high-skilled service sector, with a roughly equal quantity in low-skilled services and goods. Note that data on the allocation of ICT capital across sectors is only available between 1970 and 2015. To infer technologies from the data, as detailed in the next section, we impute values for this allocation in 1960 and in 2017. Since the allocation across sectors seems quite flat between 1970 and 1978 and between 2010 and 2015, we impose the 1970 values for 1960, and the 2015 values for 2017.

Finally, we break down employment and labor income within each sector by occupation. As discussed in the introduction, we believe that in order to understand what is driving sectoral labor productivity growth it is crucial to differentiate between occupations. Since the national accounts do not contain any information on the occupation of workers within industries, we turn to the decennial US Census and ACS data between 1960 and 2017 from IPUMS, provided by Ruggles et al. (2010), which contains information on the occupation of workers. We follow the classification of occupations into three categories by Acemoglu and Autor (2011): manual (non-routine non-cognitive), routine (both cognitive and non-cognitive) and abstract (non-routine cognitive). We implement this classification by relying on a harmonized and balanced panel of occupational codes as in Autor and Dorn (2013) and Bárány and Siegel (2018). We then classify each worker into one of these three broad occupations and into one of the three sectors defined earlier.[12] Given this classification we can calculate the share of hours worked by occupation o workers within a sector J. We measure sectoral employment shares and overall employment growth using Full Time Equivalent (FTE) employees by industry provided by the BEA.[13] To get the employment share of a sector-occupation cell, l_{oJ}, we multiply the within-sector hours share of occupation o (from the Census/ACS) by the employment share of sector J in the economy (from the BEA). We also calculate the labor income share of occupation o in sector J as:

$$\theta_{oJ} \equiv \frac{\text{earnings of occupation } o \text{ workers in sector } J}{\text{earnings of sector } J \text{ workers}}. \tag{1}$$

Relative average occupational wages within sectors can then be calculated as

$$\frac{w_{oJ}}{w_{rJ}} = \frac{\theta_{oJ}}{\theta_{rJ}} \frac{l_{rJ}}{l_{oJ}}. \tag{2}$$

Figure 4 shows the employment share of each sector (l_J) and of each sector-occupation cell (l_{oJ}, in the top row), as well as within each sector the labor income share of each occupation (θ_{oJ}, in the middle row) and the average wage of abstract and manual relative to routine occupations (w_{aJ}/w_{rJ} and w_{mJ}/w_{rJ} in the bottom row).[14]

Clearly, the share of labor income earned by routine workers declined in each sector (as seen in the middle row). This is driven by the falling employment share of routine workers (plotted in the top row), and by their wages which tend to fall relative to the other occupations (bottom row). Note that the relative average occupational

[12]See appendix A for more details on the classification of occupations and Table A1 for industries.

[13]As for the data on the components of value added, we again have to combine data based on two different industry classification systems (SIC until 1998, NAICS afterwards).

[14]In section 5.3 we consider a variant of this framework where we control for observable characteristics of workers (as one might be concerned that these are confounding the patterns of average wages). Note, the income shares we show here are informative even if there is heterogeneity amongst workers in terms of their human capital.

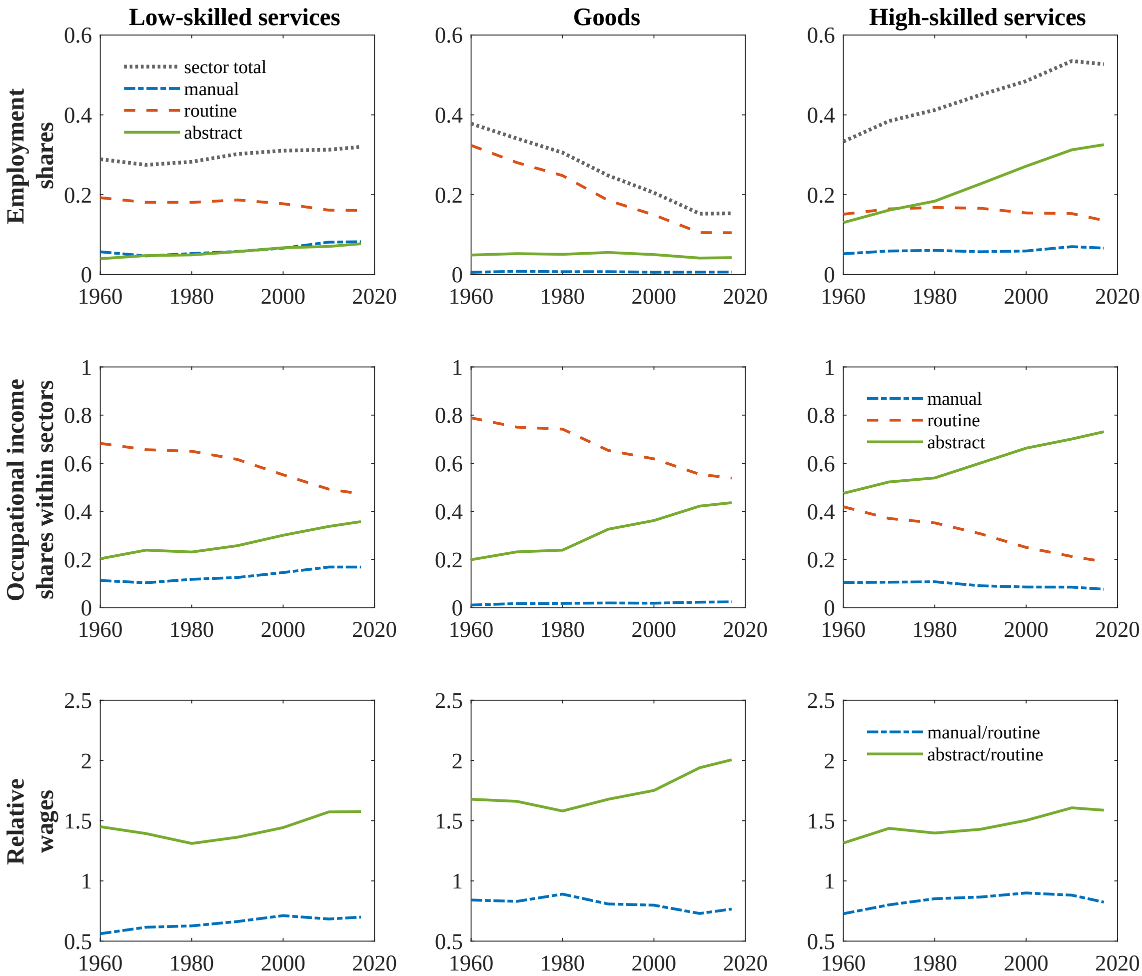

Figure 4: Sector-occupation income shares, hours shares, and relative wages

Notes: Sectoral employment shares are based on BEA data on full time equivalent workers. The data on occupational employment, income and wages 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, goods, high-skilled services) and three occupational categories (manual, routine, abstract), this figure plots in the top row the evolution of employment shares in sector-occupation cells, as well as in sectors (dark gray dotted lines), in the middle row each occupation's share in sectoral labor income, and in the bottom row the ratio of manual to routine wages and of abstract to routine wages within the given sector.

hourly wages are not equalized across sectors. This is an important observation, and we want to allow for sectoral differences in occupational wages in our model and when extracting technologies.

The top row of Figure 4 demonstrates that all of the three sectors employ workers in each of the three occupations, but at different intensities. It is therefore a possibility that the observed sectoral differences in labor productivity growth are due to differences in occupational labor use. Note that the goods sector is the most intensive in routine workers, while high-skilled services is the most intensive in abstract workers.

Now suppose that technological change increased routine workers' productivity the most, but equally across sectors. It is then conceivable that the differences in occupational intensities generate the sectoral differences in measured labor productivity growth (in terms of all workers), especially the high growth in goods. Moreover, the observed slowdown in aggregate productivity growth could be driven by the contraction of routine employment in all sectors. We evaluate the role of these mechanisms in section 4.

3 A production side framework

In order to study the drivers of sectoral labor productivity growth, we specify a production side framework. We assume a relatively flexible CES functional form for sectoral production, which allows matching the data – especially the time varying factor income shares – we documented in the previous section. Note that with CES production functions relative factor prices in equilibrium depend both on relative supplies and on relative productivities. This means the framework does not hard-wire where changes in relative wages are stemming from. Another advantage of the CES framework is that it is relatively simple and does not require too many parameters (as argued in Krusell et al. (2000)). As discussed in the introduction, we consider as inputs manual, routine and abstract occupational labor, as well as ICT and non-ICT capital. We back out the path of factor-augmenting technologies from each sector's optimality conditions, conditional on values for the various elasticities of substitution, using data on sectoral growth rates, value added, quantities and prices of factor inputs. It is important to note that we conduct this exercise making assumptions about the production side of the economy only. We do not need to take a stance on where the demand for goods and services stem from, since observing the sectoral value added is sufficient. Similarly, observing the quantities and prices of factors employed in production is sufficient and we do not need to model capital accumulation or labor supply choices. The method in this exercise is similar to Buera et al. (2018), but with a very different focus. We allow for heterogeneity in labor across occupations and want to identify the drivers of sectoral labor productivity growth.

3.1 Sectoral production

Firms in sector J combine occupational labor (manual, routine and abstract), ICT capital and non-ICT capital as inputs according to a nested CES production function. The inner most nest gives the *routine aggregate*:

$$RA_{J,t} = \left[\alpha_{rJ}(A_{rJ,t}l_{rJ,t})^{\frac{\sigma_c-1}{\sigma_c}} + \alpha_{cJ}(A_{cJ,t}c_{J,t})^{\frac{\sigma_c-1}{\sigma_c}}\right]^{\frac{\sigma_c}{\sigma_c-1}}, \tag{3}$$

where $l_{rJ,t}$ is routine labor, $c_{J,t}$ is ICT capital used in sector J, $A_{rJ,t}$ and $A_{cJ,t}$ are *sector-specific routine labor-* and respectively *ICT capital-augmenting technologies*, all at time t. We denote the time-invariant factor intensities in sector J by α_{rJ} and α_{cJ}, where $\alpha_{rJ} + \alpha_{cJ} = 1$. This inner nest combines routine labor and ICT capital with an elasticity of substitution σ_c.

The middle nest gives the *labor aggregate*:

$$LA_{J,t} = \left(\alpha_{mJ}(A_{mJ,t}l_{mJ,t})^{\frac{\rho-1}{\rho}} + \alpha_{aJ}(A_{aJ,t}l_{aJ,t})^{\frac{\rho-1}{\rho}} + \alpha_{RAJ}RA_{J,t}^{\frac{\rho-1}{\rho}}\right)^{\frac{\rho}{\rho-1}}. \tag{4}$$

In this formulation $l_{oJ,t}$ is occupation $o \in \{m, a\}$ labor used in sector J, and $A_{oJ,t} > 0$ is a *sector-specific factor-augmenting* technology term for manual and abstract occupations, all in period t. We denote the factor weights by α_{mJ}, α_{aJ} and α_{RAJ}, with $\alpha_{mJ} + \alpha_{aJ} + \alpha_{RAJ} = 1$. This nest essentially combines the different types of labor, including the most inner routine aggregate, with an elasticity of substitution ρ.

The output in sector J is produced according to:

$$Y_{J,t} = \left[\alpha_{LAJ}LA_{J,t}^{\frac{\sigma-1}{\sigma}} + \alpha_{kJ}(A_{kJ,t}k_{J,t})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{5}$$

In this formulation $k_{J,t}$ is non-ICT capital used in sector J in period t and $A_{kJ,t} > 0$ is a *sector-specific non-ICT capital-augmenting* technology in period t. We denote the time-invariant factor weights in sector J by α_{kJ} and α_{LAJ}; these sum to 1. This outer-most layer combines aggregate labor and non-ICT capital with a substitution elasticity σ.

Each CES-layer of the production function allows for factor income shares (at the sectoral level) to change over time which is one of the salient features we have documented in the data in the previous section. The most inner nest of the production function for $\sigma_c > 1$ reflects the idea that ICT is a good substitute for routine workers (which is the consensus in the literature, e.g. Autor, Levy, and Murnane (2003), Autor and Dorn (2013)). The aggregator of occupational labor inputs is based on the notion that workers in different occupations perform different tasks and are thus imperfect substitutes in production, as for instance emphasized in a task-based model of the labor market (see Acemoglu and Autor (2011)). For $\rho \in (0,1)$ occupational inputs are complements, and if ICT capital substitutes for routine workers, it complements workers of other occupations, as in Autor and Dorn (2013).[15]

Notice that the $A_{fJ,t}$ terms capture sector-specific factor-augmenting technologies in sector J in period t. The time-invariant α_{fJ} capture the intensity at which sector J uses an input factor (or input aggregate) f in production. Without further assumptions or normalization regarding the weights, these are not separately identifiable. It is the following combination of the weights and the factor-augmenting technologies which we can extract from the data relying on key equations of the model:

$$Z_{kJ,t} = \alpha_{kJ}^{\frac{\sigma}{\sigma-1}} A_{kJ,t}, \tag{6}$$

$$Z_{oJ,t} = \alpha_{LAJ}^{\frac{\sigma}{\sigma-1}} \alpha_{oJ}^{\frac{\rho}{\rho-1}} A_{oJ,t} \qquad \text{for } o \in \{m, a\}, \tag{7}$$

$$Z_{nJ,t} = \alpha_{LAJ}^{\frac{\sigma}{\sigma-1}} \alpha_{RAJ}^{\frac{\rho}{\rho-1}} \alpha_{nJ}^{\frac{\sigma_c}{\sigma_c-1}} A_{nJ,t} \qquad \text{for } n \in \{r, c\}. \tag{8}$$

It is precisely these $Z_{fJ,t}$ terms that matter for labor productivity growth. Since $\Delta \log Z_{fJ,t} = \Delta \log A_{fJ,t}$, any change in the inferred $Z_{fJ,t}$ over time reflects sector-specific factor-augmenting technological change. We refer to $Z_{fJ,t}$ as sector-specific factor-augmenting technologies.

The evolution of sectoral output is entirely driven by changes in the amount of effective factor inputs over time. How large the impact of a change in a given factor is depends on the initial importance of this factor. Below we express the gross growth rate in each nest of the production function as a weighted average of the growth in

[15]Since there is no hard evidence on elasticities of substitution between occupational labor inputs, for simplicity we assume that they are combined in the CES aggregator in this 'symmetric' way with a common elasticity. Our framework can easily accommodate other nestings of occupational labor inputs and capital, and we explore one particular alternative as a robustness check in section 5.2.

each effective input raised to the relevant power.

$$
\begin{aligned}
\frac{RA_{J,t}}{RA_{J,0}} &= \Bigg[\underbrace{\frac{(Z_{rJ,0}l_{rJ,0})^{\frac{\sigma_c-1}{\sigma_c}}}{(Z_{rJ,0}l_{rJ,0})^{\frac{\sigma_c-1}{\sigma_c}} + (Z_{cJ,0}c_{J,0})^{\frac{\sigma_c-1}{\sigma_c}}}}_{\equiv \omega_{rJ}} \left[\frac{Z_{rJ,t}l_{rJ,t}}{Z_{rJ,0}l_{rJ,0}}\right]^{\frac{\sigma_c-1}{\sigma_c}} + \underbrace{\frac{(Z_{cJ,0}c_{J,0})^{\frac{\sigma_c-1}{\sigma_c}}}{(Z_{rJ,0}l_{rJ,0})^{\frac{\sigma_c-1}{\sigma_c}} + (Z_{cJ,0}c_{J,0})^{\frac{\sigma_c-1}{\sigma_c}}}}_{\equiv \omega_{cJ}} \left[\frac{Z_{cJ,t}c_{J,t}}{Z_{cJ,0}c_{J,0}}\right]^{\frac{\sigma_c-1}{\sigma_c}} \Bigg]^{\frac{\sigma_c}{\sigma_c-1}} \\
\frac{LA_{J,t}}{LA_{J,0}} &= \Bigg[\underbrace{\frac{(Z_{mJ,0}l_{oJ,0})^{\frac{\rho-1}{\rho}}}{(Z_{mJ,0}l_{mJ,0})^{\frac{\rho-1}{\rho}} + (Z_{aJ,0}l_{aJ,0})^{\frac{\rho-1}{\rho}} + RA_{J,0}^{\frac{\rho-1}{\rho}}}}_{\equiv \omega_{mJ}} \left[\frac{Z_{mJ,t}l_{mJ,t}}{Z_{mJ,0}l_{mJ,0}}\right]^{\frac{\rho-1}{\rho}} + \underbrace{\frac{(Z_{aJ,0}l_{aJ,0})^{\frac{\rho-1}{\rho}}}{(Z_{mJ,0}l_{mJ,0})^{\frac{\rho-1}{\rho}} + (Z_{aJ,0}l_{aJ,0})^{\frac{\rho-1}{\rho}} + RA_{J,0}^{\frac{\rho-1}{\rho}}}}_{\equiv \omega_{aJ}} \left[\frac{Z_{aJ,t}l_{aJ,t}}{Z_{aJ,0}l_{aJ,0}}\right]^{\frac{\rho-1}{\rho}} \\
&\quad + \underbrace{\frac{RA_{J,0}^{\frac{\rho-1}{\rho}}}{(Z_{mJ,0}l_{mJ,0})^{\frac{\rho-1}{\rho}} + (Z_{aJ,0}l_{aJ,0})^{\frac{\rho-1}{\rho}} + RA_{J,0}^{\frac{\rho-1}{\rho}}}}_{\equiv \omega_{RAJ}} \left[\frac{RA_{J,t}}{RA_{J,0}}\right]^{\frac{\rho-1}{\rho}} \Bigg]^{\frac{\rho}{\rho-1}} \\
\frac{Y_{J,t}}{Y_{J,0}} &= \Bigg[\underbrace{\frac{LA_{J,0}^{\frac{\sigma-1}{\sigma}}}{LA_{J,0}^{\frac{\sigma-1}{\sigma}} + (Z_{kJ,0}k_{J,0})^{\frac{\sigma-1}{\sigma}}}}_{\equiv \omega_{LAJ}} \left[\frac{LA_{J,t}}{LA_{J,0}}\right]^{\frac{\sigma-1}{\sigma}} + \underbrace{\frac{(Z_{kJ,0}k_{J,0})^{\frac{\sigma-1}{\sigma}}}{LA_{J,0}^{\frac{\sigma-1}{\sigma}} + (Z_{kJ,0}k_{J,0})^{\frac{\sigma-1}{\sigma}}}}_{\equiv \omega_{kJ}} \left[\frac{Z_{kJ,t}k_{J,t}}{Z_{kJ,0}k_{J,0}}\right]^{\frac{\sigma-1}{\sigma}} \Bigg]^{\frac{\sigma}{\sigma-1}} \qquad (9)
\end{aligned}
$$

This formulation shows that the impact of a change in each effective input depends on the relative initial amount of the given effective input (to the relevant power) within the given nest. These weights depend on a combination of initial technology levels and the initial quantity of inputs. The weights ω_{rJ} and ω_{cJ}, for example, determine how large an impact a change in the effective routine labor and in the effective ICT capital has on the growth of RA_J. These weights play an important role when evaluating the impact of changes in factor quantities and in factor-augmenting technologies.

3.2 Inferring factor-augmenting technologies

The assumptions about the production side of the economy allow us to infer sector-specific factor-augmenting technologies (sector-factor specific technologies for short, the $Z_{fJ,t}$s) from observables. In addition to the sectoral production functions, we assume that there is perfect competition in all markets, such that firms take prices as given.

Here we describe in detail how we can back out the factor-augmenting technologies from the data. First, using optimality conditions for production in each sector we express relative factor-augmenting technologies within a sector and period. Second, we derive how the growth of sectoral value added pins down the evolution of technologies within each sector over time.[16] In what follows, where possible, we omit the time t subscripts to simplify the notation.

We have shown in Figure 4 that average occupational wages are not equalized across sectors, and we want to allow for this in our model. We therefore assume that wages are sector-occupation specific and denote them by w_{oJ}. Assuming further that the rental rates of ICT (R_c) and non-ICT capital (R_k) are equalized across sectors, the profit maximization problem of firms in each sector is

$$\max_{\{l_{oJ}\}, c_J, k_J} p_J Y_J - \sum_o w_{oJ} l_{oJ} - R_c c_J - R_k k_J,$$

subject to (5), where p_J denotes the price of sector J output. Optimal input use in each sector has to satisfy the following first order conditions:

$$\frac{\partial \pi_J}{\partial l_{oJ}} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho}{\rho\sigma}} Z_{oJ}^{\frac{\rho-1}{\rho}} l_{oJ}^{-\frac{1}{\rho}} - w_{oJ} = 0 \quad \text{for} \quad o \in \{m, a\}, \tag{10}$$

$$\frac{\partial \pi_J}{\partial l_{rJ}} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho}{\rho\sigma}} RA_J^{\frac{\rho-\sigma_c}{\sigma_c\rho}} Z_{rJ}^{\frac{\sigma_c-1}{\sigma_c}} l_{rJ}^{-\frac{1}{\sigma_c}} - w_{rJ} = 0, \tag{11}$$

$$\frac{\partial \pi_J}{\partial c_J} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho}{\rho\sigma}} RA_J^{\frac{\rho-\sigma_c}{\sigma_c\rho}} Z_{cJ}^{\frac{\sigma_c-1}{\sigma_c}} c_J^{-\frac{1}{\sigma_c}} - R_c = 0, \tag{12}$$

$$\frac{\partial \pi_J}{\partial k_J} = p_J Y_J^{\frac{1}{\sigma}} Z_{kJ}^{\frac{\sigma-1}{\sigma}} k_J^{-\frac{1}{\sigma}} - R_k = 0. \tag{13}$$

Inferring technologies within sectors. We can express the relative optimal demand for factor inputs from the first order conditions as a function of relative factor prices and relative technologies. We invert these to express relative technologies in terms of relative wages, rental rates and relative factor incomes within sectors.

