

Engines of Sectoral Labor Productivity Growth*

Zsófia L. Bárány[†]

Christian Siegel[‡]

March 2019

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 sector-specific routine labor augmenting technological change to be crucial. It is the most important driver of sectoral differences, and has a large and increasing contribution to aggregate labor productivity growth. Neither capital accumulation nor the occupational employment structure within sectors explains much of the sectoral differences.

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 and CEPR. Email: zsofia.barany@sciencespo.fr

[‡]University of Kent, School of Economics and Macroeconomics, Growth and History Centre. 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 2.49% in the goods sector, much higher than the 1.53% in low-skilled and the 0.72% in high-skilled services. However, there is no consensus on the origins of these differences. 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. Our results show that technological change has been very far from neutral. That we do not impose that sectoral technological change is factor-neutral, nor that factor-specific technological change is uniform across sectors, turns out to be crucial. Technologies have evolved at very differential rates, both across factors within each sector and across sectors for a given occupation or type of capital. In particular, amongst the labor-augmenting technologies those augmenting routine occupations have been growing the fastest in all sectors, but at very different rates: at 5.59% per year in goods, at 2.92% in low-skilled services and at 1.32% in high-skilled services.

Through a series of counterfactual simulations, we study the role of inputs and of technological change in labor productivity growth. We find that the single most important driver of sectoral labor productivity growth differences are the sector-specific growth rates of routine labor augmenting technologies. Without sector-specific routine labor augmenting technological change, labor productivity growth would have been almost equalized across sectors. Specifically, this type of 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. Our result that sectoral differences in technological change are crucial therefore lends support to the mechanism of Ngai and Pissarides (2007). Moreover, in terms of labor productivity growth in the aggregate, we show 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.

These counterfactuals also allow us to evaluate the role of various other channels proposed in the literature for sectoral productivity growth differences, such as differential capital intensities and capital accumulation (e.g. Acemoglu and Guerrieri (2008)), or differential sectoral intensities in occupational employment and technological change specific to occupations (Duernecker and Herrendorf (2016), Lee and Shin (2017)). We show that while 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 observed in the data. Similarly we find that in terms of occupational employment structure, differences across sectors as well as changes over time within sectors hardly matter for sectoral labor productivity differences.

In the public debate there is a growing concern about the effects of routinization and of new technologies on the labor market, and in particular on wage inequality and unemployment. To mitigate these – potentially detrimental – effects, active labor market policies, such as training programs targeted at workers of specific occupations, and policies aiming at maintaining certain industries have been advocated. However, a better understanding of the nature of technological change is required to inform this debate and to evaluate such policies. Our framework is useful in this context, as it identifies the differential paths of the various sector-and-factor specific technologies. We believe this is a necessary first step in understanding the drivers of labor demand for workers in various occupations and sectors. Our finding that occupation-augmenting technological change varies across sectors suggests that policies that target specific occupations or specific industries might be less suitable than sector-occupation specific policies.

In our model we consider different occupations as distinct labor inputs for a variety of reasons. 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. This implies that differences in the occupational

composition of sectors – a feature of the data which we show in section 2 – might affect average sectoral labor productivity growth (Duernecker and Herrendorf (2016), Lee and Shin (2017)). 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. As routine tasks are repetitive and easy to computerize, improvements in ICT knowledge or capital are likely to substitute for routine workers (Autor, Levy, and Murnane (2003)). This so-called routinization hypothesis is the main explanation for employment polarization, the shift out of routine occupations into manual (non-routine non-cognitive) and abstract (non-routine cognitive) jobs.¹ In our analysis we therefore differentiate between manual, routine and abstract labor inputs.

Our model also features capital inputs, as the accumulation of capital potentially is another important driver of labor productivity growth. If capital intensities differ across sectors, capital deepening induces structural transformation,² as argued by Acemoglu and Guerrieri (2008), and results in sectoral differences in the growth of labor productivity. As our model features capital inputs, we are able to evaluate the role of this channel. Similarly to Aum, Lee, and Shin (2018) and Eden and Gaggl (2018), we distinguish between ICT and non-ICT capital, and allow for them to have different degrees of substitutability with the various types of labor.

While observing factor inputs and output allows the computation of a neutral productivity,³ it is not possible to infer factor-augmenting technologies without making assumptions about the structure of production. Assuming a nested CES production function in all sectors and perfect competition, we infer from firm optimality conditions the sector-specific factor-augmenting technology parameters in each period. The share of income going to each factor of production and factor prices pin down rel-

¹In fact, Acemoglu and Autor (2011) argue that labor market polarization warrants to move beyond models that distinguish only between skilled and unskilled workers. In Bárány and Siegel (2018) we documented that labor market polarization in the United States started as early as 1950/1960.