The first order conditions on manual and abstract labor, (10), pin down the optimal relative labor use as:

$$\frac{l_{aJ}}{l_{mJ}} = \left(\frac{w_{mJ}}{w_{aJ}}\right)^{\rho} \left(\frac{Z_{aJ}}{Z_{mJ}}\right)^{\rho-1}. \tag{14}$$

It is optimal to use more abstract relative to manual labor in sector J if the relative manual wage, w_{mJ}/w_{aJ}, is higher. Additionally, if the term $(Z_{aJ}/Z_{mJ})^{\rho-1}$ is larger it is optimal to use relatively more abstract labor in that sector. Multiply the above by

[16] The full derivations can be found in appendix C.

w_{aJ}/w_{mJ} and re-arrange to get:

$$\frac{Z_{mJ}}{Z_{aJ}} = \frac{w_{mJ}}{w_{aJ}} \left(\frac{w_{mJ} l_{mJ}}{w_{aJ} l_{aJ}} \right)^{\frac{1}{\rho-1}} = \frac{w_{mJ}}{w_{aJ}} \left(\frac{\theta_{mJ}}{\theta_{aJ}} \right)^{\frac{1}{\rho-1}}, \tag{15}$$

where $\theta_{mJ} = (w_{mJ} l_{mJ})/(\sum_o w_{oJ} l_{oJ})$ is the share of labor income in sector J that is going to manual workers. Equation (15) shows that conditional on ρ, observing the relative wage and the relative income share of manual and abstract workers within a sector, both shown in Figure 4, allows us to infer relative occupation-augmenting technologies in that sector at a given point in time. This equation thereby also implies the relative growth rates of these technologies over time.

Similarly, from the first order conditions on routine labor and ICT capital, (11) and (12), we can express the relative demand for these factors, and consequently their relative Z as well:

$$\frac{Z_{cJ}}{Z_{rJ}} = \frac{R_c}{w_{rJ}} \left(\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}} \right)^{\frac{1}{\sigma_C - 1}}, \tag{16}$$

where $\Theta_{cJ} = (R_c c_J)/(p_J Y_J)$ is the share of income in sector J paid to ICT capital, and $\Theta_J = (R_c c_J + r_k k_J)/(p_J Y_J)$ is the share of income in sector J paid to both types of capital. This expression is very similar to the one in (15), except that it is a different elasticity of substitution that is relevant.

Expressing the remaining two relative technology levels within sectors, Z_{rJ}/Z_{mJ} and Z_{kJ}/Z_{mJ}, follows a similar principle, but is slightly more convoluted, and we delegate the details of these derivations to appendix C. Here we only explain the intuition. First, from the optimal use of routine labor relative to ICT capital, we express RA_J, the routine aggregate, in terms of routine labor only. This then allows us to express the optimal use of manual relative to routine labor within a sector, which, multiplied by w_{rJ}/w_{mJ}, gives us the relative technologies as:

$$\frac{Z_{mJ}}{Z_{rJ}} = \frac{w_{mJ}}{w_{rJ}} \left(\frac{\theta_{mJ}}{\theta_{rJ}} \right)^{\frac{1}{\rho-1}} \left[1 + \frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}} \right]^{\frac{\rho-\sigma_C}{(\sigma_C-1)(\rho-1)}}. \tag{17}$$

Next we express LA_J, the labor aggregate, in terms of manual labor only, which again allows us to express the optimal use of manual labor relative to non-ICT capital. Multiplying this by relative factor prices and re-arranging we get:

$$\frac{Z_{kJ}}{Z_{mJ}} = \frac{R_k}{w_{mJ}} \left(\frac{1}{\theta_{mJ}}\right)^{\frac{1}{\rho-1}} \left(\frac{\Theta_J - \Theta_{cJ}}{1-\Theta_J}\right)^{\frac{1}{\sigma-1}} \left(1 + \frac{\Theta_{cJ}}{1-\Theta_J}\right)^{\frac{\sigma-\rho}{(\rho-1)(\sigma-1)}} . \tag{18}$$

Thus, we showed how to infer all relative technologies within a sector and a period from observables, conditional on the elasticities ρ, σ_c and σ. Taking for example Z_{kJ} as the base technology, all other factor-augmenting technologies in sector J are proportional to Z_{kJ}, where the proportionality depends on observables in the data, and on the values of the three substitution elasticities.

Inferring technologies over time. The last step is to pin down the evolution of the Zs over time in each sector, as well as the initial values of the technologies. Until now we did not index any variable by time, as we explained how to infer the relative Zs within a period. Plugging all the optimal relative input use expressions in (5) sectoral output can be expressed as:

$$Y_{J,t} = Z_{kJ,t} k_{J,t} \left(\frac{1}{\Theta_{J,t} - \Theta_{cJ,t}}\right)^{\frac{\sigma}{\sigma-1}} .$$

The evolution of the $Z_{kJ,t}$ over time is then given by:

$$\frac{Z_{kJ,t}}{Z_{kJ,0}} = \frac{Y_{J,t}}{Y_{J,0}} \frac{k_{J,0}}{k_{J,t}} \left(\frac{\Theta_{J,t} - \Theta_{cJ,t}}{\Theta_{J,0} - \Theta_{cJ,0}}\right)^{\frac{\sigma}{\sigma-1}} . \tag{19}$$

Again in equation (19) all right-hand side variables can be observed in the data, and hence, conditional on σ, this equation gives us the growth rate of $Z_{kJ,t}$ over time. Intuitively, the change in the capital-augmenting technology reflects by how much the sector's value-added per unit of capital changed (the first two factors) and how this value-added is split between the production factors taking into account the substitution elasticity (the third factor).

Finally, we need to pin down the initial level of Zs. It is important to note that these have no impact on our conclusions regarding the drivers of sectoral labor productivity growth; they only matter for the growth rate of labor productivity in the aggregate economy.[17] We infer these initial Zs from initial sectoral prices. Using the

[17] Even for this, only the relative initial Zs matter, i.e. we could normalize one of the sectors' $Z_{fJ,0}$

above expression for sectoral output in the first order condition on non-ICT capital we get:

$$Z_{kJ,0} = \frac{R_{k,0}}{p_{J,0}} \left(\Theta_{J,0} - \Theta_{cJ,0}\right)^{\frac{1}{\sigma-1}} . \tag{20}$$

Equations (15), (16), (17), (18), (19) and (20) describe how to infer factor-augmenting technologies in each sector and in all periods. Note that equations (15) to (19) are implied by firms' cost minimization and therefore would still hold if there were imperfect competition in product markets. As such, our conclusions about the drivers of sectoral labor productivity growth would also hold if firms were charging – potentially time-varying – mark-ups. It is worth to point out that as the observed path of income shares, factor prices and quantities does not evolve the same way across sectors, there is no reason why the technologies of a given factor should evolve the same way across sectors. Thus these technologies in general will be sector-factor specific.

3.3 Implementation

To infer the sector-specific factor-augmenting technologies from the data using equations (15) to (20), we need the value of three elasticities. First, we need the elasticity of substitution between non-ICT capital and the labor aggregate, σ. The overwhelming majority of studies which estimate the elasticity of substitution between capital and labor from aggregate data finds values below one, see Table 1 in León-Ledesma, McAdam, and Willman (2010) for a recent summary.[18] Lawrence (2015) obtains estimates ranging from 0.27 to 0.96 for this elasticity in the (total) manufacturing sector. Oberfield and Raval (2014) follow a more micro approach, and estimate the elasticity of substitution between capital and labor in the US manufacturing sector by aggregating the actions of individual plants, and find a value around 0.7. Closest to our setup with sectoral CES production functions is Herrendorf et al. (2015), though we differentiate between various types of occupational labor. While they find differences across sectors, they report for the aggregate economy an elasticity of 0.84. We take this value

without loss of generality.

[18]These studies estimate jointly the elasticity of substitution and a constant growth rate of (either Hicks-neutral or factor-augmenting) technological change. As discussed in the introduction, since we do not impose any restrictions on how technologies evolve over time we cannot identify both technologies and elasticities from the data.

for our baseline parametrization, but in the robustness checks of section 5.1.1 we also explore model variants with sector-specific elasticities.

Second, we need the elasticity of substitution between ICT capital and routine labor, σ_c. While the literature has argued that routine labor and ICT capital are very good substitutes, there are surprisingly few estimates of this elasticity. Eden and Gaggl (2018) estimate a CES production function differently nested to ours, where the elasticity of substitution between ICT capital and routine labor is not constant, but it ranges between 2.1 and 3.3. Aum, Lee, and Shin (2018) calibrate industry specific elasticities between ICT capital and all types of occupational labor and find values between 1.2 and 1.8. As our baseline we set $\sigma_c = 2$, in the mid-range of these estimates.

Third, we need the elasticity of substitution between the different occupations and the routine aggregate, ρ. Goos, Manning, and Salomons (2014) estimate an elasticity of substitution of 0.9 between 21 occupations, Lee and Shin (2017) calibrate $\rho = 0.7$ and Aum et al. (2018) calibrate 0.81 both among 10 occupations, and Duernecker and Herrendorf (2016) calibrate an elasticity of 0.56 between 2 occupations. It is likely that the more coarse the occupation categories are, the lower is the elasticity of substitution. In our model with three occupational categories we therefore set $\rho = 0.6$. We summarize these values for the substitution elasticities in Table 1. While we use these values for the three elasticities as our baseline, we conduct in section 5 extensive robustness checks, also with respect to these elasticities.

Table 1: Calibrated substitution elasticities

	capital – labor, σ	**ICT capital – routine labor, σ_c**	**different occupations, ρ**
Value	0.84	2	0.6

To infer the evolution of technologies over time we need the following measures from the data for every period: sector-occupation specific wage rates ($w_{oJ,t}$), rental rates for non-ICT and ICT capital ($R_{k,t}$ and $R_{c,t}$), the income share of occupations within sectors ($\theta_{oJ,t}$), the share of sectoral value added paid to ICT capital ($\Theta_{cJ,t}$), and to both types of capital together ($\Theta_{J,t}$), the quantity of non-ICT capital by sector ($k_{J,t}$), the per worker growth rate of sectoral value added ($Y_{J,t}/Y_{J,0}$), as well as sectoral prices in the initial period ($p_{J,0}$). In section 2 we showed $\theta_{oJ,t}$, $\Theta_{J,t}$, $p_{J,t}$, and $\gamma_{J,t}$, calculated as the growth rate of real value added in sector J (shown in Figure 1) divided by the growth rate of full time equivalent workers from the BEA. Note that without loss of generality we normalize all our quantity measures by the FTE workforce, i.e. we use employment shares, the stocks of ICT and non-ICT capital per worker, growth of real value added in each sector per worker, and nominal value added per worker. In the quantitative analysis rather than using workers' self-reported income from the Census/ACS, we use the following accounting identity to obtain sector-occupation wage rates, $w_{oJ,t}$:

$$w_{oJ,t}l_{oJ,t} = Y_t^{nom} \cdot VA_{J,t}(1 - \Theta_{J,t})\theta_{oJ,t},$$

where Y_t^{nom} is nominal GDP per worker in year t and $VA_{J,t}$ is the share of value added produced in sector J (shown in Figure 1). This accounting identity ensures that the sum of all income paid to workers of different occupations within a sector is equal to the nominal labor income in that sector. Note that relative occupational wages within a sector are the same as those calculated from the micro data (see equation (28) and the discussion that follows in appendix B). Using similar accounting identities and a no arbitrage condition, we obtain $R_{k,t}$, $R_{c,t}$ and $\Theta_{cJ,t}$ from the data shown in Figure 2 and 3 as explained in appendix B. These accounting identities ensure that the sum of all factor incomes is equal to nominal value added.

4 The role of changing technologies and input use

Table 2 shows the average annual growth rate of sector-factor specific technologies between 1960 and 2017. Technological change has been uneven, within each sector across factors, as well as across sectors for a given factor (i.e. for a given occupation or a type of capital). Nonetheless some patterns can be discerned. It is obvious that among the three occupations routine labor had the highest growth rate in all sectors, roughly between 1.3 and 5.6 percent annually. Technological change augmenting manual labor was much more modest and less dispersed across sectors. Finally, technological change augmenting abstract labor varied across sectors, with negative growth rates in low- and in high-skilled services.[19] In terms of capital-augmenting technologies we find that those related to ICT increased rapidly in low-skilled services and in goods and fell in high-skilled services, while those augmenting non-ICT capital increased at a lower rate in low- and in high-skilled services and they fell in goods.[20] Our results highlight that routine workers became more productive over and beyond what is embodied in ICT capital, in line with what Aum et al. (2018) find. Technologies augmenting routine workers increased the most in all sectors, even after accounting for the increase in ICT capital (c_J) and in its productivity (Z_{cJ}). In terms of sectoral patterns, the growth rates of all factor-augmenting technologies were the highest in goods, followed by low-skilled services, except for manual labor and non-ICT capital, which had the highest growth in high-skilled services. Thus beyond the factor-specific patterns, there also seem to be sector-specific components to technological progress.

Table 2: Average annual growth rate of Zs between 1960 and 2017

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
Low-skilled services	0.26%	2.93%	-0.69%	0.85%	2.02%
Goods	0.59%	5.61%	0.98%	-1.61%	4.43%
High-skilled services	0.72%	1.32%	-2.38%	1.78%	-1.94%

These growth rates are inferred from the data using equations (15)-(19) conditional on the elasticities summarized in Table 1. For example, to understand the origins of the growth rate of non-ICT capital-augmenting technologies by sector consider the following expression (from equation (19)):

$$\frac{Z_{kJ,t}k_{J,t}/Y_{J,t}}{Z_{kJ,0}k_{J,0}/Y_{J,0}} = \underbrace{\left[\frac{\Theta_{J,t} - \Theta_{cJ,t}}{\Theta_{J,0} - \Theta_{cJ,0}}\right]}_{\text{non-ICT capital share}}^{\frac{\sigma}{\sigma-1}}. \tag{21}$$

In the goods sector the share of income paid to non-ICT capital increased substantially, which with an elasticity of substitution smaller than 1 ($\sigma = 0.84$ is our baseline) implies that effective capital relative to real output decreased in this sector. As real output

[19]These negative growth rates might be explained by a compositional change within abstract occupations in these sectors, towards more time-consuming tasks.

[20]These negative growth rates are in line with what previous literature has found. Both Antràs (2004) and Herrendorf et al. (2015) find negative capital-augmenting technological change at the aggregate level, and respectively in the manufacturing and service sectors.

increased, in turn this implies that effective non-ICT capital fell. Putting this together with the slight increase in non-ICT capital in the goods sector over time (see Table 3 for the evolution of inputs and the middle panel of Figure 1 for real sectoral value added), we obtain the result that the technology augmenting non-ICT capital in the goods sector had to fall.

We can obtain similar expressions for the growth of the technologies augmenting all other factors. For manual labor augmenting technologies combine expression (18) with (21) to obtain:

$$\frac{Z_{mJ,t}l_{mJ,t}/Y_{J,t}}{Z_{mJ,0}l_{mJ,0}/Y_{J,0}} = \underbrace{\left[\frac{1-\Theta_{J,t}+\Theta_{cJ,t}}{1-\Theta_{J,0}+\Theta_{cJ,0}}\right]^{\frac{\rho-\sigma}{(\rho-1)(\sigma-1)}}}_{LA\text{ share}} \underbrace{\left[\frac{(1-\Theta_{J,t})\theta_{mJ,t}}{(1-\Theta_{J,0})\theta_{mJ,0}}\right]^{\frac{\rho}{\rho-1}}}_{\text{manual labor share}}. \tag{22}$$

The growth in effective manual labor intensity in sector J depends on the growth in the income share of the labor aggregate (LA) and the growth in the income share of manual labor within sectoral value added.[21] This expression is similar to the one for non-ICT capital, except it contains two terms on the right hand side. This is because manual labor enters the production function through the middle nest. Manual labor-augmenting technologies impact the trade-off between manual labor and the other inputs within the labor aggregate, and they also impact the trade-off between the labor aggregate and non-ICT capital. The two terms on the right hand side of expression (22) represent these two margins. What the evolution of these two income shares reveal about the evolution of the effective manual labor input relative to output depends on the two relevant elasticities. As $\rho < 1$, an increase in the manual income share implies a reduction in the effective manual labor input intensity. If $\rho < \sigma < 1$, then an increase in the share of income going to the labor aggregate implies a decrease in the effective manual labor input intensity. Netting out the actual growth in manual labor input intensity (l_{mJ}/Y_J), we obtain the growth in manual labor-augmenting technology by sector. Table A2 in the appendix contains by sector the income share in value added of each factor in 1960 and in 2017; these values together with the evolution of factor inputs and of sectoral real value added pin down the growth rate of technologies.

[21]A similar expression can be obtained for abstract labor-augmenting technologies by combining (22) with (15).

For the growth of sector-specific routine-augmenting technologies by combining (22) with (17) we obtain:

$$\frac{Z_{rJ,t}l_{rJ,t}/Y_{J,t}}{Z_{rJ,0}l_{rJ,0}/Y_{J,0}} = \underbrace{\left[\frac{1-\Theta_{J,t}+\Theta_{cJ,t}}{1-\Theta_{J,0}+\Theta_{cJ,0}}\right]^{\frac{\rho-\sigma}{(\rho-1)(\sigma-1)}}}_{LA\text{ share}} \underbrace{\left[\frac{(1-\Theta_{J,t})\theta_{rJ,t}+\Theta_{cJ,t}}{(1-\Theta_{J,0})\theta_{rJ,0}+\Theta_{cJ,0}}\right]^{\frac{\sigma_c-\rho}{(\sigma_c-1)(\rho-1)}}}_{RA\text{ share}} \underbrace{\left[\frac{(1-\Theta_{J,t})\theta_{rJ,t}}{(1-\Theta_{J,0})\theta_{rJ,0}}\right]^{\frac{\sigma_c}{\sigma_c-1}}}_{\text{routine labor share}}. \tag{23}$$

As routine labor enters the production in the inner-most nest, the evolution of three income shares is informative about the evolution of the effective routine intensity. These three income shares are: that of the labor aggregate (LA), of the routine aggregate (RA) and of routine labor.[22] Let us consider the evolution of all three of these income shares. The first one, the LA share, is the labor income share plus the share of income going to ICT capital. The evolution of this is very similar to the evolution of the labor income share, as the quantity of ICT capital and the share of income paid to it is very low (albeit increasing rapidly). The LA share decreased in low-skilled services and in goods, while it increased slightly in high-skilled services between 1960 and 2017. As the share of labor income going to routine labor decreased steeply in all sectors (see Figure 4) the second and third income shares decreased in all sectors quite substantially. As our baseline is that $\rho < \sigma < 1 < \sigma_c$, the exponents on the first two income shares are negative, while on the last one it is positive. The last two terms in (23) thus go in opposite directions, but the increase implied by the middle term more than offsets the decrease implied by the last term in all sectors. Moreover routine employment declined in all sectors, which implies that routine labor-augmenting technologies had to increase substantially in all sectors.

In the above discussion we made use of how inputs changed between 1960 and 2017. Most of these numbers can be seen in Figure 4, for completeness we summarize all of these in Table 3. We also add a column which contains the gross growth rate of total employment in the sector, as these changes are important in the calculation of labor productivity growth.

Output and labor productivity in each sector increased over time because of changes in effective factor inputs. As discussed in section 3.1 the impact of a change in the

[22] A symmetric expression can be obtained for ICT capital by combining (23) with (16).

Table 3: Gross growth rate of inputs between 1960 and 2017

	total	occupations			capital	
	labor	manual	routine	abstract	non-ICT	ICT
Low-skilled services	1.107	1.443	0.834	1.945	1.393	192.820
Goods	0.406	1.115	0.324	0.867	1.002	94.280
High-skilled services	1.582	1.274	0.897	2.500	1.985	80.410

Notes: The numbers are the ratio of the 2017 to the 1960 input quantity (per full time equivalent worker), so for example, the second cell shows $l_{mL,2017}/l_{mL,1960}$.

amount of effective input on sectoral output depends on the weight that the given factor receives in the initial period, as defined in equation (9). Table 4 shows the weights for 1960. Weights within a nest are informative; a larger weight implies that a change in the effective input, due to changes in the technology or in the quantity of the input, has a larger impact on the evolution of the given nest.

Table 4: Initial factor weights

	ω_{kJ}	ω_{LAJ}	ω_{mJ}	ω_{aJ}	ω_{RAJ}	ω_{rJ}	ω_{cJ}
Low-skilled services	0.413	0.587	0.113	0.203	0.684	0.994	0.006
Goods	0.368	0.632	0.012	0.199	0.789	0.995	0.005
High-skilled services	0.501	0.499	0.101	0.458	0.441	0.916	0.084

In what follows we study the drivers of sectoral labor productivity growth in detail, by computing average sectoral labor productivity growth rates between 1960 and 2017 for various counterfactual scenarios. First, we assess the importance of the various forms of technological change. To do this, we take factor inputs from the data and fix technologies at counterfactual values. In addition, we use a factor model to identify common sector and occupation components in the sector-occupation-specific labor-augmenting technologies. Second, to quantify the role of changing input use and of differences in occupational employment shares across sectors, we use the Zs as extracted from the data, and fix factor inputs at counterfactual levels. Comparing these two sets of counterfactuals to each other sheds light on whether changing inputs or evolving technologies are more important. The comparison within a set of counterfactuals where we fix just some of the inputs or just some of the technologies informs us which particular inputs and types of technological change matter the most. Finally we evaluate the implications of these channels for aggregate labor productivity growth.