²That differential sectoral intensities in production factors complementary with labor, coupled with aggregate growth in these factors, could lead to structural transformation was first proposed by Caselli and Coleman (2001) in the context of human capital.

³This 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.

ative technologies within each sector. The evolution of real value added by sector pins down the within sector changes in technology over time. To implement this, we combine data from the U.S. Bureau of Economic Analysis (BEA) and EU KLEMS 2017 to get sectoral value added and its components, sectoral prices, sectoral employment and capital (by type). Importantly, we need information on the occupations of workers within each sector, which is not available from the BEA or EU KLEMS. For this we use the US Census and American Community Survey (ACS) data between 1960 and 2017.

Similarly to us, Aum et al. (2018) also model sectoral production as a function of occupational labor and ICT and non-ICT capital, but the focus of their paper and their modeling choices are quite different to ours. They find in their calibrated model that computerization and routinization are distinct and important drivers of the recent productivity slowdown, but that sectoral TFP differences have only a limited role. We rely on factor income shares (among other observables) from the data to infer technologies, and while we also find a distinct role for improvements in routine-labor and ICT capital augmenting technologies, our results highlight that sectoral differences in factor-augmenting technological change are important.

Our results have implications for the supply side drivers of structural transformation. We do not find evidence for capital accumulation to be important, the mechanism suggested in Acemoglu and Guerrieri (2008). Similar to Herrendorf, Herrington, and Valentinyi (2015), our findings suggest that differences in technological progress across sectors are key, as proposed in Ngai and Pissarides (2007). According to Herrendorf et al. (2015) it is the differences in labor-augmenting technological progress across sectors that are crucial. As the composition of occupational employment varies across sectors, these sectoral differences potentially could be due to technological change biased only across occupations. However, in our framework we show that this is not the case and that there are substantial sectoral differences in occupation-augmenting technological change.

Similarly to Katz and Murphy (1992) and Krusell, Ohanian, Ríos-Rull, and Violante (2000) we assume a (nested) CES production function with different types of labor inputs. Both these papers focus on skilled and unskilled labor and impose a specific process for factor-biased technological change – this is what allows them to estimate

the elasticity of substitution. In contrast, we consider occupational labor inputs and we do not impose any restrictions on technological change. Similar in methodology to Caselli (2005) and Caselli and Coleman (2006), we extract factor productivities from the data taking values for the elasticities from the literature. Our methodology is also close to Buera, Kaboski, and Rogerson (2015) in relying on optimality conditions to infer technological change from the data. We find that technological change is very far from neutral, echoing the general conclusions of Caselli (2016).⁴

Our paper relates to the recent literature that connects the phenomena of structural change and polarization across occupations. Duernecker and Herrendorf (2016) show in a two-sector two-occupation model that unbalanced occupational productivity growth by itself provides dynamics consistent with structural change and with the trends in occupational employment. Lee and Shin (2017) allow for occupation-specific productivity growth and find that their calibrated model can quantitatively account for polarization as well as for structural change, and in an extension find a limited role for sector-specific technological change. In Bárány and Siegel (2018) we show that forces behind structural change, i.e. differences in productivity growth across sectors, lead to polarization of wages and employment at the sectoral level, which in turn imply polarization in occupational outcomes. Relative to these papers, the key difference is that we do not a priori restrict technological change to be biased in a particular way, and we find that technological change has been biased both across occupations and across sectors.

Several papers have emphasized the role of sectoral productivity differences for aggregate productivity (e.g. Duarte and Restuccia (2010), Duernecker, Herrendorf, and Valentinyi (2017), and Duarte and Restuccia (2019)). We add to this literature in two ways. First, we analyze the driving forces behind sectoral labor productivity growth. Second, we study the role of sectoral inputs and technologies in aggregate labor productivity growth. We find that also in the aggregate, technological change is much more important than input use for labor productivity growth. Moreover, we show that the contribution of labor-augmenting – and in particular sector-specific routine-

⁴While 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).

augmenting – technological change has increased over time.

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 demonstrate that our results are robust to alternative values for the elasticities of substitution and when controlling for workers' human capital. The final section concludes.

2 Factor use and factor income shares by sector

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.⁵ 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.⁶ It is common to split services, as already in 1947 the service industries as a whole constituted 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 ser-

⁵The industry categories in this dataset are based on the North American Industry Classification System (NAICS)).