4.1 The role of technological change

Figure 5 shows the average annual labor productivity growth in the three sectors over 1960-2017. The first set of bars is the actual data, which is perfectly reproduced by our baseline model, showing that the goods sector had with 2.5% the highest labor productivity growth, whereas it was 1.5% in low-skilled services and 0.7% in high-skilled services. The subsequent sets of bars show the results of various counterfactuals in which we fix some of the factor-augmenting technologies (the $Z_{fJ,t}$) at their 1960 values, but let inputs and other technologies vary over time as extracted from the data. Comparing the implied sectoral labor productivity growth (and their differences) to the data informs us about the importance of the technological change that we shut down, which is described below each set of bars.

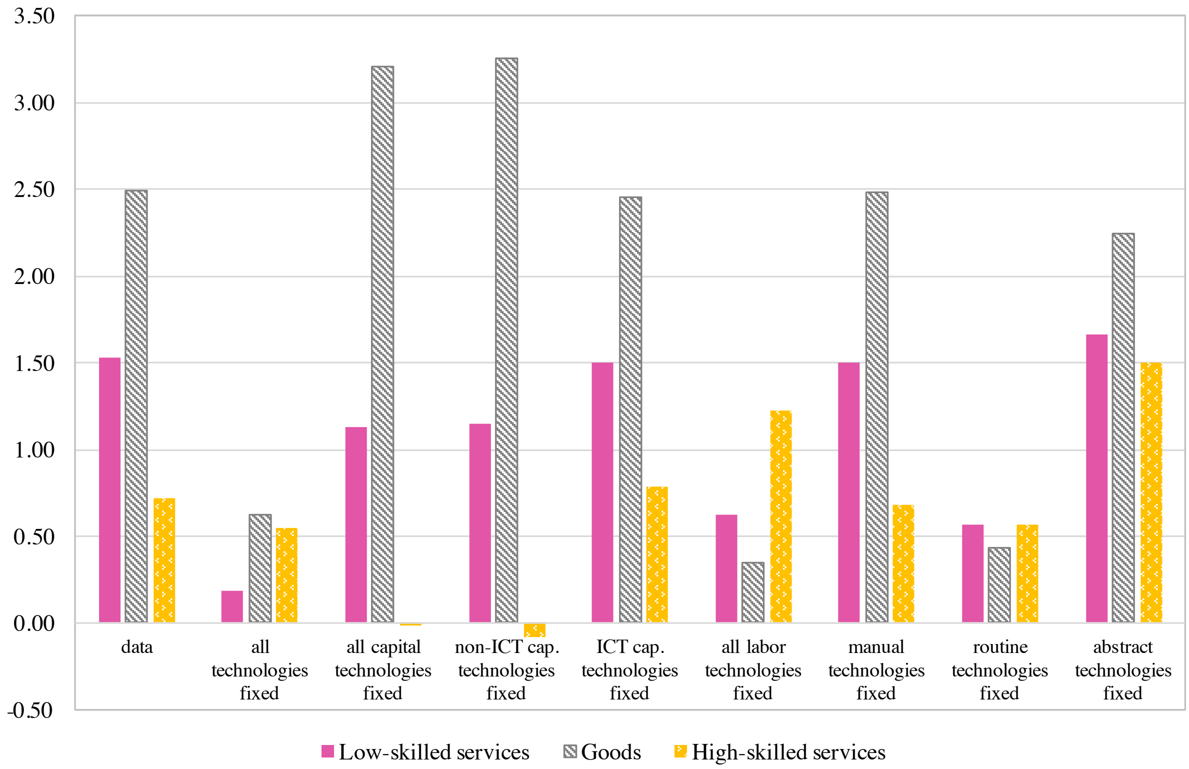

Figure 5: Average sectoral labor productivity growth with fixed technologies

Each set of bars shows the average annual labor productivity growth rate (in percent) over 1960-2017 for the three sectors (low-skilled services in pink solid, goods in gray striped, and high-skilled services in yellow patterned). The first set of bars shows the growth rates in the data, and the subsequent sets show counterfactual growth rates when holding technologies augmenting the factors listed below the bars at their 1960 level, with all inputs as well as all other technologies evolving as in the data.

Absent any change in factor-augmenting technologies, but just due to capital accumulation and employment reallocation, as the second set of bars ('all technologies fixed') shows, there is hardly any growth in labor productivity in low-skilled services and in goods production and only very small differences across sectors. Evolving technologies explained at least 75% of labor productivity growth in low-skilled services, 55% in goods and 24% in high-skilled services.[23] This clearly demonstrates that technological progress was crucial for the level of labor productivity growth as well as for its sectoral differences. High-skilled services seem to be somewhat of an exception; in this sector changes in factor input use must have been crucial.

To see whether this is due to capital-augmenting technological change, we next fix just the productivity parameters of ICT and non-ICT capital. Comparing the results of the third set ('all capital technologies fixed') to the data reveals that (sector-specific) capital-augmenting technological change has increased labor productivity growth in low- and high-skilled services, but lowered it in goods, thus acting to reduce differences in labor productivity growth across sectors. This demonstrates that capital-augmenting technological change was not the driver of the differences across sectors observed in the data. When distinguishing further between technological change in the two types of capital (the next two sets of bars), we see that these results are mainly driven by the evolution of non-ICT capital's productivity, and not by ICT capital. These results can be understood from Table 2 and Table 4. As the weight on ICT capital is very low in each sector, what happens to ICT capital augmenting technologies has little impact in all sectors. The weight on non-ICT capital is much higher, and thus the evolution of this type of technological change is more important. As non-ICT capital augmenting technologies decreased in the goods sector, fixing them at their 1960 level would imply an increase in labor productivity growth in goods, while the opposite holds in the other two sectors, with a much larger magnitude in high-skilled services since it has the highest growth in Z_{kJ} as well as the highest weight, ω_{kJ}.

In the last four counterfactuals we first fix all labor-augmenting technologies at their 1960 level, and then in turn fix only manual, only routine or only abstract labor augmenting technologies (within each sector). The results show that without any improvements in labor-augmenting technologies the magnitude of and the differences

[23] These numbers are the minimum of the fraction of the data predicted when fixing all inputs, and of one minus the fraction predicted when fixing all technologies.

between sectoral labor productivity growth would have been very far from the data. This highlights that technological change augmenting labor is key. We break this up further to study the role of technologies augmenting the various occupations. We find that routine labor augmenting technological change was a first-order determinant of labor productivity growth in low-skilled services and in goods, explaining at least 59 and 74 percent respectively. It explains at least 21 percent of labor productivity growth in high-skilled services, again pointing to the importance of changing factor input use in this sector, which we investigate in section 4.2.[24] Sector-specific routine labor augmenting technological change is also the single most important driver of sectoral differences; without it labor productivity growth would have been almost equalized across sectors. While changes in abstract labor augmenting technologies have contributed to sectoral differences in labor productivity growth to some extent, manual labor augmenting technologies hardly had any impact on the level of and on the differences in sectoral labor productivity growth. These results as well can be understood from Table 2 and Table 4. Both the weight on manual labor and the growth in technologies augmenting manual labor are quite low in all sectors, thus shutting down this technological change has only a small impact on labor productivity growth. Shutting down abstract labor-augmenting technologies has the largest impact in high-skilled services, as there the weight on and the change in the technology augmenting abstract labor are both large. The weight on routine labor, as well as the growth rate of technologies augmenting it, are high in all sectors, implying that shutting down routine-augmenting technological change has a large impact in all sectors. However, its impact is smallest in high-skilled services, which compared to the other two sectors has the lowest weight on this component and experienced the slowest growth in routine-augmenting technologies.

[24] To obtain these numbers we conducted an additional counterfactual, where we fixed everything at the 1960 level except for $Z_{rJ,t}$ which evolved as extracted from the data. We report the minimum of the fraction of the data predicted by this additional counterfactual, and of one minus the fraction predicted when shutting down only the change in $Z_{rJ,t}$ (the 'routine tech. fixed' counterfactual of Figure 5).

4.1.1 The role of sector and occupation components in labor-augmenting technological change

As we found such an important role for labor – and in particular for routine labor – augmenting technological change we investigate this further. In light of the sector and factor patterns visible in Table 2, we want to understand whether the effect of labor-augmenting technologies can be assigned to occupation-specific or to sector-specific components. We want to know, for example, where exactly the effects of sector-specific routine labor augmenting technological change are stemming from; is it the differences across sectors or the growth differential relative to the other occupations that is more important?

To decompose the changes of technologies augmenting labor in all sector-occupation cells, we set up a factor model.[25] In particular we regress the change in log cell technologies between each consecutive period on a (time-varying) sector effect ($\gamma_{J,t}$), an occupation effect ($\delta_{o,t}$), and a time effect (β_t) in the following way

$$\Delta \ln Z_{oJ,t} \equiv \ln Z_{oJ,t} - \ln Z_{oJ,t-1} = \beta_t + \gamma_{J,t} + \delta_{o,t} + \varepsilon_{oJ,t}, \qquad (24)$$

where we use weights $\Omega_{oJ,t}$ to reflect the relative importance of the sector-occupation cell.[26] We restrict both the average sector effect and the average occupation effect to be zero, which effectively implies that β_t captures the average labor augmenting technological change across all cells between period $t-1$ and t.[27]

Based on the results of (24), we compute counterfactual series for $\Delta \ln Z_{oJ,t}$, from (i) the neutral component alone ($\widehat{\beta}_t$), (ii) the neutral and sector-specific components ($\widehat{\beta}_t + \widehat{\gamma}_{J,t}$) which we call 'sector-only', (iii) the neutral and occupation-specific components ($\widehat{\beta}_t + \widehat{\delta}_{o,t}$) which we call 'occupation-only', and (iv) from all components (everything but $\widehat{\varepsilon}_{oJ,t}$), to which we refer as the 'full factor' prediction. In the appendix we show

[25] In macroeconomics factor models have been also used to study how country-level outcomes depend on sector and country factors, for instance in Stockman (1988), Ghosh and Wolf (1997) and Koren and Tenreyro (2007).

[26] The weights we use are the cells' average labor income between period $t-1$ and t, $\Omega_{oJ,t} = \frac{VA_{J,t}(1-\Theta_{J,t})\theta_{oJ,t} + VA_{J,t-1}(1-\Theta_{J,t-1})\theta_{oJ,t-1}}{\sum_{o,J}(VA_{J,t}(1-\Theta_{J,t})\theta_{oJ,t} + VA_{J,t-1}(1-\Theta_{J,t-1})\theta_{oJ,t-1})}$. The results are very robust to alternatives, such as using cell employment shares, or using year $t-1$ or year t shares, rather than averages.

[27] To be more precise these restrictions are: $\sum_o \sum_J \Omega_{oJ,t}\gamma_{J,t} = 0$ and $\sum_J \sum_o \Omega_{oJ,t}\delta_{o,t} = 0$ for every t.

in Figure A1 the path of sector-occupation technology changes over time as extracted from the data as well as those predicted from the various components.

To gauge how much of the variation in cell productivities the neutral, sector- and occupation-specific components can explain jointly and separately, we calculate the following *distance measure* between the extracted and the various predicted $\Delta \ln Z_{oJ}$:

$$D = \frac{\sum_{o,J,t} \Omega_{oJ,t}(\widehat{\Delta \ln Z}_{oJ,t} - \Delta \ln Z_{oJ,t})^2}{\sum_{o,J,t} \Omega_{oJ,t}(\Delta \ln Z_{oJ,t} - \overline{\Delta \ln Z})^2}.$$

This measure captures the variation in the extracted productivity changes that the various components cannot account for. It is always positive and the smaller it is, the closer the predictions are to the data. It is worth to note that this measure is closely related to the R^2, and in certain cases, including the 'full factor' and the 'neutral' prediction, it exactly equals $1 - R^2$.[28]

	neutral	**full factor**	**sector**	**occupation**
Distance measure	0.703	0.033	0.228	0.407

The above table shows the distance measure for the alternative series. It is immediately clear that the neutral prediction explains rather little of the variation (29.7 percent), while the full factor prediction explains almost all of the variation (96.7 percent) in the extracted technologies. The latter also implies that the part that is idiosyncratic to the sector-occupation cell accounts for only 3.3% of the variation. The distance measures of both the 'sector-only' and of the 'occupation-only' predictions are much larger than that of the 'full factor' prediction, whose explanatory power hence comes from both types of components.[29]

[28]The R^2 is defined as

$$R^2 = \frac{\sum_{o,J,t} \Omega_{oJ,t}(\widehat{\Delta \ln Z}_{oJ,t} - \overline{\Delta \ln Z})^2}{\sum_{o,J,t} \Omega_{oJ,t}(\Delta \ln Z_{oJ,t} - \overline{\Delta \ln Z})^2},$$

and $R^2 = 1 - D$ if the predictor is unbiased, $\sum_{o,J,t} \Omega_{oJ,t}\widehat{\Delta \ln Z}_{oJ,t} = \overline{\Delta \ln Z}$, and if the independent variables are uncorrelated with the error term, $corr(\Delta \ln Z - \widehat{\Delta \ln Z}, \Delta \ln Z) = 0$. These conditions only hold for the 'full-factor' and the 'neutral' series, and in these cases $D = 1 - R^2$.

[29]In appendix E.1 we conduct this analysis for a range of the elasticity of substitution between the occupational labor inputs. For larger values of ρ the distance measure of the neutral, the sector and the full factor component is larger, while of the occupation component it is smaller.

The results from this decomposition imply that the growth of labor-augmenting technologies is very well described as the sum of neutral, sector-specific and occupation-specific components. This holds not only in terms of explained variation of Z_{oJ}, but also for the components' contributions to sectoral labor productivity growth. Figure 6

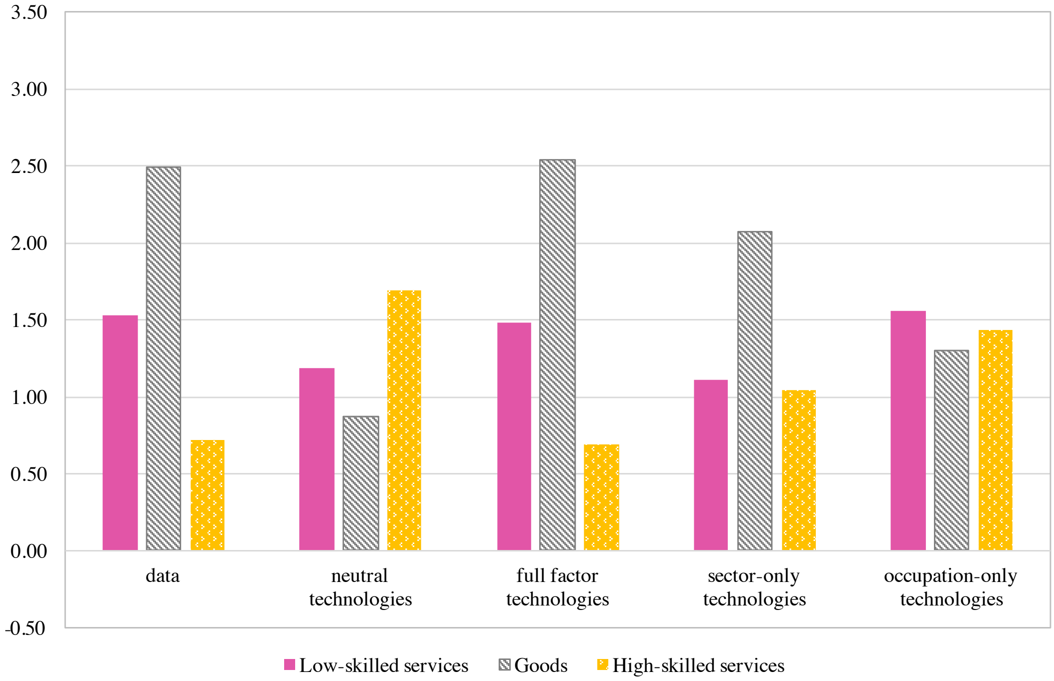

Figure 6: Role of occupation and sector components in sectoral labor productivity

Each set of bars shows the average annual labor productivity growth rate (in percent) over 1960-2017 for the three sectors (low-skilled services in pink solid, goods in grey striped, and high-skilled services in yellow patterned). The first set of bars shows the growth rates in the data, and the subsequent sets show growth rates when feeding in counterfactual labor-augmenting technologies obtained from (24) based on the components listed below the bars, with all inputs as well as all capital augmenting technologies evolving as in the data.

shows the results for counterfactuals that evaluate the role of the different components of labor-augmenting technological change for sectoral labor productivity growth. Here all inputs and capital-augmenting technologies evolve as in the data, but we feed in the counterfactual technologies based on the components listed below the bars. In this bar chart, the closer is a set of bars to the data, the better the given component explains the growth rates of sectoral labor productivity, and the less important are the omitted components. Not surprisingly, neutral labor-augmenting technological change does not reproduce the data. The counterfactual based on the 'full factor' prediction, on the other hand, replicates the observed labor productivity growth rates well. The last two counterfactuals show that neither the sector nor the occupation components by themselves are enough to generate all aspects of the data. The occupation component alone fails to generate the level and the differences of growth rates across sectors, whereas the sector component alone gets closer in terms of these aspects but quantitatively falls short. The fact that the occupation-only model does not reproduce sectoral labor productivity growth differences suggests that differences in technological change across sectors are important, supporting the mechanism of Ngai and Pissarides (2007). Occupation-specific components of technological change (emphasized in Duernecker and Herrendorf (2016) and Lee and Shin (2017)) – once accounting for the possibility of sector-specific changes – alone are not enough. Overall this analysis reveals that both sector and occupation components are important drivers of labor productivity growth at the sectoral level.

4.2 The role of changing input use

We now turn our attention to the role of production factors. Figure 7 shows the results of various counterfactuals in which we fix some inputs at their 1960 values, but let all other inputs and the factor productivities ($Z_{fJ,t}$) vary over time as extracted from the data between 1960 and 2017. The last set of bars shows a different type of counterfactual. Here we assign in each period identical occupational employment shares to all sectors, thus both their initial level and their evolution over time is the same. Again comparing the results implied by the counterfactual to the actual data gives a sense of the importance of the changing use of the fixed input(s).

In the second set of bars ('all inputs fixed') we fix all inputs at their 1960 level. Keeping all inputs at their initial level results in lower labor productivity growth in all sectors. This implies that the reallocation of labor and the accumulation of capital had a positive effect on labor productivity growth in all three sectors. While the size of this effect varied across sectors, the ranking of sectors in terms of labor productivity growth was not affected by changing input use. However, absent capital accumulation and employment reallocation across sector-occupation cells, there would have been hardly any difference between the productivity growth in goods and in low-skilled services. This highlights that changing input use is important for the level of labor

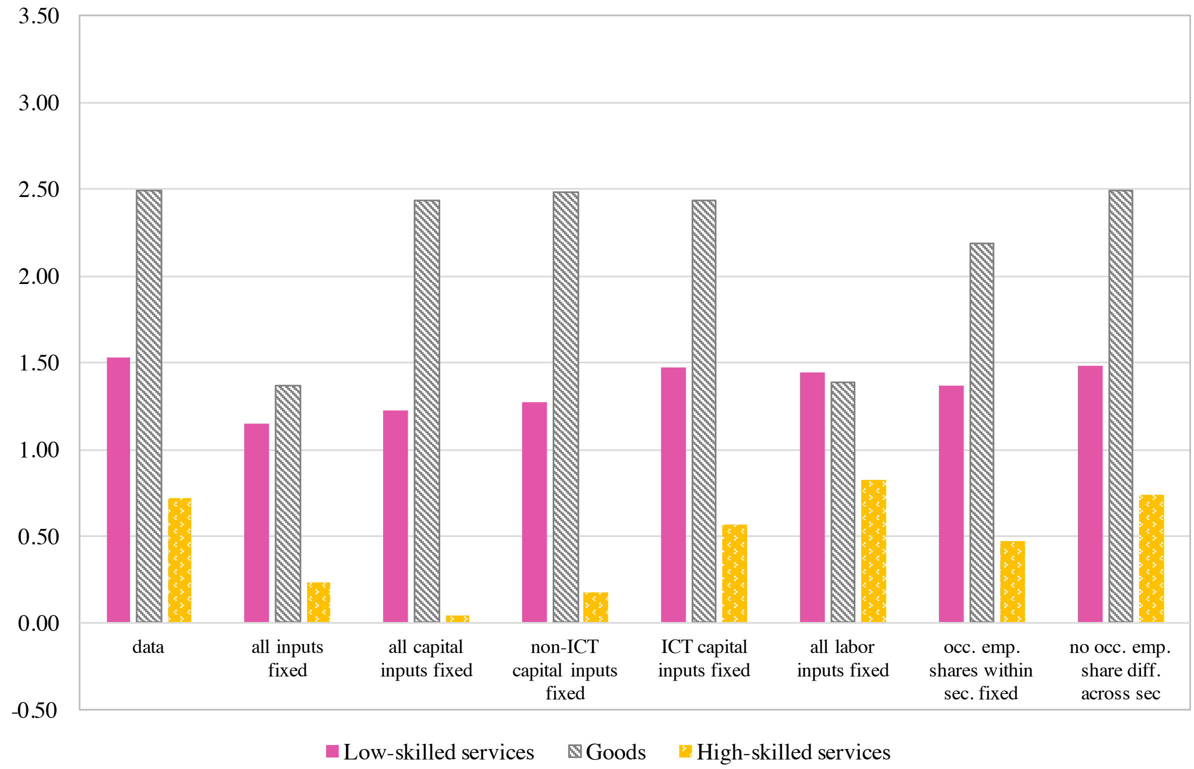

Figure 7: Average sectoral labor productivity growth with fixed factor inputs

Each set of bars shows the average annual labor productivity growth rate (in percent) over 1960-2017 for the three sectors of the economy (low-skilled services in pink solid, goods in gray striped, and high-skilled services in yellow patterned). The first set of bars shows the growth rates in the data, and the subsequent sets show counterfactual growth rates when holding the inputs listed below the bars at their 1960 level (or share), and with all other inputs as well as technologies evolving as in the data. In the last set of bars we assign identical occupational employment shares to all sectors and let everything else evolve as in the data.

productivity growth, as well as for its differences across sectors. On the other hand, comparing these results to those of fixed technologies in Figure 5 (second set of bars) highlights that evolving technologies matter much more than changing inputs for sectoral growth rates, except in the high-skilled service sector. In this sector changing input use is more important.

The next three counterfactuals shed light on the role of capital accumulation. With both types of capital inputs fixed at their 1960 level ('all capital inputs fixed'), the growth rate in all sectors falls short of the data, on average by 39 percent.[30] This effect is the most pronounced in high-skilled services, where absent capital accumulation there would have been hardly any growth in labor productivity. Capital accumula-

[30] Labor productivity growth in low-skilled services would have been 80% of its actual value, in goods 98%, and in high-skilled services 5%, the simple average of this is 61%, i.e. 39 % lower than in the data.

tion resulted in smaller sectoral differences in labor productivity growth, but without altering the ranking of sectors. This suggests that capital deepening, which was differential across sectors, was important for the level of labor productivity growth, but was not the main driver of sectoral differences. In particular, if capital deepening was the source of structural transformation, as argued in Acemoglu and Guerrieri (2008), then shutting it down should result in a larger reduction in productivity growth in the goods sector compared to services, which is not what we find. Comparing the counterfactual where we shut down only non-ICT capital with the one where we shut down only ICT capital accumulation shows that non-ICT capital had a larger and less uniform effect on labor productivity growth across sectors. These results can be understood from Table 3, which shows the gross growth of each input between 1960 and 2017, and Table 4, which shows the weight that a change in any effective input carries. While ICT capital increased tremendously, its weight is very small in each sector, explaining the relatively small impact it had on labor productivity growth. Both the weight on and the increase in non-ICT capital input was the largest in high-skilled services and the smallest in goods. This explains why the impact of fixing non-ICT capital at its 1960 level has a very small impact on labor productivity growth in goods and a much larger impact in high-skilled services.

In the last three counterfactuals we study the role of labor allocation across sector-occupation cells. We first fix all labor inputs at their 1960 values ('all labor inputs fixed'). The resulting productivity growth rate falls considerably short of the data in goods, in low-skilled services only marginally, whereas in high-skilled services it is slightly higher than in the data. Hence, absent employment reallocations, sectoral differences in labor productivity growth are not in line with the data. Overall this highlights that changing labor use was important for the level of growth in goods and for sectoral differences in labor productivity growth. To understand where these results come from note that fixing all labor inputs at the 1960 levels shuts down two margins: (i) changes in total sectoral employment and (ii) reallocations between occupational employment within sectors.

We investigate this second channel in the penultimate set of bars where we fix the share of occupations within each sector at initial ratios ('occ. emp. shares within sec.

fixed') but let the overall employment share of each sector (as well as all other inputs and technologies) evolve as in the data. This counterfactual yields growth rates that are lower than, but quite close to, the actual data. This shows that shifts in the occupational employment structure within sectors had positive, yet only very modest, effects on sectoral labor productivity growth, but hardly any effect on sectoral differences. In contrast, changes in overall sectoral employment, the first channel, play a large role and are thus the dominating channel in the counterfactual fixing all labor inputs. As can be seen in Table 3, the total employment share in goods declined to less than half the initial value, while it increased by about 10 and 60 percent in low- and high-skilled services respectively. These changes clearly impact both the numerator (sectoral output) and the denominator (sectoral employment) of labor productivity, but –holding the other inputs fixed– have a larger impact on the denominator. Thus, shifts in employment from goods to services increased labor productivity growth in goods and lowered it in the two service sectors.

In the last set of bars we impose the same occupational structure in each sector, which we let evolve in the same way as the occupational composition of the aggregate economy. The results of this counterfactual hardly differ from the data. This implies that the differences in occupational intensities across sectors did not generate, nor contribute to, the sectoral differences in labor productivity growth observed in the data. Our findings from the last two counterfactuals make it unlikely that occupational productivity growth differences and sectoral differences in occupational employment structure are driving structural transformation, as implied by Duernecker and Herrendorf (2016), Lee and Shin (2017).

To summarize our findings so far, both changing inputs and changing technologies have been important for the observed sectoral labor productivity growth, with technologies playing a larger role. We find that both capital accumulation and capital-augmenting technological change acted to reducing differences in labor productivity growth across sectors. When isolating the effects of changing technologies by production factors, we see that labor-augmenting technological change had the largest role, and in particular (sector-specific) routine labor-augmenting technological change.

4.3 Implications for aggregate labor productivity growth

Table 5: Average annual growth rate of Zs between 1960-1990 and 1990-2017

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
1960-1990					
Low-skilled services	0.32%	1.08%	-1.51%	2.86%	3.38%
Goods	-2.40%	3.69%	-1.73%	0.63%	4.97%
High-skilled services	-0.79%	-1.18%	-4.13%	3.94%	-3.47%
1990-2017					
Low-skilled services	0.19%	5.01%	0.23%	-1.33%	0.54%
Goods	4.01%	7.78%	4.08%	-4.04%	3.84%
High-skilled services	2.43%	4.18%	-0.40%	-0.58%	-0.21%

We established that while capital accumulation was important for the level of labor productivity growth, especially in the high-skilled service sector, technological change seems to have been a more important determinant of both the level of and the sectoral differences in labor productivity growth. We also showed that the key driver was sector-specific routine-augmenting technological change. In what follows we study whether these findings hold for labor productivity growth in the aggregate economy. In addition, we investigate whether the importance of the various drivers changed over time. First, however we show in Table 5 the average annual growth rate of sector-factor augmenting technologies for two sub-periods, 1960-1990 and 1990-2017. The sector-specific and factor-specific patterns are very similar to those shown in Table 2. Comparing the earlier to the more recent period shows that technological change augmenting each type of labor accelerated over time (for all occupations in all sectors but for Z_{mL}, the growth rate of which remained virtually constant), while technological change augmenting either type of capital decelerated (except for ICT-augmenting capital in high-skilled services). This suggests that the relative importance of capital- vs labor-augmenting technologies for labor productivity growth has changed over the last decades.

In Figure 8 we show average annual labor productivity in the whole economy between 1960-2017 and in two sub-periods, 1960-1990 and 1990-2017 in the data and for several counterfactuals. Note that a larger difference between data and counterfactual implies a larger role for the component that we shut down. Comparing the 'all inputs

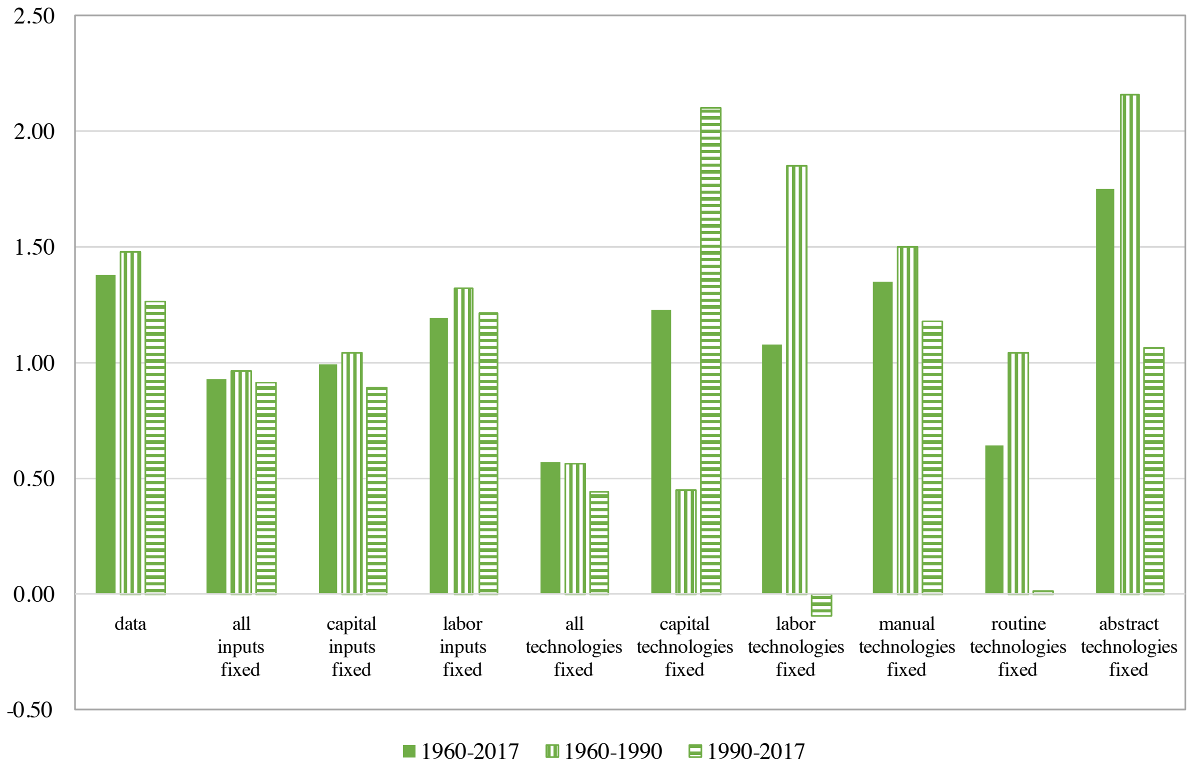

Figure 8: The role of inputs and technology in aggregate labor productivity growth

Each set of bars shows the average annual labor productivity growth rate of the economy (in percent) over 1960-2017 (solid), and over two sub-periods, 1960-1990 (vertically striped) and 1990-2017 (horizontally striped). The first set of bars shows the growth rates in the data, and the subsequent sets show counterfactual growth rates when holding the inputs or the technologies augmenting the factors listed below the bars at their initial level, and allowing all other inputs as well as all other technologies to evolve as in the data.

fixed' and the 'all technologies fixed' counterfactual with the data, it is evident that technological change was more important for labor productivity growth than changing input use for the entire period (with technologies explaining at least 59%, and inputs at least 33%), as well as for both sub periods. In terms of input use, capital accumulation ('capital inputs fixed') played a larger role than labor reallocation across sector-occupation cells ('labor inputs fixed'). In particular, without capital accumulation labor productivity growth would have been about a third lower over the entire period, as well as in both sub-periods. This is in line with the findings of the growth accounting literature for GDP *per worker* (i.e. labor productivity) growth, see for example Aghion and Howitt (2007).

Analyzing the effect of different technologies for aggregate labor productivity growth it becomes clear that capital augmenting technologies were more important between

1960-1990, while labor augmenting technologies played a larger role in 1990-2017.[31] Finally, looking at the respective role of sector-specific technologies augmenting the three occupations, routine-augmenting technological change stands out as the one contributing the most to aggregate labor productivity growth, explaining at least 54% of labor productivity growth. Moreover, its role became substantially more pronounced over the time period studied. Absent routine labor augmenting technological change growth would have been about 30% lower between 1960 and 1990, while between 1990 and 2017 there would have been hardly any growth.

5 Robustness checks and extensions

In this section we show that our results are very robust to alternative values for the various substitution elasticities, including a model variant with sector-specific substitution elasticity between capital and the labor aggregate. Furthermore, we establish that our main results are robust to changing the production function to a different nesting of inputs. We also describe how to control for observable worker characteristics in our framework and demonstrate that our conclusions are robust to accounting for worker efficiency. Finally, we note that the results are robust to including the self-employed in the analysis and that the results for the aggregate goods sector are very similar to the one for an industry sector that does not include agriculture.

5.1 Alternative substitution elasticities

So far we showed results from our framework based on three elasticities, $\sigma = 0.84$, $\sigma_c = 2$ and $\rho = 0.6$. In this subsection we briefly summarize how our results are affected when we change these elasticities to alternative values. We consider all combinations of $\sigma \in \{0.65, 0.75, 0.84\}$, $\sigma_c \in \{1.5, 2, 2.5\}$ and $\rho \in \{0.5, 0.6, 0.7\}$.[32] The general

[31] That in the period 1960-1990 labor productivity growth would have been higher absent labor augmenting technological change, and that between 1990-2017 it would have been higher without capital augmenting technological change, reflects the numbers smaller than 1 in Table 2.

[32] As discussed in section 3.3 the majority of studies finds values less than one for the elasticity of substitution between capital and labor, and our baseline of $\sigma = 0.84$ is towards the upper end of estimates. Therefore, we discuss how our results change with lower values, 0.75 and 0.65. For σ_c we consider two alternative values: 1.5 in the midrange of values calibrated in Aum et al. (2018), and 2.5 in the midrange

conclusions are that all of our results are extremely robust. It is important to keep in mind that our baseline framework under any parameterization matches all data targets perfectly. As such, alternative values for these elasticities of substitution lead to different series of the inferred technologies. We show in Table A3 in the appendix the range of the annual growth rates of the various extracted sector-specific factor-augmenting technologies for all combinations of the different elasticities that we consider. This table shows that the general patterns described in section 4 for Table 2 remain the same, with similar patterns arising both across factors and across sectors.[33]

Looking at equations (21) and (22) and (23), it is easy to understand why this is the case. In all of these equations what changes on the right hand side are the elasticities, and the only thing that adjusts on the left hand side are the technology terms, the Zs. For the growth in non-ICT capital augmenting technologies, as seen in (21) the key is the sign of $\sigma - 1$, which is negative for all σ values considered, thus implying similar results. For manual and abstract labor augmenting technologies, which are pinned down by (22) (and its analogue), the signs of $\rho - 1$ and of $\sigma - 1$ matter. For routine labor-augmenting technologies and ICT capital-augmenting in addition also $\rho - \sigma$ all matter, as can be seen in (23) (and its analogue). The sign of the first two is negative for all elasticities considered, while the sign of $\rho - \sigma$ is mostly negative, except for the case when $\rho = 0.7$ and $\sigma = 0.65$. Yet since in all sectors the changes in the LA-share is small compared to the changes in the occupational labor income shares and the RA-share (see Table A2), changes in the sign of $\rho - \sigma$ have only a limited effect on our results. This discussion demonstrates that beyond the specific elasticity values considered in this robustness analysis, our results are robust to a wider range of elasticities. These ranges encompass most of the values which have been considered in the literature.

Figures A2, A3, A4 and A5 are analogues to our main result figures. We show for each counterfactual the ranges of labor productivity growth implied by all combinations of elasticities considered. Most of our results qualitatively do not depend on the parameterization of the elasticities, and are even quantitatively very similar, except

of the values implied by the estimation in Eden and Gaggl (2018). For ρ we consider a value below and one above the baseline value, but below one, thus continuing to assume complementarity between occupational labor inputs.

[33] In fact, beyond the ranges shown in Table A3, across the different sets of parameterizations even the ranking of technology growth rates within each sector is by large the same.

for the role of sector- vs. occupation-components. As Table A4 in the appendix shows, the larger is ρ,[34] the larger is the distance measure both of the full factor and of the sector-only technologies, and the smaller is the distance measure of the occupation-only technologies. For larger elasticities, the ranking of sectors in terms of labor productivity growth under 'sector-only' technologies is less in line with the data, and under 'occupation-only' technologies it is more in line with the data. Thus we find that the respective role of sector- and occupation-components is sensitive to this elasticity, but the observation that we need both type of components to match the data holds for all elasticities, as Figure A3 demonstrates.

5.1.1 Sectoral heterogeneity in elasticities between capital and labor

We also consider a model variant where the elasticity of substitution between capital and the labor aggregate differs across sectors, as papers estimating this elasticity have found differences across industries (e.g. Oberfield and Raval (2014), Lawrence (2015)). Most papers focus however only on non-service industries. One exception is Herrendorf et al. (2015) which finds 0.75 for services. As such we set $\sigma^L = \sigma^H = 0.75$ for both of our service sectors. Our goods sector contains both agriculture and manufacturing, therefore we set a value of $\sigma^G = 0.9$, in between their estimates of 0.8 for manufacturing and 1.58 for agriculture.

As we infer the technologies by sector and we just showed that our results are robust to altering the common σ parameter, one should not expect large differences compared to our baseline. Indeed the table mimics the patterns of our baseline quite closely. Appendix Figure A6 compares the effects of the various channels with those in the baseline for sectoral labor productivity growth, showing that our results are very robust. The only noticeable difference is quantitative: the effect of technologies is somewhat more pronounced, and in particular with sector-specific elasticities it seems that the role of labor-augmenting technologies in aggregate labor productivity growth is slightly larger.

[34]Note, changing the value of ρ does not affect the growth rate of the Z_{kJ}s at all.

5.2 Different Nesting

In our baseline we assumed that manual and abstract workers and the routine aggregate are combined in a symmetric fashion in the sectoral production. As a further robustness check we consider a different nesting, proposed by vom Lehn (2019), where the routine aggregate is first combined with abstract labor according to elasticity of substitution ρ_a, giving the *abstract aggregate*:

$$AA_{J,t} = \left((Z_{aJ,t} l_{aJ,t})^{\frac{\rho_a - 1}{\rho_a}} + RA_{J,t}^{\frac{\rho_a - 1}{\rho_a}} \right)^{\frac{\rho_a}{\rho_a - 1}}, \tag{25}$$

and the abstract aggregate is then combined with manual labor under a constant elasticity of substitution ρ_m, which gives the *labor aggregate*:

$$LA_{J,t} = \left((Z_{mJ,t} l_{mJ,t})^{\frac{\rho_m - 1}{\rho_m}} + AA_{J,t}^{\frac{\rho_m - 1}{\rho_m}} \right)^{\frac{\rho_m}{\rho_m - 1}}. \tag{26}$$

This labor aggregate is then combined with non-ICT capital exactly as in our baseline. Based on this production structure, we repeat our exercise, following the steps described in appendix E.2 and setting the additional elasticities $\rho_a = 0.31$ and $\rho_m = 1.49$, following from vom Lehn (2019).

While the growth rates for factor-augmenting technologies under this alternative nesting naturally are different (see Table A6), their implications in the counterfactuals are very similar to our baseline, both for sectoral and for aggregate labor productivity growth, as shown in Figure A7. While there are some small quantitative differences, for example shutting down routine-augmenting technological change implies a larger reduction in labor productivity growth in the goods sector but a smaller reduction in the two service sectors, the qualitative predictions are exactly the same. This demonstrates that our conclusions are very robust to alternative assumptions about the production structure.

5.3 Allowing for efficiency units of labor in production

In our baseline framework we measured occupational labor inputs as (shares of) hours worked, implicitly assuming that all workers are equally efficient, both within and across periods. A potential concern with this setup is that the evolution of workers' human capital over time might confound the growth rates of technologies that we inferred. To address this, we estimate each worker's efficiency units from a Mincer log wage regression on worker characteristics, including a polynomial in potential experience, education, gender and race, using the IPUMS Census/ACS data. From the estimates we construct average efficiency units of labor in each sector-occupation cell, $\overline{e}_{oJ,t}$ and wages per efficiency units of labor, as we explain in appendix E.3.[35]

To incorporate efficiency units of labor into the model, we assume that firms choose $n_{oJ,t} \equiv \overline{e}_{oJ,t} l_{oJ,t}$ in each period, instead of just hours worked ($l_{oJ,t}$). This implies that we need to use wages per efficiency unit of labor in equations (15) to (20) to infer sector-factor technologies, whereas the measurement of all other variables remains unchanged.

Figure A8 in the appendix plots the alternative series for the relative wages within sectors. The resulting patterns for relative occupational wages within a sector are very similar,[36] whether accounting for efficiency units or not, though their levels are somewhat different. Since we identify the within-sector ratios of occupational productivities precisely from these relative wages, the general conclusions about the inferred technological change are very similar, as shown in Table A8. Given that the series of the factor-augmenting technologies (by sector) in the model with efficiency units of labor are so similar to the baseline model, and in fact for the capital inputs coincide, the implications for sectoral labor productivity are very similar too. Figure A9 in the appendix shows the role of individual inputs and technologies in this model variant alongside the baseline results. While there are very small quantitative differences, qualitatively they have the very same implications.

[35]We construct this in two different ways, by including/not-including the residuals from the Mincer wage regression in $\overline{e}_{oJ,t}$. Note that, even though we calculate sector-occupation wage rates from our accounting identity (see equation (41) in the appendix) as before, the relative wages within sectors are the same as those implied by the the Mincer wage regression.

[36]From 2000 onwards, in high-skilled services there is somewhat of a divergence between relative average ('raw') wages and relative wages controlling for workers' characteristics.