⁶Services 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.

vices, 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




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.

shows the evolution of value added shares, and it is clear that while 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. In terms of real value added, depicted in the middle panel, low-skilled services grew the fastest, and goods the slowest. This divergence between the path of nominal and real quantities can be understood from the steep increase in the price of high-skilled services relative to the other two sectors and from the near constancy of the relative price of low-skilled services to goods, 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 BEA. We calculate the labor income share as:⁷

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

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

⁷This definition of the labor income share excludes proprietors' income. We choose to do this for two reasons. First, Elsbj, 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.

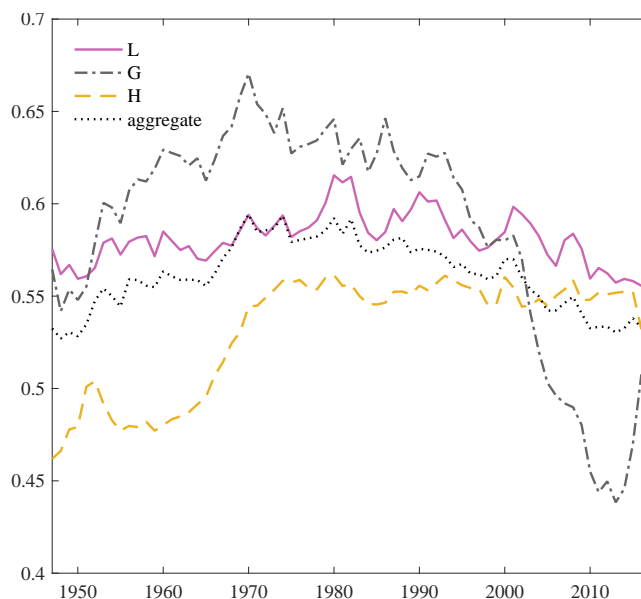


Figure 2: Labor income share by sector

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.

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.⁸ Figure 2 plots the evolution of the labor income share by sector as well as for the aggregate economy. The labor income share in the economy as a whole increased until the early 1970s, which was followed by a virtually equal reduction thereafter.⁹ 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

⁸Herrendorf et al. (2015) also combine data on the labor income and employment shares across different industries based on the SIC and the NAICS classification.

⁹When 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.

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 computer capital from the BEA we include *Information processing equipment* and *Software*, while traditional capital comprises of all other non-residential capital.¹⁰ 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 traditional 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.




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

Notes: The left and middle panels are computed based on data from the BEA Fixed Asset Accounts, while the data for the right panel is calculated from EU KLEMS.

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

¹⁰Traditional 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.

types of capital by industry (based on the International Standard Industrial Classification of All Economic Activities (ISIC)) from 1970 onwards. When constructing the allocation of computer 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 computer capital stock in each sector, and shows that most of the computer 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 computer 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 Barany and Siegel (2018). We then classify each worker into one of these three broad occupations and into one of the three sectors defined earlier.¹¹ 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.¹² To get the employment share of a sector-occupation cell, l_{oJ} , we multiply the within-sector hours share of occupation

¹¹See Appendix A.1 for more details on the classification of occupations and Table A1 for industries.

¹²As 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).

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:

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

Relative average occupational wages within sectors can then be calculated as

$$\frac{w_o}{w_r} = \frac{\theta_{oJ} l_{rJ}}{\theta_{rJ} l_{oJ}}.$$

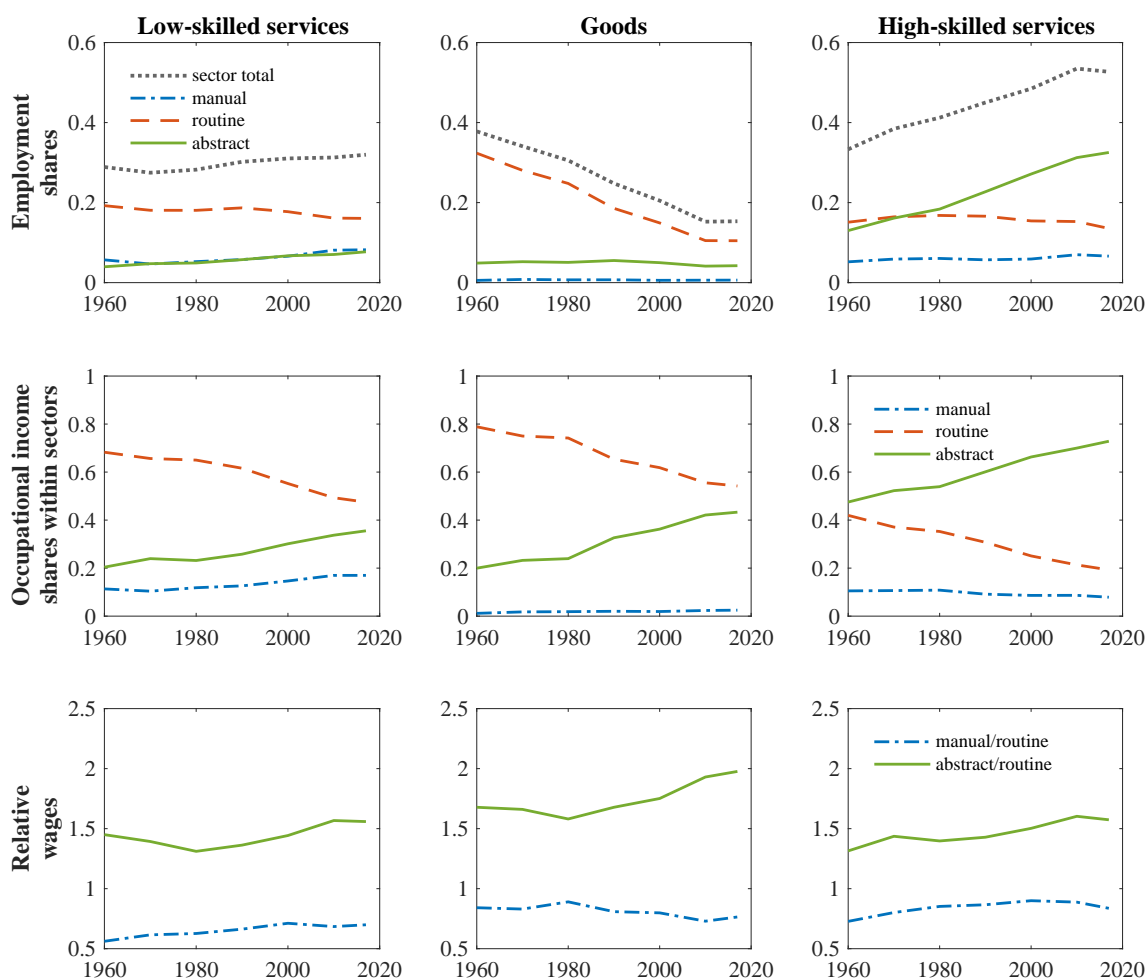


Figure 4: Sector-occupation income shares, hours shares, and relative wages

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.

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).¹³

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 hourly wages are not equalized across sectors.

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 input 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 role of changing factor inputs and changing factor-augmenting technologies, we specify a production side framework. We assume a relatively flexible CES functional form for sectoral production, which allows matching the data we documented in the previous section. Note that with CES production functions rela-

¹³In 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.

tive 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)). 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. (2015), 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

There are three sectors in the economy which respectively produce low-skilled services (L), goods (G), and high-skilled services (H). Firms in each sector combine occupational labor (manual, routine and abstract), computer capital and traditional capital as inputs according to the following CES production function:

$$Y_{J,t} = \left[\left(\sum_{o=m,a} (\alpha_{oJ,t} l_{oJ,t})^{\frac{\rho-1}{\rho}} + \left[(\alpha_{rJ,t} l_{rJ,t})^{\frac{\sigma_c-1}{\sigma_c}} + (\alpha_{cJ,t} c_{J,t})^{\frac{\sigma_c-1}{\sigma_c}} \right]^{\frac{\sigma_c(\rho-1)}{(\sigma_c-1)\rho}} \right)^{\frac{\rho(\sigma-1)}{(\rho-1)\sigma}} + (\alpha_{kJ,t} k_{J,t})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

for $J \in \{L, G, H\}$. (2)

In this formulation $l_{oJ,t}$ is occupation $o \in \{m, r, a\}$ labor, $c_{J,t}$ is computer capital and $k_{J,t}$ is traditional capital used in sector J , and $\alpha_{fJ,t} > 0$ is a *sector-specific factor-augmenting* technology term for each production factor, all in period t . The production function is of a nested CES form, where the most inner level is the combination of routine la-

bor and computer capital with an elasticity of substitution σ_c . Next the different types of labor, including the most inner routine aggregate, are aggregated according to an elasticity of substitution ρ , and the outer-most layer combines aggregate labor and traditional 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 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)). In line with Autor and Dorn (2013), the most inner nest is such that the use of ICT capital most directly affects routine workers, whereas according to the middle nest it would complement workers in the other occupations when $0 < \rho < 1$.¹⁴

It is worth to note that the initial value of technology $\alpha_{fJ,0}$ reflects the initial productivity as well as the intensity at which sector J uses input f , whereas any subsequent change over time, $\alpha_{fJ,t}/\alpha_{fJ,0}$, reflects sector-factor augmenting technological change.¹⁵ This formulation of the production function is very flexible and does not impose any restrictions on the nature of technological change. In particular, it does not require taking a stance on whether labor-augmenting technological change is specific to sectors or occupations.¹⁶ It also allows for systematic co-movements and can

¹⁴Since 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.