5.4 Further Robustness Checks

Finally we establish that our results are robust to handling two aspects of the data differently. First, we show in appendix E.4 that including self-employed in the analysis, which in our baseline were dropped from the micro data on employment shares and wages and excluded from the computation of the sectoral labor income shares, does hardly matter for the results (see Figure A11 in particular). Second, in appendix E.5 we break up the broad goods sector and differentiate between agriculture and 'industry', i.e. manufacturing, mining and construction. Here we find that the results for the engines of labor productivity growth in industry are very similar to what we found for the overall goods sector, while the agricultural sector looks somewhat different with a much larger role for improvements in non-ICT capital technology. However, as the results for industry are so similar to those for goods and agriculture accounts only for a very small share of goods and of GDP, the implications for aggregate labor productivity growth allowing for a separate agricultural sector are virtually identical, again pointing to the robustness of our main results.

6 Conclusion

In this paper we analyze the drivers of sectoral labor productivity growth in the United States over 1960–2017, combining detailed Census/ACS data with sectoral data from the BEA and EU KLEMS. We propose and implement a novel approach to extract sector-specific factor-augmenting technologies from observed changes in factor prices, factor shares, value added shares and sectoral growth in real value added over time. Key in our approach is that we distinguish between occupational labor inputs and that we do not impose a priori assumptions about whether technological change occurs at the sector or at the factor level. Our results show that the growth rates of factor-augmenting technologies differed not only across the various occupations and types of capital, but also for given production factors across sectors. Had we not taken this very flexible approach of allowing technologies to evolve at the sector-factor level, we would not have been able to identify these patterns.

Through a range of counterfactual exercises we find that most of labor productivity growth, both at the sector level and in the aggregate, was due to technological change. In particular we show that sector-specific routine labor augmenting technological change was crucial, explaining at least 54% of labor productivity growth in the aggregate. Changing occupational employment shares within sectors and capital accumulation both had a positive effect on the level, but neither contributed to the sectoral differences in labor productivity growth observed in the data. Furthermore, differences in occupational structure across sectors did not explain any of the sectoral patterns of labor productivity growth.

While we establish that the rate at which occupational labor-augmenting technologies evolved differs both across sectors and occupations, we also identify common components using a factor model. We find that occupation and sector components jointly explain 96.7 percent of occupational labor-augmenting technological changes, and that in measured sectoral labor productivity growth both components of technological change are crucial. One implication of this finding is that the growth rate of sector-occupation technologies is well approximated by the sum of the relevant sector- and occupation-component.

Overall, our results highlight that sector-specific routine-augmenting technological change has been the key determinant of labor productivity growth over 1960-2017 in the US economy, and that its contribution has accelerated in more recent decades.

Our finding that occupation-specific technological change varies across sectors is novel. As such there are no theories for this, but we believe there are at least three possible, complementary, explanations. First, as we discussed in the introduction, an occupation's productivity and its evolution may very naturally depend on the sector of work. Second, sectoral differences in firm size or organizational structure might result in differential effects of new technologies across sectors. Finally, as we consider relatively broad occupational categories, there still might be some compositional differences across sectors left in terms of finer occupational categories.

In this paper we did not investigate the reasons for sectoral differences in occupation-augmenting technologies, but rather evaluated their role in labor productivity growth. Our analysis highlights the need to better understand why routine labor-augmenting technologies have been growing at different rates across sectors.

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.

Philippe Aghion and Peter Howitt. Capital, innovation, and growth accounting. *Oxford Review of Economic Policy*, 23(1):79–93, 2007.

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

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. Biased Technological Change and Employment Reallocation. Working paper, Sciences Po and University of Kent, October 2019.

Kirsten Jäger (The Conference Board). EU KLEMS Growth and Productivity Accounts 2017 release. – Description of Methodology and General Notes, September 2017.

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

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

Francesco Caselli. Chapter 9 Accounting for Cross-Country Income Differences. volume 1 of *Handbook of Economic Growth*, pages 679 – 741. Elsevier, 2005.

Francesco Caselli. *Technology Differences over Space and Time*. Princeton University Press, 2016.

Francesco Caselli and Wilbur John II Coleman. The World Technology Frontier. *American Economic Review*, 96(3):499–522, June 2006.

Margarida Duarte and Diego Restuccia. The Role of the Structural Transformation in Aggregate Productivity. *The Quarterly Journal of Economics*, 125(1):129–173, 2010.

Margarida Duarte and Diego Restuccia. Relative Prices and Sectoral Productivity. *Journal of the European Economic Association*, forthcoming, 2019.

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

Georg Duernecker, Berthold Herrendorf, and Ákos Valentinyi. Structural Change within the Service Sector and the Future of Baumol's Disease. Working Paper DP12467, CEPR, 2017.

Maya Eden and Paul Gaggl. On the Welfare Implications of Automation. *Review of Economic Dynamics*, 29:15–43, July 2018.

Michael W L Elsby, Bart Hobijn, and Ayşegül Şahin. The decline of the U.S. labor share. *Brookings Papers on Economic Activity*, (FALL 2013):1–52, 2013.

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.

Douglas Gollin. Getting income shares right. *Journal of Political Economy*, 110(2):458–474, 2002.

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

Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi. Two Perspectives on Preferences and Structural Transformation. *American Economic Review*, 103(7):2752–89, 2013.

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

Karabarbounis, Loukas and Neiman, Brent. The global decline of the labor share. *The Quarterly Journal of Economics*, 129(1):61–103, 2014.

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

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

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

Robert Z Lawrence. Recent Declines in Labor's Share in US Income: A Preliminary Neoclassical Account. Working Paper 21296, National Bureau of Economic Research, June 2015.

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.

Miguel A. León-Ledesma, Peter McAdam, and Alpo Willman. Identifying the Elasticity of Substitution with Biased Technical Change. *American Economic Review*, 100(4), September 2010.

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

Ezra Oberfield and Devesh Raval. Micro Data and Macro Technology. Working Paper 20452, National Bureau of Economic Research, September 2014.

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

U.S. Bureau of Economic Analysis. Proprietors' income with inventory valuation and capital consumption adjustments: Farm [b042rc1a027nbea]. *retrieved from FRED, Federal Reserve Bank of St. Louis*, 2020a.

U.S. Bureau of Economic Analysis. Nonfarm proprietors' income by industry. *retrieved from U.S. Bureau of Economic Analysis*, 2020b.

Christian vom Lehn. Labor market polarization, the decline of routine work, and technological change: A quantitative analysis. *Journal of Monetary Economics*, 01 2019.

Appendix

A Classification

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.

We combine four different industry classification systems, the NAICS, the SIC, the ISIC and the IND1990. Table A1 summarizes our categorization in terms of each system.

B Data Appendix

Capital targets. To back out all Zs we need the rental rate of non-ICT, R_k, and of ICT capital, R_c, the share of income going to both types of capital, Θ_J, and to ICT capital alone, Θ_{cJ}, as well as the amount of non-ICT capital in each sector, k_J. As discussed in the main text, we obtain the labor income share in each sector, $1 - \Theta_J$, from the BEA as the compensation of employees over gross value added. Starting from data on current-cost net stock and quantity indices for fine capital categories from the BEA, we calculate for non-ICT and ICT capital real quantity (q_k and q_c) and price indices (p_k and p_c) using the cyclical expansion procedure. Due to the quantity index normalization of the BEA, these are both normalized to be 1 in 2009. Thus, we assume that the real quantity of non-ICT and ICT capital in 2009 is equal to the share of non-ICT and ICT

Table A1: Classification of industries into sectors

	NAICS	SIC	ISIC	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; Repair of motor vehicles & motorcycles (G) • Transportation & storage (H) • Arts, entertainment & recreation (R) • Accommodation & food serv. activities (I) • Other serv. activities (S) • Activities of households as employers; undifferentiated goods- & serv. producing activities of households for own use (T)	• 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 (A) • Mining & quarrying (B) • Construction (F) • Total manufacturing (C)	• 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	• Electricity, gas & water supply (D-E) • Information & communication (J) • Financial & insurance activities (K) • Real estate activities (L) • Professional, scientific, technical, Administrative & support serv. activities (M-N) • Education (P) • Health & social work (Q) • Public administration & defence; • compulsory social security (O)	• Utilities • Communications • Finance, insurance & real estate • Professional serv. High-skilled business serv. • Public administration

capital in the current-cost net stock of capital in 2009. Multiplying these 2009 values with the quantity indices (q_k, q_c) we get the time series of the real quantity of non-ICT and ICT capital. Dividing both by the number of full-time equivalent workers we get the model equivalent of k and c. We calculate annual depreciation rates for both types of capital δ_k and δ_c from the BEA data by dividing the sum of current-cost depreciation of fixed assets of all non-ICT (or ICT) capital with the sum of current cost net stock of these same fixed assets. The depreciation rate of non-ICT capital is fairly stable at around 5.5 percent annually, whereas of ICT capital the depreciation rate increases from 15.5 percent to 28 percent.

Nominal sectoral value added multiplied by the sector's capital income share should be equal to the value of total sectoral capital income. This results in the following accounting identity:

$$R_k k + R_c c = Y^{nom} \cdot \sum_J VA_J \Theta_J, \tag{27}$$

where Y^{nom} denotes nominal GDP per full-time equivalent worker, VA_J is sector J's nominal value-added share, and Θ_J is sector J's capital income share, all obtained from the BEA. Furthermore, we assume a no-arbitrage condition on the rate of returns to non-ICT and ICT capital:

$$\frac{R_c + (1-\delta_c)p_c'}{p_c} = \frac{R_k + (1-\delta_k)p_k'}{p_k},$$

where p_k' denotes the price of non-ICT capital in the next year. From these two equations we can calculate in each period the rental rates of non-ICT and of ICT capital, R_k and R_c.

We calculate the allocation of ICT capital across sectors from EU KLEMS between 1970 and 2015, as the share of nominal capital stock in millions of national currency in each sector, $\tilde{c}_J$, with $\sum_J \tilde{c}_J = 1$. The amount of real ICT capital (per worker) in each sector is then obtained as $c_J = c \cdot \tilde{c}_J$. The share of income going to ICT capital in each sector, Θ_{cJ}, is then pinned down by the accounting identity: $R_c c_J = Y^{nom} \cdot VA_J \Theta_{cJ}$. The amount of non-ICT capital in each sector, k_J, can then be calculated from (27).

Sector-occupation cell wages. In our quantitative model, we use workers' self-reported income in the Census/ACS to compute θ_{oJ} as in (1), but do not use it to calculate hourly wages. Instead we use an accounting identity to back out wages. This is to ensure that in the model the sum of all factor income is equal to value added, which we get from the BEA data. Nominal sectoral value added multiplied by the sector's labor income share should be the value of total sectoral labor income. This income in turn is split across the various occupations. The accounting identity therefore is that labor income of occupation o workers in sector J satisfies

$$w_{oJ}l_{oJ} = Y^{nom} \cdot VA_J(1 - \Theta_J)\theta_{oJ}, \tag{28}$$

where Y^{nom}, VA_J and Θ_J are as defined earlier, and θ_{oJ} denotes the share of sector J labor income that occupation o workers earn. Note that within sectors relative wages depend only on the relative θs and occupational employment shares, and therefore is equal to the relative wage observed in the Census/ACS data.

Income shares. Table A2 contains by sector the income share of each factor from sectoral value added in 1960 and in 2017.

Table A2: Share of income in value added in 1960 and in 2017

	occupations			capital		aggregates	
	manual	routine	abstract	non-ICT	ICT	LA	RA
1960							
Low-skilled services	0.066	0.399	0.119	0.413	0.002	0.587	0.402
Goods	0.007	0.496	0.126	0.368	0.002	0.632	0.499
High-skilled services	0.050	0.201	0.228	0.501	0.019	0.499	0.220
2017							
Low-skilled services	0.094	0.264	0.200	0.425	0.017	0.575	0.281
Goods	0.012	0.270	0.219	0.483	0.015	0.517	0.286
High-skilled services	0.041	0.103	0.392	0.429	0.035	0.571	0.138

C Derivations

In this subsection we show how the Zs can be expressed as a function of observables. In the first step we show the derivation of Zs within a period, and hence we omit the time subscripts. In the main text we showed the derivation of Z_{mJ}/Z_{aJ} and Z_{cJ}/Z_{rJ}. Here we show the derivation of Z_{mJ}/Z_{rJ} and Z_{kJ}/Z_{mJ}.

In these derivations we repeatedly use that at the optimum relative effective input use can be expressed as

$$\frac{Z_{cJ}c_J}{Z_{rJ}l_{rJ}} = \left(\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right)^{\frac{\sigma_c}{\sigma_c-1}}, \tag{29}$$

which follows from multiplying the relative optimal input use with the relative Zs. Using the above expression implies that at the optimum we can express the routine aggregate as:

$$RA_J = \left[(Z_{rJ}l_{rJ})^{\frac{\sigma_c-1}{\sigma_c}} + (Z_{cJ}c_J)^{\frac{\sigma_c-1}{\sigma_c}}\right]^{\frac{\sigma_c}{\sigma_c-1}} = Z_{rJ}l_{rJ}\left[1+\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right]^{\frac{\sigma_c}{\sigma_c-1}}. \tag{30}$$

Plugging this into the first order condition on routine labor, (11), and dividing with the FOC on manual labor, (10), and re-arranging we get:

$$\frac{l_{rJ}}{l_{mJ}} = \left[1+\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right]^{\frac{\rho-\sigma_c}{\sigma_c-1}} \left(\frac{w_{mJ}}{w_{rJ}}\right)^{\rho} \left(\frac{Z_{rJ}}{Z_{mJ}}\right)^{\rho-1}.$$

Multiplying the above with w_{rJ}/w_{mJ} and substituting in θ_{rJ}/θ_{mJ} we obtain (17):

$$\frac{Z_{mJ}}{Z_{rJ}} = \frac{w_{mJ}}{w_{rJ}}\left[1+\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right]^{\frac{\rho-\sigma_c}{(\sigma_c-1)(\rho-1)}} \left(\frac{\theta_{mJ}}{\theta_{rJ}}\right)^{\frac{1}{\rho-1}}.$$

Next we express the labor aggregate as:

$$LA_J = \left[\sum_{o=m,a}(Z_{oJ}l_{oJ})^{\frac{\rho-1}{\rho}} + RA_J^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}} = Z_{mJ}l_{mJ}\left[\frac{1}{\theta_{mJ}}\left(1+\frac{\Theta_{cJ}}{1-\Theta_J}\right)\right]^{\frac{\rho}{\rho-1}}, \tag{31}$$

using (30) and substituting in $\left(\frac{Z_{rJ}l_{rJ}}{Z_{mJ}l_{mJ}}\right)^{\frac{\rho-1}{\rho}} = \frac{\theta_{rJ}}{\theta_{mJ}}\left[1+\frac{\Theta_{cJ}}{(1-\theta_J)\theta_{rJ}}\right]^{\frac{\sigma_c-\rho}{(\sigma_c-1)\rho}}$ and $\left(\frac{Z_{aJ}l_{aJ}}{Z_{mJ}l_{mJ}}\right)^{\frac{\rho-1}{\rho}} = \frac{\theta_{aJ}}{\theta_{mJ}}$ (obtained similarly to (29)), and that $\sum_o \theta_{oJ} = 1$. Plugging the expression for LA_J

into the FOC for manual labor, (10), and dividing by the FOC on non-ICT capital, (13), and re-arranging we get:

$$\frac{k_J}{l_{mJ}} = \left(\frac{w_{mJ}}{R_k}\right)^{\sigma} \left(\frac{Z_{kJ}}{Z_{mJ}}\right)^{\sigma-1} \left[\frac{1}{\theta_{mJ}} \left(1 + \frac{\Theta_{cJ}}{1-\Theta_J}\right)\right]^{\frac{\rho-\sigma}{\rho-1}}.$$

Multiplying with R_k/w_{mJ} and re-arranging we get equation (18):

$$\frac{Z_{kJ}}{Z_{mJ}} = \frac{R_k}{w_{mJ}} \left(\frac{1}{\theta_{mJ}}\right)^{\frac{1}{\rho-1}} \left(\frac{\Theta_J - \Theta_{cJ}}{1-\Theta_J}\right)^{\frac{1}{\sigma-1}} \left(1 + \frac{\Theta_{cJ}}{1-\Theta_J}\right)^{\frac{\sigma-\rho}{(\rho-1)(\sigma-1)}}.$$

Finally we express sectoral output as a function of observables. Using the expression on LA_J (31) and substituting that $\left(\frac{Z_{kJ}k_J}{Z_{mJ}l_{mJ}}\right)^{\frac{\sigma-1}{\sigma}} = \frac{\Theta_J-\Theta_{cJ}}{(1-\Theta_J)\theta_{mJ}} \left[\frac{1}{\theta_{mJ}} \left(1 + \frac{\Theta_{cJ}}{1-\Theta_J}\right)\right]^{\frac{\sigma-\rho}{(\rho-1)\sigma}}$ (obtained similarly to (29)) we can express sectoral output as:

$$Y_J = \left(LA_J^{\frac{\sigma-1}{\sigma}} + (Z_{kJ}k_J)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} = Z_{kJ}k_J \left(\frac{1}{\Theta_J - \Theta_{cJ}}\right)^{\frac{\sigma}{\sigma-1}}$$

which is the expression in the main text.

D Decomposing labor-augmenting technological change

Figure A1 shows the path of sector-occupation technology changes (between each consecutive period) as extracted from the data, as well as the different predicted productivities based on the components derived from the factor model. The 'full factor' prediction (green solid line) is quite close to the data (red solid line with marker), illustrating that the contribution of technological change idiosyncratic to the sector-occupation cell is very small. For some cells, the 'occupation-only' predictions (the yellow dashed-dotted line) gives a good account of the data, whereas for others the 'sector-only' predictions (the blue dashed line) are closer. The neutral predictions (gray dotted line) give only minor changes for some cells (e.g. in the goods sector), whereas for others it is relatively close to the data (rH cell for example).

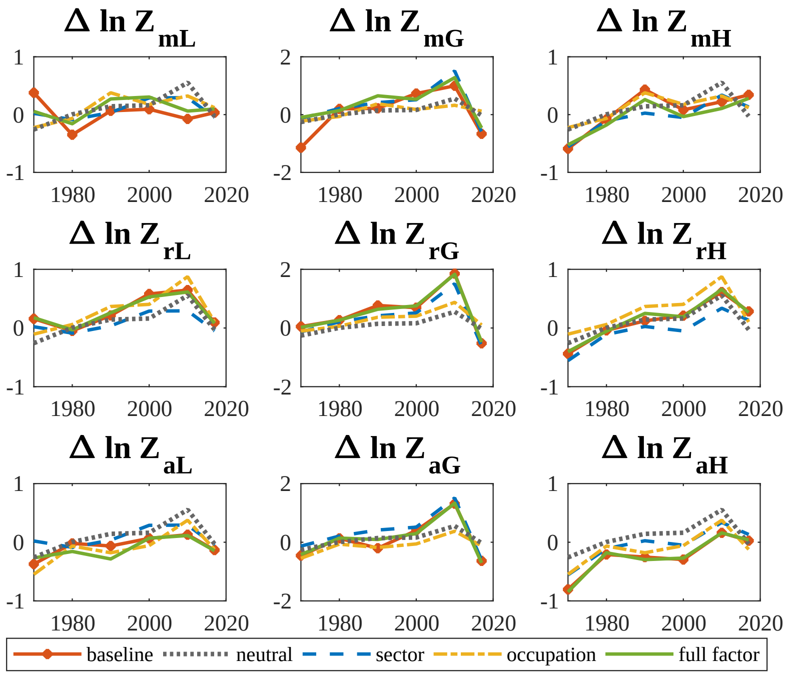

Figure A1: Baseline and counterfactual cell productivities

The solid red line with the marker shows the decennial change in the log of sector-occupation technologies, as calculated from the data. The other lines show the counterfactual paths, based on the neutral (gray dotted), the sector-specific (blue dashed line), the occupation-specific (yellow dashed-dotted), or sector- and occupation-specific (green solid line) components.

E Robustness checks and extensions

E.1 Alternative and heterogeneous substitution elasticities

We provide more detailed results for the robustness checks discussed in the main text, by contrasting the results from our baseline analysis with those of the alternative elasticity values. Table A3 shows in the top rows the average annual growth rates of the factor augmenting technologies in each sector in our baseline ($\sigma = 0.84, \rho = 0.6, \sigma_c = 2$). The segments below shows the range of each of these growth rates across the 27 alternative calibrations we run in our joint sensitivity analysis ($\sigma \in \{0.65, 0.75, 0.84\} \otimes \rho \in \{0.5, 0.6, 0.7\} \otimes \sigma_c \in \{1.5, 2, 2.5\}$). These ranges display similar key features to those which we highlighted in the discussion of Table 2.

Similarly Figures A2, A3, A4 and A5 show the results for sectoral labor productivity growth in the baseline along with the range of predictions across the parameters considered in the sensitivity analysis (shown as black error bars). These figures demonstrate that all the results from the baseline are very robust across all the alter-

Table A3: Average annual growth rate of Zs over 1960–2017 under alternative parameters

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
			Baseline:		
Low-sk. serv.	0.26	2.93	-0.69	0.85	2.02
Goods	0.59	5.61	0.98	-1.61	4.43
High-sk. serv.	0.72	1.32	-2.38	1.78	-1.94
			Sensitivity Ranges:		
Low-sk. serv.	[-0.41, 0.59]	[2.39, 3.47]	[-1.59, -0.22]	[0.85, 1.03]	[0.17, 6.96]
Goods	[-1.67, 1.24]	[3.94, 6.19]	[-1.31, 1.65]	[-1.61, -0.01]	[1.40, 9.63]
High-sk. serv.	[0.43, 2.04]	[0.43, 3.18]	[-2.96, -1.24]	[0.83, 1.78]	[-3.09, 1.69]

Notes: All numbers are annualized growth rates in percent.

native parametrizations.