¹⁵An alternative, isomorphic way of writing the production function in (2) is

$$Y_{J,t} = \left[\left(\sum_{o=m,a} x_{oJ} (A_{oJ,t} l_{oJ,t})^{\frac{\rho-1}{\rho}} + \left[x_{rJ} (A_{rJ,t} l_{rJ,t})^{\frac{\sigma_c-1}{\sigma_c}} + x_{cJ} (A_{cJ,t} c_{J,t})^{\frac{\sigma_c-1}{\sigma_c}} \right]^{\frac{\sigma_c(\rho-1)}{(\sigma_c-1)\rho}} \right)^{\frac{\rho(\sigma-1)}{(\rho-1)\sigma}} + x_{kJ} (A_{kJ,t} k_{J,t})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

where x_{fJ} are constant weights and $A_{fJ,t}$ are sector-factor technologies that can change over time. The two formulations are equivalent since one can rewrite $\alpha_{fJ,t} = x_{fJ}^{\frac{\eta}{\eta-1}} A_{fJ,t}$ (where η is the relevant elasticity depending on the layer of the CES nest). In this sense the $\alpha_{fJ,t}$ terms comprise of a fixed weight and a changing sector-specific factor augmenting technology. We are interested in changes in technology over time, which – due to the weights being constant – are equal in the two formulations, $\Delta \log \alpha_{fJ,t} = \Delta \log A_{fJ,t}$.

¹⁶It is easy to conceive that some technologies improve a given occupation’s productivity in a similar way regardless of the sector of work. For example an accountant’s productivity has increased by the advent of computers, though potentially more so in sectors characterized by larger firms. There are also occupations which – though similar – perform different tasks depending on the sector of work.

capture general purpose technologies, sector-specific innovations, occupation-biased technological change, or changes specific to the sector-factor cell.

3.2 Inferring factor-augmenting technologies

The assumptions about the production side of the economy allow us to infer factor-augmenting technologies (the $\alpha_{f,J,t}$ 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.¹⁷ In what follows, where possible, we omit the time t subscripts to simplify the notation.

In line with the data we have shown in Figure 4, we assume that wages are sector-occupation specific, we denote these by w_{oJ} . Assuming further that the rental rates of computer (R_c) and traditional 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 (2), where p_J denotes the price of sector J output. Optimal input use in each

Ford's Model T is a good example: by introducing the moving assembly line in production, rather than the then usual hand crafting, the productivity of workers directly producing the car increased, later leading to a spillover to workers in other car producers. This did not have a concurrent effect on other production workers. In this sense the introduction of assembly lines in car manufacturing can be regarded as a sector-occupation-specific productivity change.

¹⁷The derivations can be found in appendix A.3.

sector has to satisfy the following first order conditions:

$$\frac{\partial \pi_J}{\partial l_{oJ}} = p_J Y_J^{\frac{1}{\sigma}} (LA)^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma} - 1} \alpha_{oJ}^{\frac{\rho}{\rho-1}} l_{oJ}^{-\frac{1}{\rho}} - w_{oJ} = 0 \quad \text{for } o \in \{m, a\}, \quad (3)$$

$$\frac{\partial \pi_J}{\partial l_{rJ}} = p_J Y_J^{\frac{1}{\sigma}} (LA)^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma} - 1} [RA]^{\frac{\sigma_c}{\sigma_c-1} \frac{\rho-1}{\rho} - 1} \alpha_{rJ}^{\frac{\sigma_c-1}{\sigma_c}} l_{rJ}^{-\frac{1}{\sigma_c}} - w_{rJ} = 0, \quad (4)$$

$$\frac{\partial \pi_J}{\partial c_J} = p_J Y_J^{\frac{1}{\sigma}} (LA)^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma} - 1} [RA]^{\frac{\sigma_c}{\sigma_c-1} \frac{\rho-1}{\rho} - 1} \alpha_{cJ}^{\frac{\sigma_c-1}{\sigma_c}} c_J^{-\frac{1}{\sigma_c}} - R_c = 0, \quad (5)$$

$$\frac{\partial \pi_J}{\partial k_J} = p_J Y_J^{\frac{1}{\sigma}} \alpha_{kJ}^{\frac{\sigma-1}{\sigma}} k_J^{-\frac{1}{\sigma}} - R_k = 0, \quad (6)$$

where we define the *routine aggregate* as $RA = (\alpha_{rJ} l_{rJ})^{\frac{\sigma_c-1}{\sigma_c}} + (\alpha_{cJ} c_J)^{\frac{\sigma_c-1}{\sigma_c}}$ and the *labor aggregate* as $LA = (\alpha_{mJ} l_{mJ})^{\frac{\rho-1}{\rho}} + (\alpha_{aJ} l_{aJ})^{\frac{\rho-1}{\rho}} + [RA]^{\frac{\sigma_c}{\sigma_c-1} \frac{\rho-1}{\rho}}$.

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, (3), pin down the optimal relative labor use as:

$$\frac{l_{aJ}}{l_{mJ}} = \left(\frac{w_{mJ}}{w_{aJ}} \right)^{\rho} \left(\frac{\alpha_{aJ}}{\alpha_{mJ}} \right)^{\rho-1}. \quad (7)$$

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 $(\alpha_{aJ}/\alpha_{mJ})^{\rho-1}$ is larger it is optimal to use relatively more abstract labor in that sector. Multiply the above by w_{aJ}/w_{mJ} and re-arrange to get:

$$\frac{\alpha_{mJ}}{\alpha_{aJ}} = \frac{w_{mJ}}{w_{aJ}} \left(\frac{w_{mJ} l_{mJ}}{w_{aJ} l_{aJ}} \right)^{\frac{1}{\rho-1}} = \frac{w_{mJ}}{w_{aJ}} \left(\frac{\theta_{mJ}}{\theta_{aJ}} \right)^{\frac{1}{\rho-1}}, \quad (8)$$

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 (8) 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.

Similarly, from the first order conditions on routine labor and computer capital, (4) and (5), we can express the relative demand for these factors, and consequently their relative α as well:

$$\frac{\alpha_{cJ}}{\alpha_{rJ}} = \frac{R_c}{w_{rJ}} \left(\frac{\Theta_{cJ}}{(1 - \Theta_J)\theta_{rJ}} \right)^{\frac{1}{\sigma_c - 1}}, \quad (9)$$

where $\Theta_{cJ} = (R_c c_J)/(p_J Y_J)$ is the share of income in sector J paid to computer 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 (8), except that it is a different elasticity of substitution that is relevant.

Expressing the remaining two relative technology levels within sectors, α_{rJ}/α_{mJ} and α_{kJ}/α_{mJ} , follows a similar principle, but is slightly more convoluted, and we delegate the details of these derivations to appendix A.3. Here we only explain the intuition. First, from the optimal use of routine labor relative to ICT capital, we express RA , 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:

$$\frac{\alpha_{mJ}}{\alpha_{rJ}} = \frac{w_{mJ}}{w_{rJ}} \left(\frac{\theta_{mJ}}{\theta_{rJ}} \right)^{\frac{1}{\rho - 1}} \left[1 + \frac{\Theta_{cJ}}{(1 - \Theta_J)\theta_{rJ}} \right]^{\frac{\rho - \sigma_c}{(\sigma_c - 1)(\rho - 1)}}. \quad (10)$$

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

$$\frac{\alpha_{kJ}}{\alpha_{mJ}} = \frac{R_k}{w_{mJ}} \left(\frac{1}{\theta_{mJ}} \right)^{\frac{1}{\rho - 1}} \left(\frac{\Theta_J - \Theta_{cJ}}{1 - \Theta_J} \right)^{\frac{1}{\sigma - 1}} \left(1 + \frac{\Theta_{cJ}}{1 - \Theta_J} \right)^{\frac{\sigma - \rho}{(\rho - 1)(\sigma - 1)}}. \quad (11)$$

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 α_{kJ} as the base technology, all other factor-augmenting technologies in sector J are proportional to α_{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 α s 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 α s within a period. Plugging all the optimal relative input use expressions in (2) sectoral output can be expressed as:

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

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

$$\frac{\alpha_{kJ,t}}{\alpha_{kJ,0}} = \gamma_{J,t} \frac{k_{J,0}}{k_{J,t}} \left(\frac{\Theta_{J,t} - \Theta_{cJ,t}}{\Theta_{J,0} - \Theta_{cJ,0}} \right)^{\frac{\sigma}{\sigma-1}}, \quad (12)$$

where $\gamma_{J,t}$ denotes the growth of sectoral real value added between the initial and period t in the data. Again in equation (12) all right-hand side variables can be observed in the data, and hence, conditional on σ , this equation gives us the growth rate of $\alpha_{kJ,t}$ over time.