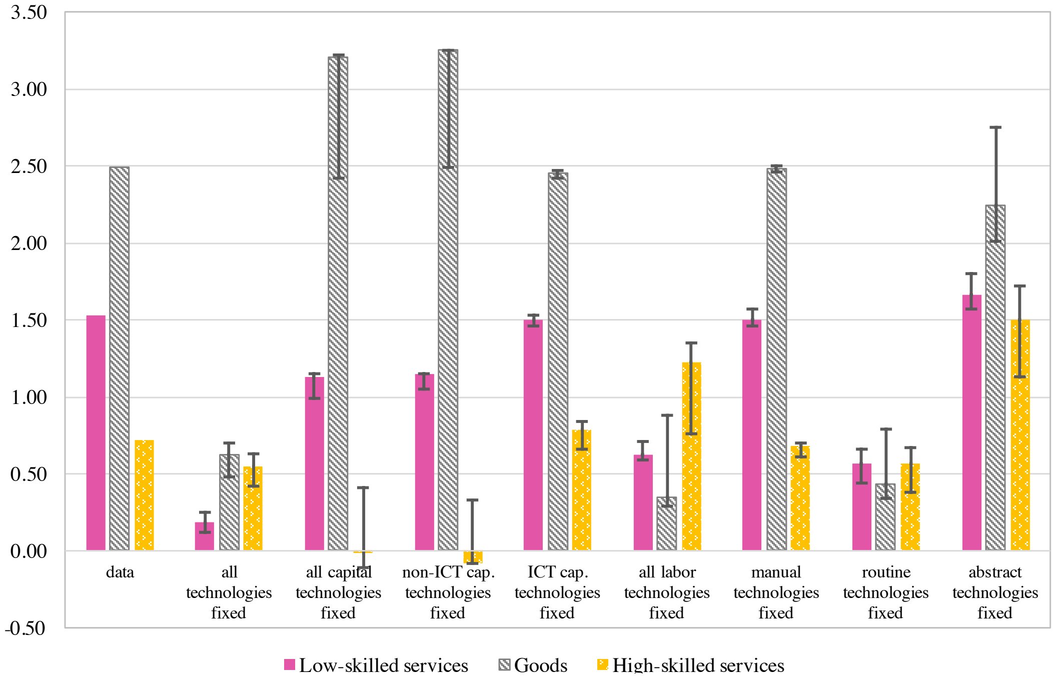

Figure A2: Average sectoral labor productivity growth with fixed technologies

This figure shows the role of the different factor-augmenting technologies in sectoral labor productivity growth for different elasticities, along with error bars that represent the range of results across the sensitivity analysis.

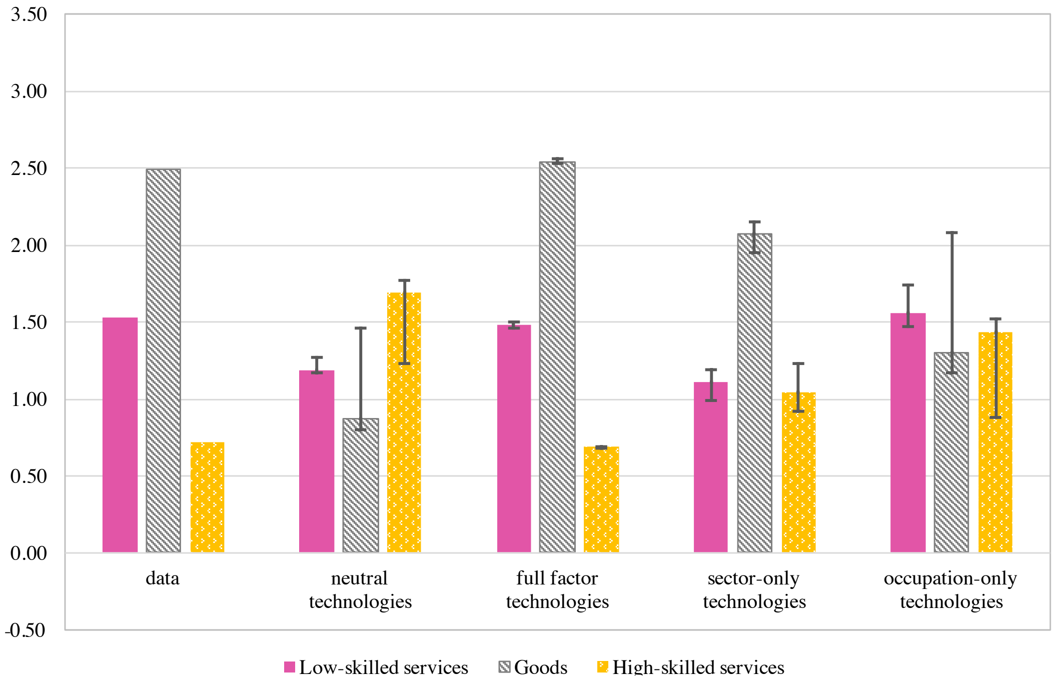

Figure A3: Average sectoral labor productivity growth with alternative technologies

This figure shows the role of the various components of labor-augmenting technologies in sectoral labor productivity growth for our baseline calibration along with error bars that represent the range of results across the sensitivity analysis.

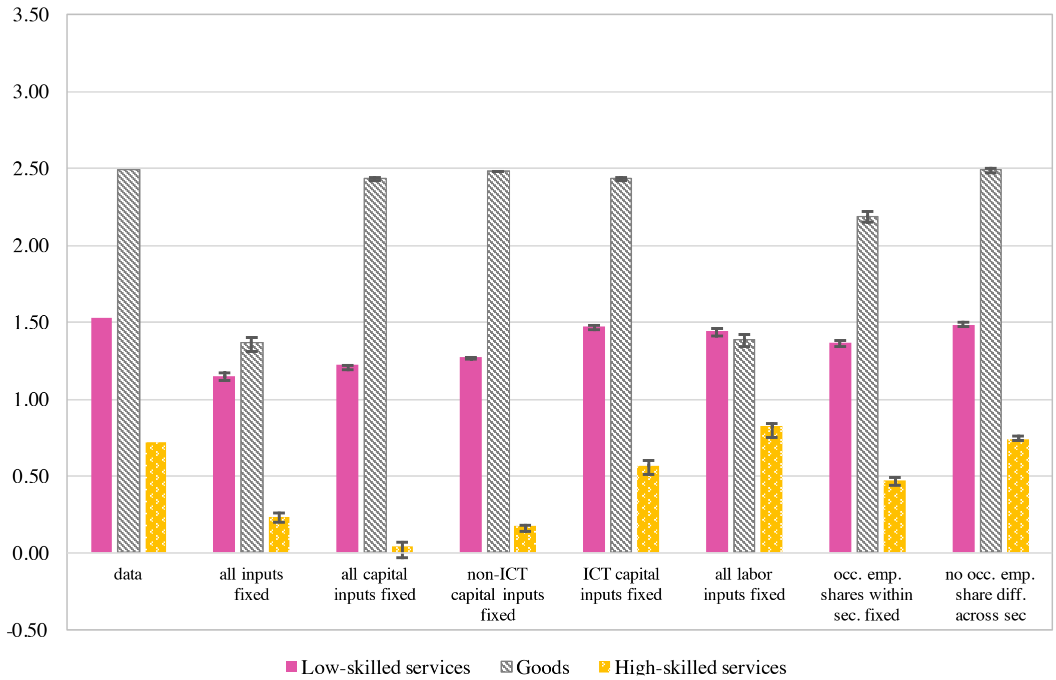

Figure A4: Average sectoral labor productivity growth with fixed inputs

This figure shows the role of changing factor input use in sectoral labor productivity growth for different elasticities, along with error bars that represent the range of results across the sensitivity analysis.

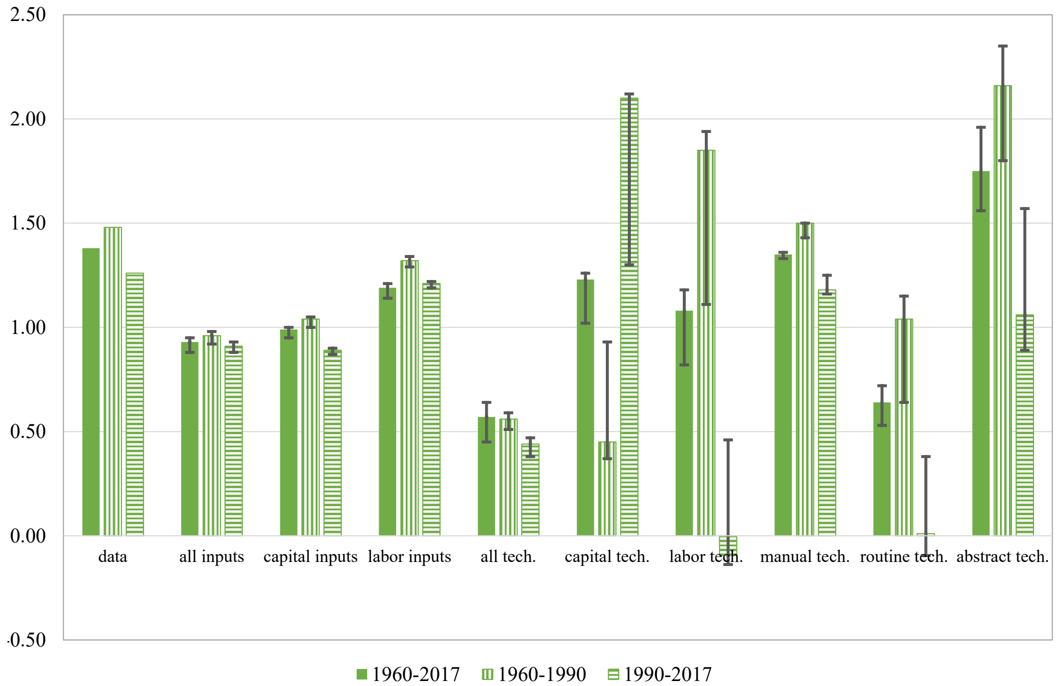

Figure A5: Counterfactual aggregate labor productivity growth in different periods

This figure shows the role of various inputs and technologies in aggregate labor productivity growth between 1960-2017, 1960-1990 and 1990-2017 for our baseline calibration along with error bars that represent the range of results across the sensitivity analysis.

Additionally, Table A4 shows for the range of ρ values which have been considered in the literature the distance measure between the changes in sector-occupation cell technologies inferred from the data and the predictions based on the various components of the factor model. This table shows that the distance measures of the predictions based on the neutral, on the sector and on the occupation components vary quite a bit with the value of ρ. If the elasticity of substitution between different occupations is low then the sector components play a larger role, while if ρ is high, then the occupation components are more important. However, the full factor prediction reproduces the data quite well for all values of ρ considered.

Table A4: Distance measure of the different predictions

ρ	**neutral**	**full factor**	**sector**	**occupation**
0.5	0.674	0.024	0.151	0.455
0.6	**0.703**	**0.033**	**0.228**	**0.407**
0.7	0.751	0.049	0.361	0.339
0.8	0.833	0.076	0.590	0.247
0.9	0.943	0.111	0.914	0.168

Sector-specific substitution elasticity of non-ICT capital and labor aggregate

We also consider sector-specific values of the elasticity of substitution between non-ICT capital and labor. The average annual growth rates of the impliedsector-specific factor augmenting technologies are shown in Table A5. Both the sector- and the factor-specific patterns that we found in our baseline analysis are maintained.

Table A5: Average annual growth rate of Zs 1960–2017, sector-specififc σ^J

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
1960-2017					
Low-skilled services	0.18%	2.84%	-0.77%	0.97%	1.94%
Goods	1.92%	7.00%	2.32%	-3.34%	5.81%
High-skilled services	1.26%	1.87%	-1.86%	1.15%	-1.42%

Figure A6 shows the robustness of the model to allowing for different σ^J across sectors in terms of its implications for sectoral labor productivity. In this figure the column on the left shows the baseline results, and the one on the right the results with heterogeneity across sectors.

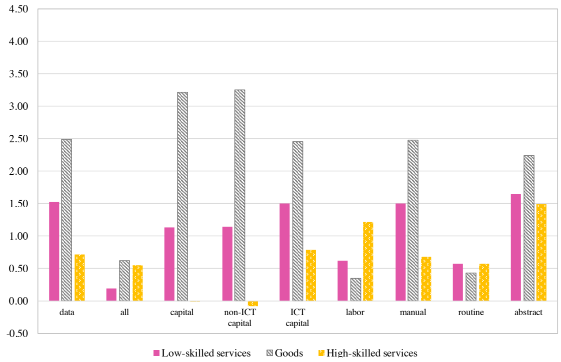

(a) Role of technologies, baseline

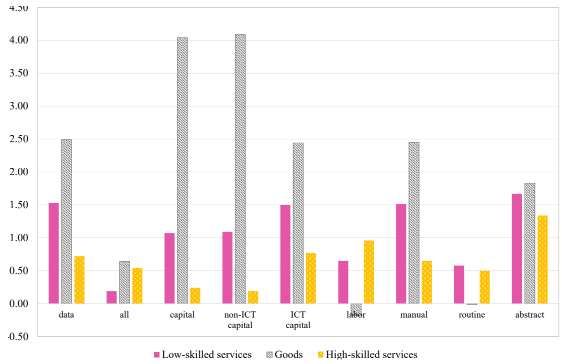

(b) Role of technologies, σ^J

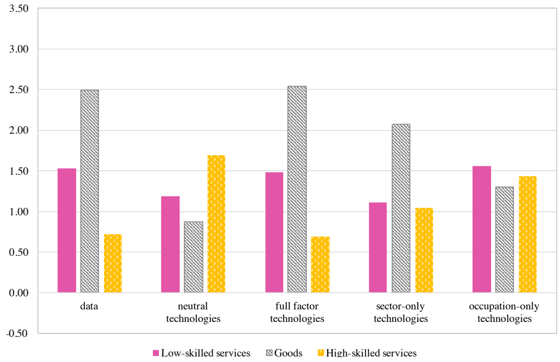

(c) Role of components, baseline

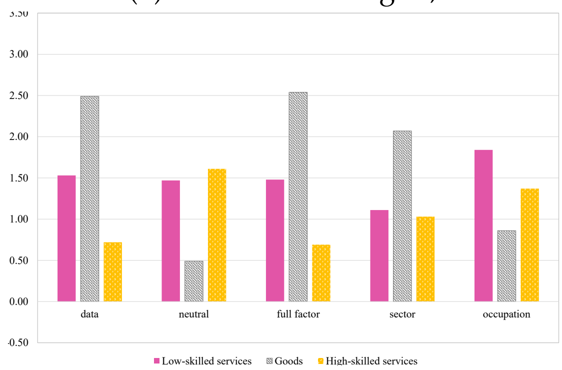

(d) Role of components, σ^J

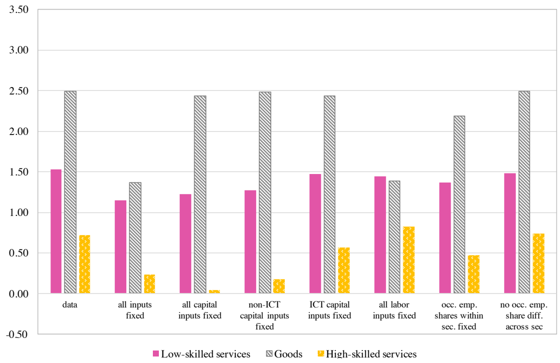

(e) Role of inputs, baseline

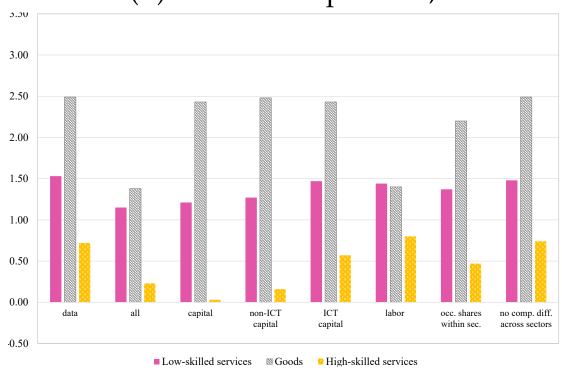

(f) Role of inputs, σ^J

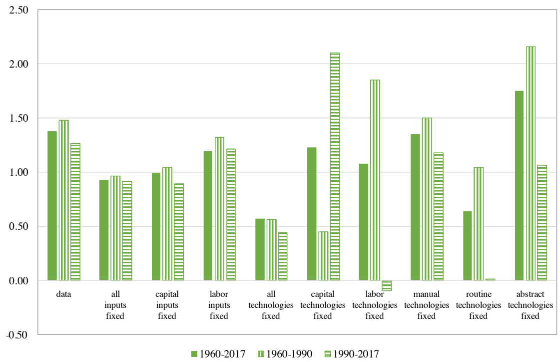

(g) Aggregate labor prod., baseline

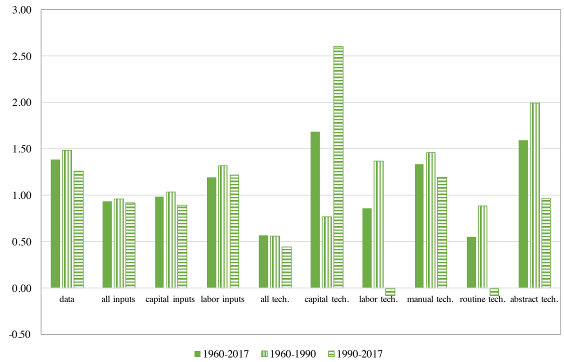

(h) Aggregate labor prod., σ^J

Figure A6: Baseline vs sector specific σ^J

This figure shows the differences between the magnitude of the various channels when considering sector-specific σs relative to the baseline. The values used are $\sigma^L = \sigma^H = 0.75$ and $\sigma^G = 0.9$. In the baseline these are all set to 0.84.

E.2 Different Nesting

Given the production structure described in section 5.2, the each sector's representative firm's problem is given by

$$\max_{\{l_{oJ}\},c_J,k_J} p_J Y_J - \sum_o w_{oJ} l_{oJ} - R_c c_J - R_k k_J.$$

The first order conditions that optimal input use in each sector has to satisfy are:

$$\frac{\partial \pi_J}{\partial l_{mJ}} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho_m}{\rho_m \sigma}} Z_{mJ}^{\frac{\rho_m-1}{\rho_m}} l_{mJ}^{-\frac{1}{\rho_m}} - w_{mJ} = 0, \tag{32}$$

$$\frac{\partial \pi_J}{\partial l_{aJ}} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho_m}{\rho_m \sigma}} AA_J^{\frac{\rho_m-\rho_a}{\rho_a \rho_m}} Z_{aJ}^{\frac{\rho_a-1}{\rho_a}} l_{aJ}^{-\frac{1}{\rho_a}} - w_{aJ} = 0, \tag{33}$$

$$\frac{\partial \pi_J}{\partial l_{rJ}} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho_m}{\rho_m \sigma}} AA_J^{\frac{\rho_m-\rho_a}{\rho_a \rho_m}} RA_J^{\frac{\rho_a-\sigma_c}{\sigma_c \rho_a}} Z_{rJ}^{\frac{\sigma_c-1}{\sigma_c}} l_{rJ}^{-\frac{1}{\sigma_c}} - w_{rJ} = 0, \tag{34}$$

$$\frac{\partial \pi_J}{\partial c_J} = p_J Y_J^{\frac{1}{\sigma}} LA_J^{\frac{\sigma-\rho_m}{\rho_m \sigma}} AA_J^{\frac{\rho_m-\rho_a}{\rho_a \rho_m}} RA_J^{\frac{\rho_a-\sigma_c}{\sigma_c \rho_a}} Z_{cJ}^{\frac{\sigma_c-1}{\sigma_c}} c_J^{-\frac{1}{\sigma_c}} - R_c = 0, \tag{35}$$

$$\frac{\partial \pi_J}{\partial k_J} = p_J Y_J^{\frac{1}{\sigma}} Z_{kJ}^{\frac{\sigma-1}{\sigma}} k_J^{-\frac{1}{\sigma}} - R_k = 0. \tag{36}$$

The first order conditions on routine labor and ICT capital give exactly the same relative technology levels as in our baseline, (16), and the expression for the routine aggregate under optimal input use is unchanged. The first order conditions on routine and abstract labor using the expression on RA_J, (30), give:

$$\frac{l_{rJ}}{l_{aJ}} = \left(1 + \frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right)^{\frac{\rho_a-\sigma_c}{\sigma_c-1}} \left(\frac{w_{aJ}}{w_{rJ}}\right)^{\rho_a} \left(\frac{Z_{rJ}}{Z_{aJ}}\right)^{\rho_a-1}.$$

Multiplying both sides by w_{rJ}/w_{aJ} and re-arranging we get:

$$\frac{Z_{aJ}}{Z_{rJ}} = \frac{w_{aJ}}{w_{rJ}} \left(1 + \frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right)^{\frac{\rho_a-\sigma_c}{(\sigma_c-1)(\rho_a-1)}} \left(\frac{\theta_{aJ}}{\theta_{rJ}}\right)^{\frac{1}{\rho_a-1}}. \tag{37}$$

Using optimal input use we can express the abstract aggregate as:

$$AA_{J,t} = Z_{aJ} l_{aJ} \left[1 + \left(1 + \frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right) \frac{\theta_{rJ}}{\theta_{aJ}}\right]^{\frac{\rho_a}{\rho_a-1}}.$$

The first order conditions on manual and abstract labor (using the above expression for AA_J) give:

$$\frac{l_{mJ}}{l_{aJ}} = \left(\frac{Z_{mJ}}{Z_{aJ}}\right)^{\rho_m-1}\left(\frac{w_{aJ}}{w_{mJ}}\right)^{\rho_m}\left[1+\left(1+\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right)\frac{\theta_{rJ}}{\theta_{aJ}}\right]^{\frac{\rho_a-\rho_m}{\rho_a-1}}.$$

Multiplying both sides by w_{mJ}/w_{aJ} and re-arranging we get:

$$\frac{Z_{aJ}}{Z_{mJ}} = \frac{w_{aJ}}{w_{mJ}}\left[1+\left(1+\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}}\right)\frac{\theta_{rJ}}{\theta_{aJ}}\right]^{\frac{\rho_a-\rho_m}{(\rho_a-1)(\rho_m-1)}}\left(\frac{\theta_{aJ}}{\theta_{mJ}}\right)^{\frac{1}{\rho_m-1}}. \quad (38)$$

Using optimal input use we can express the labor aggregate as:

$$LA_{J,t} = Z_{mJ}l_{mJ}\left[\frac{1}{\theta_{mJ}}\left(1+\frac{\Theta_{cJ}}{1-\Theta_J}\right)\right]^{\frac{\rho_m}{\rho_m-1}}.$$

Again we use the FOCs of non-ICT capital and manual labor to get:

$$\frac{k_J}{l_{mJ}} = \left(\frac{Z_{kJ}}{Z_{mJ}}\right)^{\sigma-1}\left(\frac{w_{mJ}}{R_J}\right)^{\sigma}\left[\frac{1}{\theta_{mJ}}\left(1+\frac{\Theta_{cJ}}{1-\Theta_J}\right)\right]^{\frac{\rho_m-\sigma}{\rho_m-1}}$$

Multiplying both sides by R/w_{mJ} and re-arranging gives us:

$$\frac{Z_{mJ}}{Z_{kJ}} = \frac{w_{mJ}}{R}\theta_{mJ}^{\frac{1}{\rho_m-1}}\left(\frac{(1-\Theta_J)}{\Theta_J-\Theta_{cJ}}\right)^{\frac{1}{\sigma-1}}\left(1+\frac{\Theta_{cJ}}{1-\Theta_J}\right)^{\frac{\rho_m-\sigma}{(\rho_m-1)(\sigma-1)}}. \quad (39)$$

Finally, we express sectoral output as

$$Y_J = Z_{kJ}k_J\left(\frac{1}{\Theta_J-\Theta_{cJ}}\right)^{\frac{\sigma}{\sigma-1}},$$

which pins down the growth rate of Z_{kJ} in each sector as in our baseline, (19). Initial technology levels are also pinned down exactly as before, (20). All sector-factor technologies can then be extracted from the data using expressions (16), (37), (38), (39), (19), and (20), for given values for each elasticity of substitution. We take the values of the new elasticities from vom Lehn (2019), these are $\rho_a = 0.31, \rho_m = 1.49$, while we keep all other elasticities at their baseline values.