Finally, we need to pin down the initial level of α s. 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.¹⁸ We infer these initial α s from initial sectoral prices. Using the above expression for sectoral output in the first order condition on traditional capital we get:

$$\alpha_{kJ,0} = \frac{R_{k,0}}{p_{J,0}} (\Theta_{J,0} - \Theta_{cJ,0})^{\frac{1}{\sigma-1}}. \quad (13)$$

Equations (8), (9), (10), (11), (12) and (13) describe how to infer factor-augmenting technologies in each sector and in all periods. Note that equations (8) to (12) 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.

¹⁸Even for this, only the relative initial α s matter, i.e. we could normalize one of the sectors' $\alpha_{kJ,0}$ without loss of generality.

3.3 Implementation

To infer the sector-specific factor-augmenting technologies from the data using equations (8) to (13), 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.¹⁹ 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 for our baseline parametrization, but in the robustness checks of section 5.2 we also explore model variants with sector-specific elasticities.

Second, we need the elasticity of substitution between computer capital and routine labor, σ_c . While the literature has argued that routine labor and computer 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 computer capital and routine labor is not constant, but it ranges between 2.14 and 3.27. Aum et al. (2018) calibrate industry specific elasticities between ICT capital and all types of occupational labor and find values between 1.21 and 1.84. 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 $\rho = 0.70$ and Aum et al. (2018) calibrate 0.81 both among 10 occupations, and Duernecker and

¹⁹These 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.

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 $\rho = 0.6$. 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.

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 traditional and computer 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 computer capital ($\Theta_{cJ,t}$), and to both types of capital together ($\Theta_{J,t}$), the quantity of traditional capital by sector ($k_{J,t}$), the per worker growth rate of sectoral value added ($\gamma_{J,t}$), 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 traditional 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 (16) and the discussion that follows in appendix A.2). 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 A.2. These accounting identities ensure that the sum of all factor incomes is equal to nominal value added.

4 The role of changing input use and technologies

Table 1 shows the average annual growth rate of sector-factor augmenting technologies between 1960 and 2017, as well as for two sub-periods, 1960-1990 and 1990-2017. Technological change has been uneven, within each sector across factors, as well as for a given factor across sectors. Nonetheless some patterns can be discerned. Looking at the average growth rates over the entire period, it is obvious that among all occupations routine labor had the highest productivity growth in all sectors, between 1.32 and 5.59 percent annually. Technological change augmenting manual labor was much more modest and less dispersed across sectors, with rates between 0.25 and 0.67 percent. Finally, technological change augmenting abstract labor varied across sectors, with negative growth rates in L and in H . These negative growth rates might be explained by a compositional change within abstract occupations in these sectors, towards more time-consuming tasks. In terms of capital-augmenting technologies we find that those related to ICT increased rapidly in L and in G and fell in H , while those augmenting traditional capital increased at a lower rate in L and in H and they fell in G . While these negative growth rates might be surprising, they are in line with what previous literature has found.²⁰ In terms of sectoral patterns, the growth rates of all factor-augmenting technologies were the highest in G , followed by L , except for manual labor and traditional capital, which had the highest growth in H . Thus beyond the factor-specific patterns, there also seem to be sector-specific components to technological progress.

Our results highlight that routine workers became more productive over and beyond what is embodied in ICT capital. 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 (α_{cJ}), suggesting that there is routinization beyond computerization, in line with what Aum et al. (2018) find.

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 α_{mL} , the growth rate of which remained virtually constant), while technological

²⁰Both 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.

Table 1: Average annual growth rate of α s for various periods between 1960 and 2017

	occupations			capital	
	manual	routine	abstract	non-ICT	ICT
1960-2017					
<i>L</i>	1.0025	1.0292	0.9933	1.0085	1.0200
<i>G</i>	1.0058	1.0559	1.0100	0.9839	1.0439
<i>H</i>	1.0067	1.0132	0.9763	1.0178	0.9803
1960-1990					
<i>L</i>	1.0032	1.0108	0.9849	1.0286	1.0338
<i>G</i>	0.9760	1.0369	0.9827	1.0063	1.0497
<i>H</i>	0.9921	0.9882	0.9587	1.0394	0.9653
1990-2017					
<i>L</i>	1.0017	1.0500	1.0027	0.9867	1.0049
<i>G</i>	1.0398	1.0775	1.0413	0.9596	1.0376
<i>H</i>	1.0231	1.0417	0.9962	0.9942	0.9974

change augmenting either type of capital decelerated (except for ICT-augmenting capital in *H*). This suggests that the relative importance of capital- vs labor-augmenting technologies for labor productivity growth has changed over the last decades.