Table A6 shows the technology growth rates when nesting the labor inputs in a different way. The growth rates for non-ICT capital are exactly the same as before, all other growth rates are different to some extent. Most of the sector and factor patterns are maintained. However, technological progress in manual occupations is much higher than before. This should not be surprising, as manual labor input is a good substitutes now for all other labor inputs, hence to obtain the same relative employment and wage growth we need higher productivity growth than in our baseline.

Table A6: Average annual growth rate of Zs 1960–2017, with different nesting

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
1960-2017					
Low-skilled services	3.31%	1.91%	-0.07%	0.85%	1.00%
Goods	6.68%	4.83%	2.32%	-1.61%	3.64%
High-skilled services	-1.82%	0.39%	-1.47%	1.78%	-2.86%

The counterfactual sectoral and aggregate labor productivity growth numbers are, however, very similar to our baseline. While there are some small quantitative differences, for example shutting down routine labor productivity growth implies a larger reduction in labor productivity growth in the goods sector, but a smaller reduction in the two service sectors, the qualitative predictions are exactly the same.

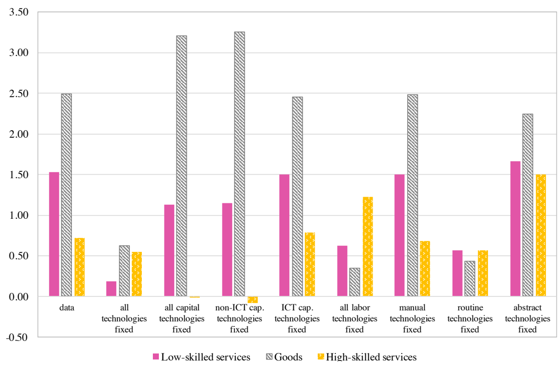

(a) Role of technologies, baseline

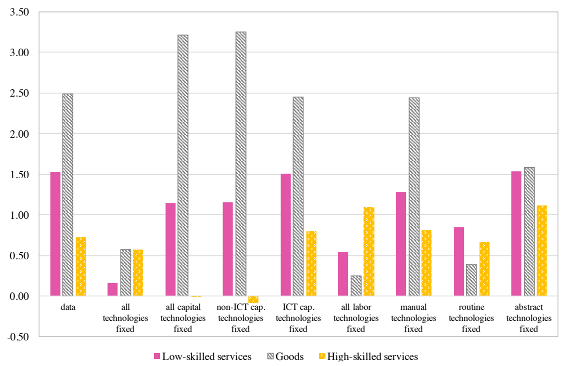

(b) Role of technologies, different nesting

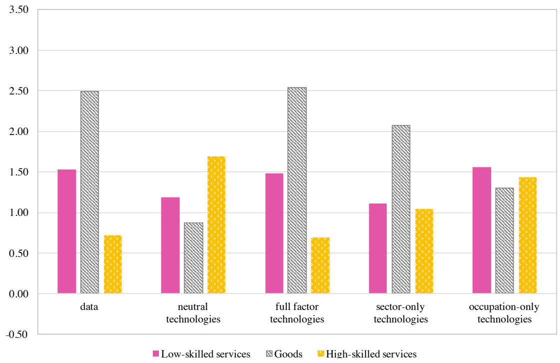

(c) Role of components, baseline

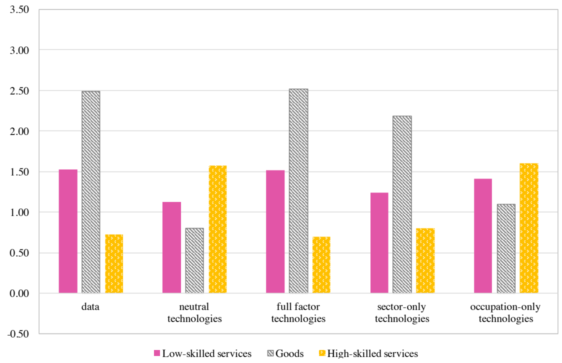

(d) Role of components, different nesting

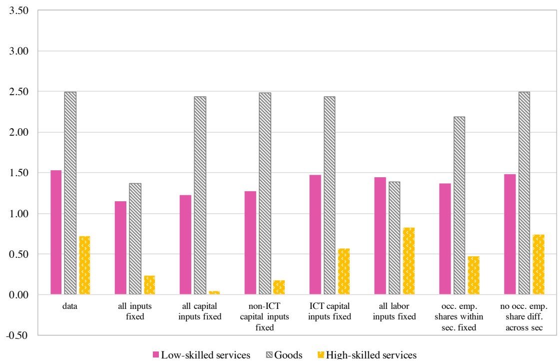

(e) Role of inputs, baseline

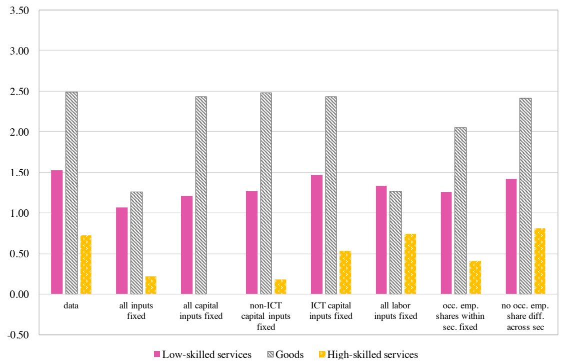

(f) Role of inputs, different nesting

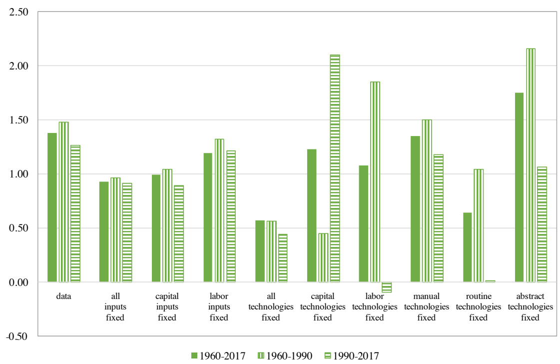

(g) Aggregate labor prod., baseline

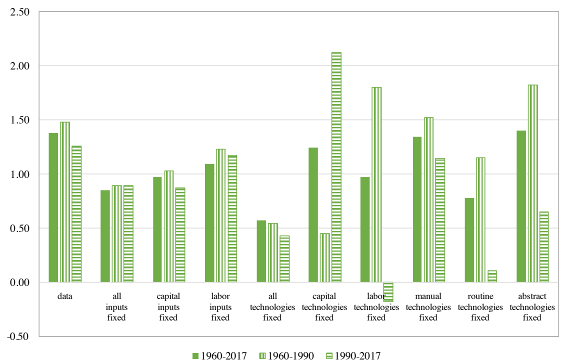

(h) Aggregate labor prod., different nesting

Figure A7: Baseline vs different nesting of labor inputs

This figure shows the differences between the magnitude of the various channels with a different nesting of labor inputs relative to the baseline.

E.3 Allowing for efficiency units of labor

To control for workers' skills, we employ the following Mincer wage regression

$$\log w_{ioJt} = \delta_{oJt} + \beta' X_{it} + \varepsilon_{ioJt}, \tag{40}$$

where δ_{oJt} are occupation-sector-time effects and X_{it} is a vector of worker characteristics. From this regression we can back out both an occupation-sector wage in year t that is not confounded by changes in composition of worker characteristics, X_it, as well as an estimate of the average efficiency units a worker in occupation o and sector J has in year t. 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, interacted with a dummy for college education and with a gender dummy, as well as a dummy for foreign-born and non-white race. Note that for our model to match the average hourly wages by sector-occupation cell in every period ($\overline{w}_{oJ,t}$), we need to assign the cell-year average of the exponent of the residuals from (40) to either the average wage per efficiency units or to the average efficiency units per hour worked. Thus we have two options. Either we construct the sector-occupation cell efficiency units per hour, $\overline{e}^1_{oJ,t}$, as the average of $\widehat{e}^1_{ioJt} = \exp(\beta' X_{it})$ within the sector-occupation-year cell. In this case the implied sector-occupation-year unit wages are given as $\widehat{w}^1_{oJ,t} = \overline{w}_{oJ,t}/\overline{e}^1_{oJ,t}$. Alternatively we construct sector-occupation cell efficiency wages per hour, $\widehat{w}^2_{oJ,t} = \exp(\delta_{oJt})$. The implied average sector-occupation-year efficiency units per hour worked are then $\overline{e}^2_{oJ,t} \equiv \overline{w}_{oJ,t}/\widehat{w}^2_{oJ,t}$.

We use the equivalent of (28) to get sector-occupation wages per efficiency unit ($\widetilde{w}_{oJ,t}$):

$$\widetilde{w}^M_{oJ,t} l_{oJ,t} \overline{e}^M_{oJ,t} = Y^{nom}_t \cdot VA_{J,t}(1 - \Theta_{J,t})\theta_{oJ,t}, \tag{41}$$

where $\overline{e}^M_{oJ,t}$ is the average sector-occupation efficiency units per hour worked in period t (according to method $M = 1, 2$). The within sector relative wages implied by the accounting identity are:

$$\frac{\widetilde{w}^M_{oJ,t}}{\widetilde{w}^M_{rJ,t}} = \frac{\frac{\theta_{oJ,t}}{l_{oJ,t}\overline{e}^M_{oJ,t}}}{\frac{\theta_{rJ,t}}{l_{rJ,t}\overline{e}^M_{rJ,t}}} = \frac{\frac{\overline{w}_{oJ,t}}{\overline{e}^M_{oJ,t}}}{\frac{\overline{w}_{rJ,t}}{\overline{e}^M_{rJ,t}}} = \frac{\widehat{w}^M_{oJ,t}}{\widehat{w}^M_{rJ,t}},$$

where the last equality follows as both our methods ensure that we match each cell's average hourly wage. Thus in this formulation – just as in the baseline – the within-sector relative wages per efficiency units obtained from the accounting identity are the same as those implied by the Mincer wage regression.

Table A7: Sector-occupation efficiency units of labor 1960–2017

Sector	Occupation	1960	1970	1980	1990	2000	2010	2017
Low-skilled services	manual	1.798	1.751	1.594	1.660	1.717	1.731	1.719
	routine	2.067	1.993	1.863	1.929	1.989	2.015	1.989
	abstract	2.310	2.248	2.117	2.168	2.246	2.272	2.245
Goods	manual	2.203	2.113	2.004	1.986	2.020	2.043	2.042
	routine	2.107	2.055	1.945	2.024	2.083	2.151	2.120
	abstract	2.457	2.499	2.455	2.507	2.593	2.668	2.621
High-skilled services	manual	2.032	1.978	1.892	1.994	2.043	2.052	2.043
	routine	1.943	1.877	1.819	1.900	1.953	2.011	2.048
	abstract	2.404	2.398	2.340	2.425	2.481	2.511	2.505

(a) fitted efficiency units, $\overline{e}^1$

Sector	Occupation	1960	1970	1980	1990	2000	2010	2017
Low-skilled services	manual	2.242	2.166	1.880	1.904	1.979	2.104	2.089
	routine	2.332	2.254	2.141	2.225	2.297	2.502	2.488
	abstract	2.799	2.685	2.539	2.577	2.663	2.795	2.819
Goods	manual	2.436	2.383	2.310	2.276	2.375	2.493	2.457
	routine	2.382	2.313	2.226	2.301	2.376	2.563	2.500
	abstract	2.799	2.817	2.762	2.847	2.981	3.148	3.152
High-skilled services	manual	2.290	2.249	2.168	2.286	2.360	2.512	2.512
	routine	2.091	2.078	2.065	2.177	2.272	2.479	2.585
	abstract	2.761	2.746	2.646	2.766	2.873	3.026	3.090

(b) residual efficiency units, $\overline{e}^2$

Table A7 shows efficiency units by sector-occupation over time for the two methods. While there is a level difference between the efficiency units directly fitted and the ones backed out as a residual from wages, the two methods give very similar patterns for the evolution of each sector-occupation cell's average efficiency over time.

In the variant of the model with efficiency units of labor, firms choose $n_{oJ,t} \equiv \overline{e}_{oJ,t} l_{oJ,t}$ in each period, instead of just hours worked ($l_{oJ,t}$). This implies that we need to use wages per efficiency unit of labor in equations (15) to (20), but the measurement of all other variables remains the same as in the baseline model. Figure A8 plots the alternative series for the relative wages within sectors. The dotted lines show method 1 and the dashed lines show method 2 applied in (41), and the solid lines show our baseline (of wages per hour worked from (28)). Note that all alternative lines qualitatively show the same patterns, most are also quantitatively very close. Given that

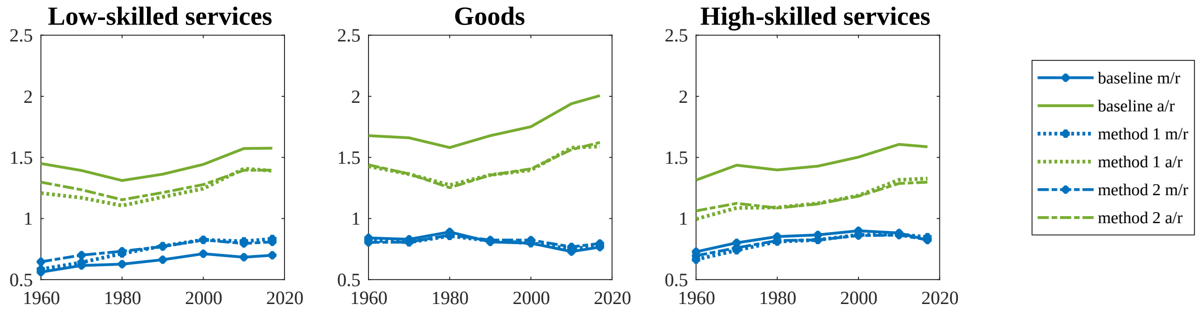

Figure A8: Comparison of relative wages

Notes: This figure plots the relative hourly wages of manual (blue with marker) and abstract (green) compared to routine workers within each sector over time for three alternative ways to compute wages: (i) from the baseline model without efficiency units (equation (28), solid lines), (ii) from fitted efficiency units ((41) based on method 1, dotted lines), (iii) from fitted efficiency wages ((41) based on method 2, dashed lines).

the relative wage path are similar to those in our baseline, it is not surprising that our results are robust to controlling for skills.

Given the series of wages per efficiency unit of labor, $\widetilde{w}_{oJ,t}^{M}$ we constructed for the two methods $M = 1, 2$, and all the other data we use in the main part of the paper, we use again our methodology to infer the factor-augmenting technologies in each sector. Table A8 shows the average annual change in the labor-augmenting technologies over 1960–2017. We do not report the results for the technology of ICT and of non-ICT capital here, as these are exactly the same as in the baseline model because they are independent of how labor income is split. Equations (15) to (19) imply that differences in the measurement of wage growth over time result in differential growth rates in the labor-augmenting technologies, but do not affect the growth rates of Z_{cJ} or Z_{kJ}.

Table A8: Average annual growth rate of Zs over 1960–2017 accounting for efficiency units of labor

	based on fitted effiency units, $\bar{e}^1$			based on residual effiency units, $\bar{e}^2$		
	manual	routine	abstract	manual	routine	abstract
Low-sk. serv.	0.34%	2.99%	-0.64%	0.38%	2.81%	-0.70%
Goods	0.72%	5.60%	0.87%	0.57%	5.52%	0.77%
High-sk. services	0.71%	1.23%	-2.45%	0.56%	0.95%	-2.57%

Notes: The change in the capital inputs' technologies (the Z_{cJ}s and Z_{kJ}s) is exactly the same as in Table 2 and not shown here.

Comparing Table A8 to Table 2 reveals that in both variants of the model with efficiency units the resulting growth rates of labor-augmenting technologies are very similar to the baseline model of the main text, both in terms of the ranking of growth in Z_{oJ} but also quantitatively. This is perhaps not that surprising given that we established in Figure A8 already that the relative occupational wages within a sector do not change much when we control for workers' characteristics. Since we identify the within-sector ratios of occupational productivities precisely from this ratio, but the across-time changes from objects that do not depend on the measurement of wages or efficiency units, the general conclusions about inferred technological change do not change when we measure the labor inputs in terms of hours worked times efficiency units.

Since the series of the factor-augmenting technologies (by sector) in the model with efficiency units of labor are so similar to the baseline model, and in fact for the capital inputs coincide, the implications for sectoral labor productivity are very similar too. While there are very small quantitative differences when studying the role of individual inputs or technologies, qualitatively they have the very same implications. Figure A9 shows this for the model variant based on fitted efficiency units, $\overline{e}^1$.

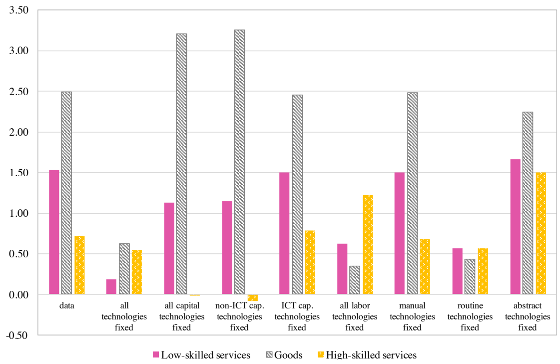

(a) Role of technologies, baseline

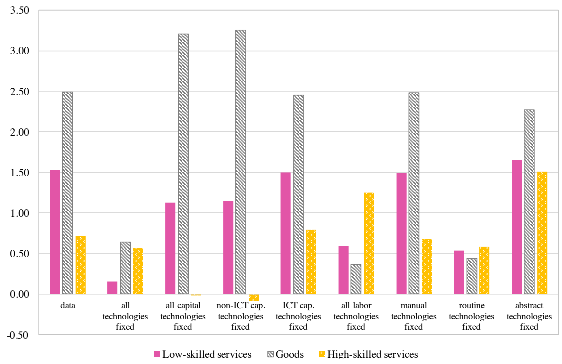

(b) Role of technologies, efficiency units

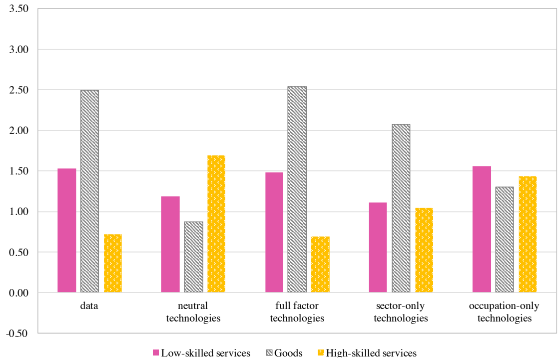

(c) Role of components, baseline

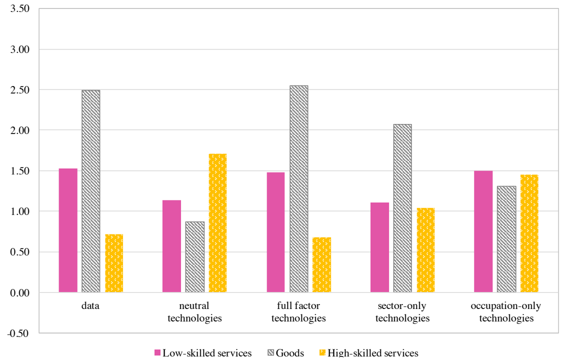

(d) Role of components, efficiency units

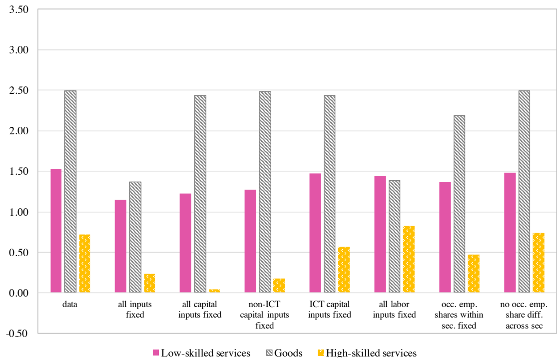

(e) Role of inputs, baseline

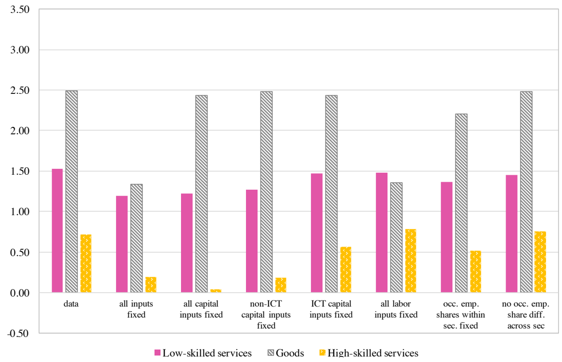

(f) Role of inputs, efficiency units

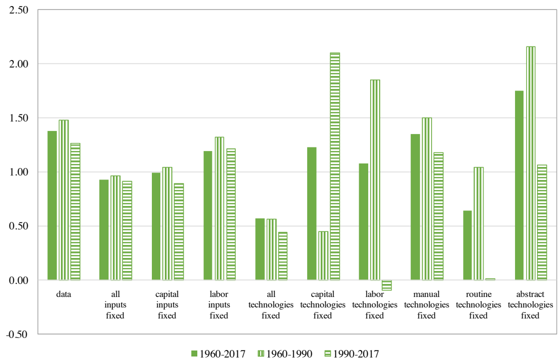

(g) Aggregate labor prod., baseline

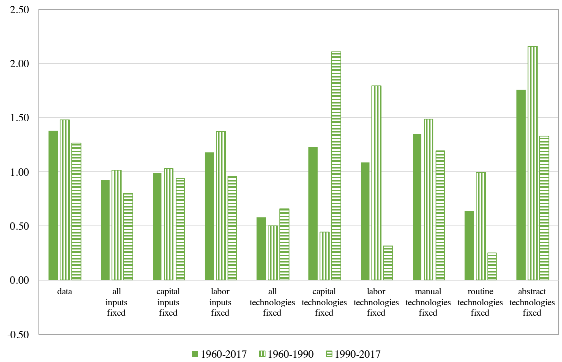

(h) Aggregate labor prod., efficiency units

Figure A9: Baseline vs efficiency unit model

This figure shows the differences between the magnitude of the various channels when considering the model with efficiency units relative to the baseline.

E.4 Including the self-employed

In our baseline specification we excluded the self-employed from (i) sectoral employment and (ii) sectoral labor income shares both calculated from BEA data, as well as (iii) from hours worked and occupational labor income shares within sectors (as well as the implied hourly wages) calculated from Census/ACS data. Our choice was driven by the fact that in the data it is nearly impossible to distinguish the labor and capital income of the self-employed, leading to problems in calculating (ii) and (iii). A large fraction of the self-employed do not report their earnings in the Census/ACS, and in the national accounts only the total proprietors' income is reported without specifying which part of it is labor and which is capital income. Below we detail how we deal with each of these. We calculate sectoral employment (and sectoral employment shares) by including the number of self-employed by sector in addition to full time equivalent employees.

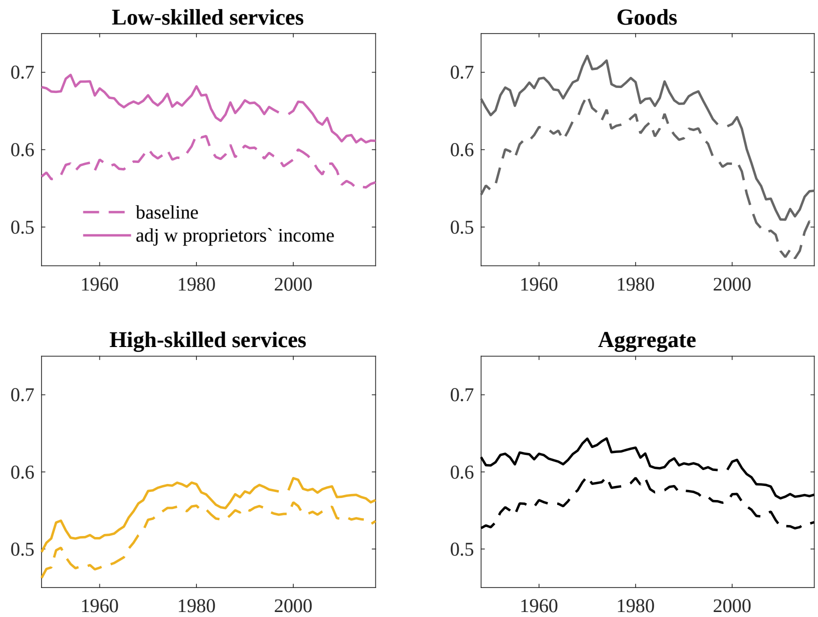

Figure A10: Labor income share adjusted with proprietors' income by sector

Notes: Labor income share as in our baseline (dashed lines, as in 2) and adjusted with proprietors' income (solid lines).

To calculate sectoral labor income shares including the self-employed, we split proprietors' income into capital and labor income in the same proportion as the rest of sectoral income is split between capital and labor income, as suggested in Gollin (2002) and Elsby et al. (2013). To obtain proprietors' income by our three sectors we combined data on Nonfarm Proprietors' Income by Industry (downloaded directly from the BEA, Tables 6.12.B, C, D, U.S. Bureau of Economic Analysis (2020b)) with Farm Proprietors' Income with inventory valuation and capital consumption adjustments (downloaded from FRED U.S. Bureau of Economic Analysis (2020a)). In Figure A10 we plot our baseline labor income shares (compensation of employees/value added, dashed lines) by sector and for the overall economy, $1-\Theta_J$, as well as the labor income shares adjusted with the proprietors' income (with solid lines), $1-\widetilde{\Theta}_J$.

The difference between the two sectoral labor income shares, $\Theta_J - \widetilde{\Theta}_J$ gives the fraction of income that accrues to the self-employed as labor income by sector. We use the information from the Census/ACS on hours worked to impute the share of this by occupation among the self-employed. To impute relative occupational labor income shares of the self-employed, we make the assumption that the relative wage of any two occupations within a sector is the same among employees and among the self-employed:

$$\frac{\theta^{se}_{oJ}}{\theta^{se}_{mJ}} = \frac{l^{se}_{oJ}}{l^{se}_{mJ}}\frac{w_{oJ}}{w_{mJ}},$$

where l^{se}_{oJ} are the hours worked by self-employed in occupation o in sector J, and w_{oJ}/w_{mJ} are relative occupational employee wages within a sector as calculated in (2). While similar, this assumption is slightly less restrictive than assuming that the wages of employees and the self-employed are the same. The occupational share of labor income within sectors, taking into account the self-employed, can then be calculated as:

$$\widetilde{\theta}_{oJ} = \frac{(1-\Theta_J)\theta_{oJ} + (\Theta_J - \widetilde{\Theta}_J)\theta^{se}_{oJ}}{1-\widetilde{\Theta}_J},$$

where θ_{oJ} is calculated as in (1) from the Census data for the employees only.

Table A9: Average annual growth rate of Zs 1960–2017, including the self-employed

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
1960-2017					
Low-skilled services	0.53%	3.25%	1.32%	-0.56%	2.45%
Goods	-0.30%	5.37%	2.71%	-2.58%	4.35%
High-skilled services	1.04%	1.59%	-1.69%	1.24%	-0.37%

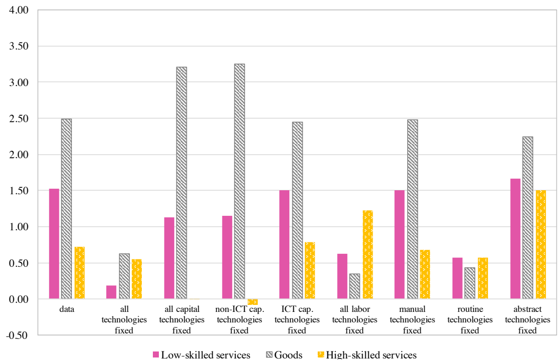

(a) Role of technologies, baseline

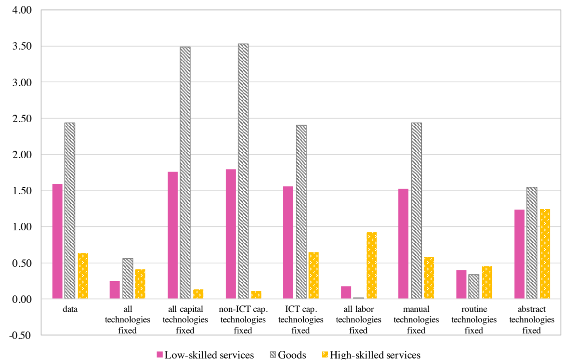

(b) Role of technologies, including the SE

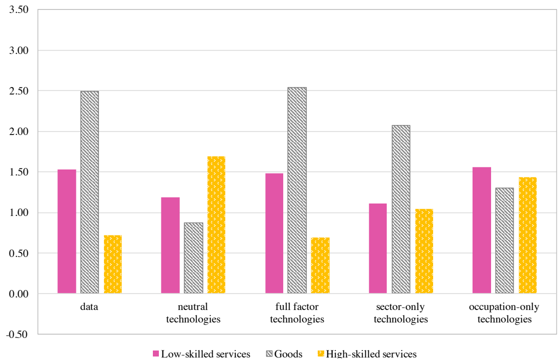

(c) Role of components, baseline

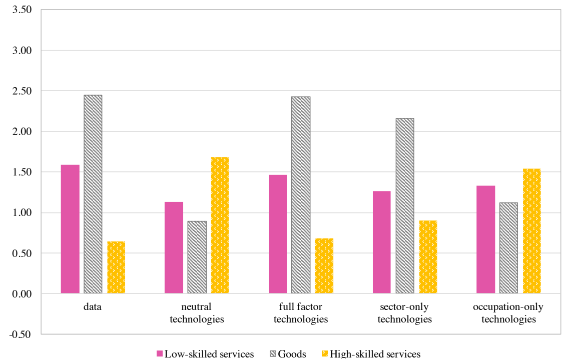

(d) Role of components, including the SE

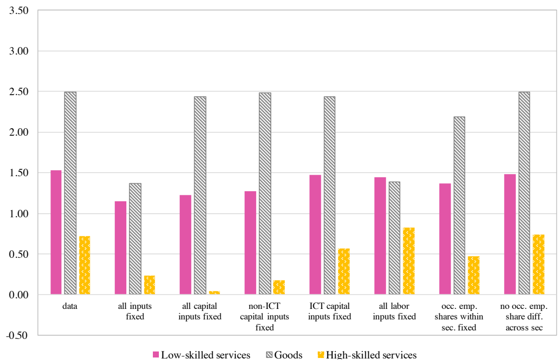

(e) Role of inputs, baseline

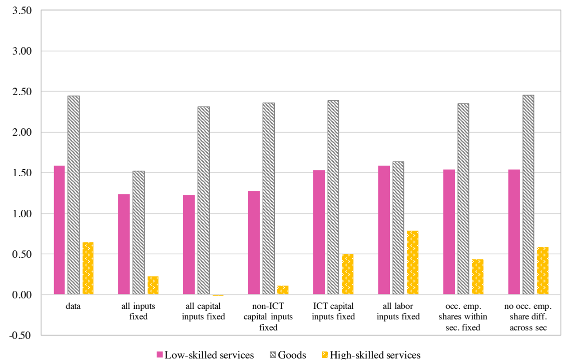

(f) Role of inputs, including the SE

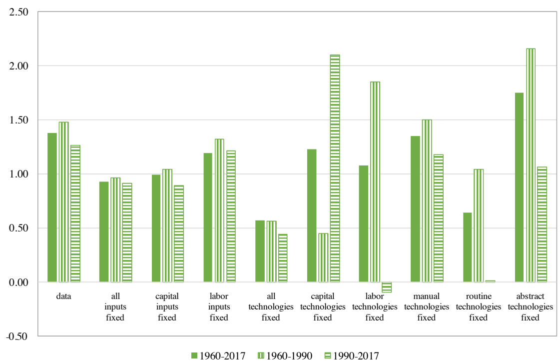

(g) Aggregate labor prod., baseline

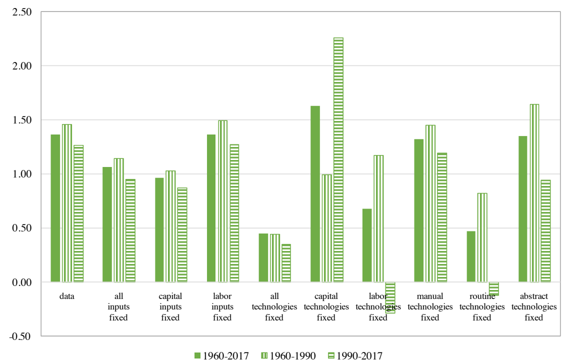

(h) Aggregate labor prod., including the SE

Figure A11: Baseline vs including the self-employed (SE)

This figure shows the differences between the magnitude of the various channels when including the self-employed relative to the baseline.

E.5 Separating Agriculture from Industry

In our main analysis we grouped various industries together into the goods sector. In this subsection we break this up and differentiate between agriculture and industry, comprising of manufacturing, mining and construction. and agriculture. We use the same methodology and parameterization of elasticities as in our baseline model. The results for low- and high-skilled services are identical to the baseline by construction, so focus here only on those for industry and agriculture. The results for the industry sector are very similar to what we found for the overall goods sector, while agricultural looks somewhat different. While the growth rates for some of the factor-augmenting technologies are different for industry and the broad goods sector (comparing Table A10 to Table 2), these differences occur in factors that carry low initial weights (see Table 4) and have therefore little consequence for the various drivers of labor productivity growth. As Figure A12 shows the role of technologies, including of their components, and of inputs for productivity growth in the broad goods sector (left column) and in the narrower goods sector (right column) are very similar. Moreover, the implications for aggregate labor productivity growth allowing for a separate agricultural sector are virtually identical.

Table A10: Average annual growth rate of Zs 1960–2017, separate agriculture

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
Industry	5.33%	6.00%	1.43%	-3.02%	4.86%
Agriculture	-11.17%	0.21%	-5.57%	5.25%	-4.67%

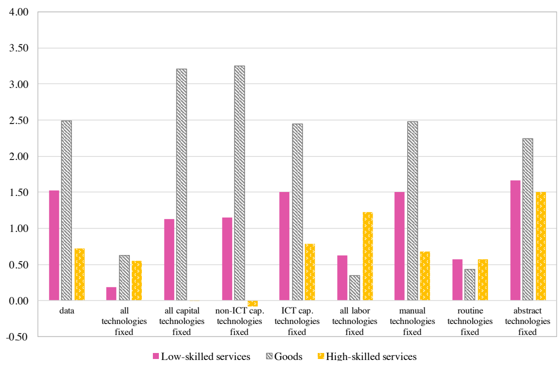

(a) Role of technologies, baseline

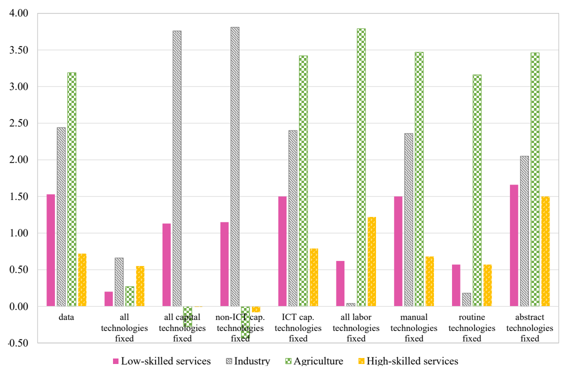

(b) Role of technologies, separate agriculture

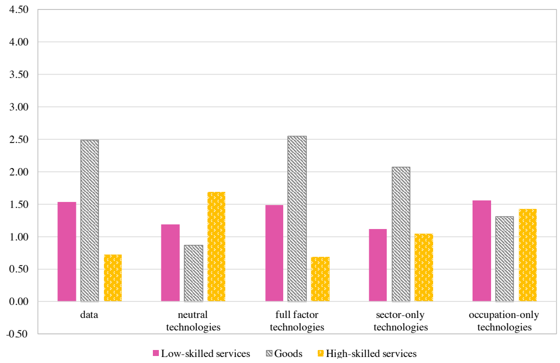

(c) Role of components, baseline

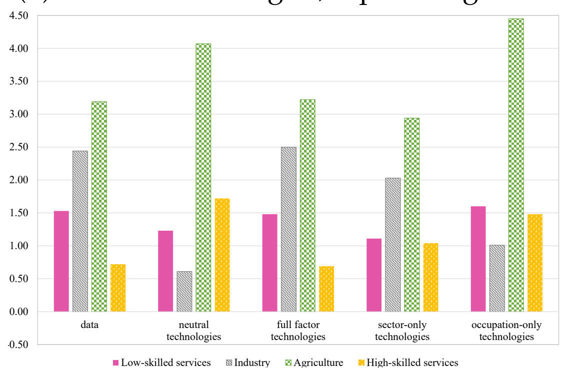

(d) Role of components, separate agriculture

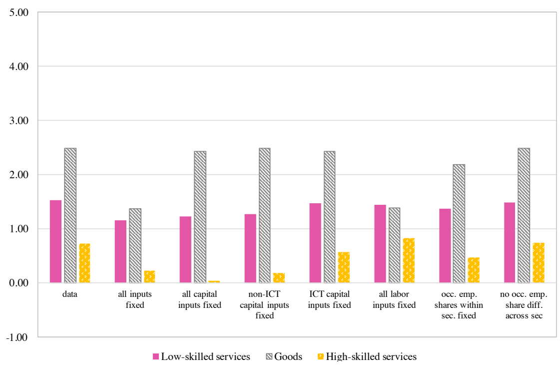

(e) Role of inputs, baseline

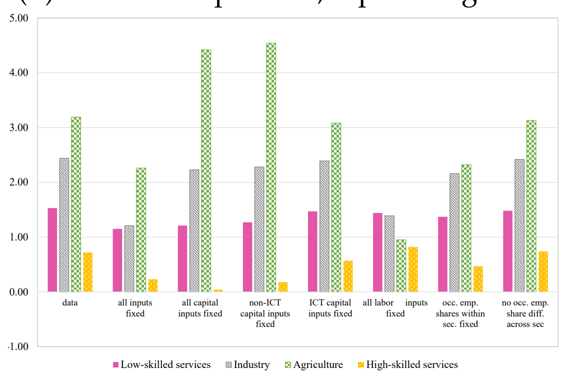

(f) Role of inputs, separate agriculture

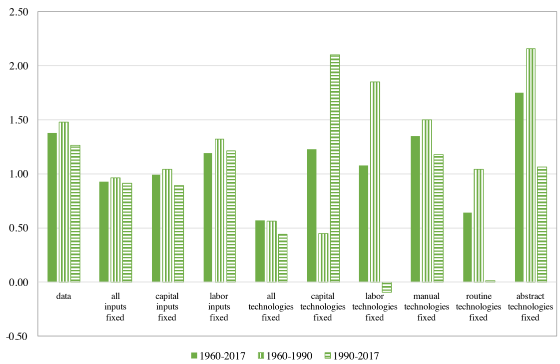

(g) Aggregate labor prod., baseline

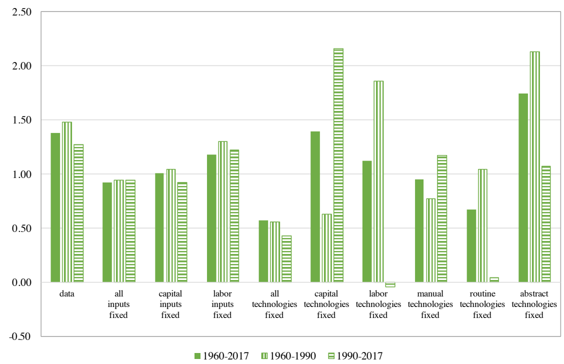

(h) Aggregate labor prod., sep. agriculture

Figure A12: Baseline classification vs separate agriculture

This figure shows the differences between the magnitudes of the various channels when separating agriculture relative to the baseline classification of sectors.