To the extent that positive labor productivity growth in the data has stemmed from improvements in technologies, Table 1 implies that this was most likely due to improvements in routine labor-augmenting technologies. It is worth to note that the ratios of the growth rates of routine-augmenting technologies across sectors are very similar to those of measured labor productivity (1.53% in *L*, 2.49% in *G* and 0.72% in *H*). How the growth rate of individual factor-augmenting technologies affects sectoral labor productivity depends on the intensity at which the various factors are used. As shown in Figure 4, sector *G* has had the highest intensity in routine workers, which could have amplified the effects coming from the differential evolution of sector-occupation cell technologies (as α_{rG} grew the most). Thus the sectoral differences both in the growth rate of routine-augmenting technologies and in the occupational composition of employment, as well as their interaction could be behind the sectoral differences in labor productivity growth.

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 quantify the role of changing input

use and of differences in occupational employment shares across sectors. To do this, we use the α s as extracted from the data, and fix factor inputs at counterfactual levels. Second, to assess the importance of technological change, we use the actual factor inputs as measured in the data, but fix the α s at their 1960 values. 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 changing input use

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 2.49% the highest labor productivity growth, whereas in low-skilled services it was 1.53% and in high-skilled services 0.72%. The subsequent sets of bars show the results of various counterfactuals in which we fix the inputs listed below the bars at their 1960 values, but let all other inputs and the factor productivities ($\alpha_{fJ,t}$) vary over time as extracted from the data (except for the last set, where we assign identical occupational employment shares to all sectors). Comparing the implied sectoral labor productivity growth (and their differences) to the data gives a sense of the importance of the changing use of the fixed input(s).

In the second set of bars ('all') we fix all inputs at their 1960 level. Keeping all inputs at their initial level results in much 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.




Figure 5: 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.

This highlights that changing input use is important for the level of labor productivity growth, as well as for its differences across sectors.

The next three counterfactuals shed light on the role of capital accumulation. With both types of capital inputs fixed at their 1960 level ('capital'), the growth rate in all sectors falls short of the data, on average by 39 percent.²¹ 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 accumulation resulted in smaller sectoral differences, but without altering the ranking of sectors in terms of labor productivity growth. 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

²¹Labor productivity growth in *L* would have been 80% of its actual value, in *G* 98%, and in *H* 5%, the simple average of this is 61%, i.e. 39 % lower than in the data.

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.

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 ('labor'). 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 G and for sectoral differences. In the last two counterfactuals we investigate whether this was driven by differences in the occupational employment structure, either over time within sectors or across sectors.

In the penultimate set of bars, we fix the share of occupations within each sector at initial ratios ('occ. shares within sec.') but let the overall employment share of each sector (as well as all other inputs and technologies) evolve as in the data. In this case we obtain 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 only modest positive effects on sectoral labor productivity growth, but hardly any effect on sectoral differences.

In the last set of bars rather than fixing an input at the sectors' initial level (or share), 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.

4.2 The role of technological change

We now turn our attention to technological change and conduct a series of counterfactuals in which we fix some of the factor-augmenting technologies but let all inputs vary as in the data between 1960 and 2017. Again comparing the results implied by the counterfactual to the actual data informs us about the importance of the technological change that we shut down. The first set of bars in Figure 6 shows the data and the subsequent sets of bars fix the technologies (the $\alpha_{fJ,t}$) of the factors listed below the bars at their 1960 level.




Figure 6: 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') shows, there is hardly any growth in labor productivity in low-skilled services and in goods production and only very small differences across sectors. This clearly demonstrates that technological progress was crucial for the level of labor productivity growth as well as

for its sectoral differences. Furthermore, comparing the results of fixed technologies to those of fixed inputs in Figure 5 (second set of bars) highlights that evolving technologies matter much more than changing inputs, both for sectoral growth rates and their differences. In particular evolving technologies explained at least 76% of labor productivity growth in L , 55% in G and 33% in H .²² High-skilled services thus seem to be somewhat of an exception; in this sector capital accumulation was 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 ('capital') 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 sectoral differences. 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, we see that these results are mainly driven by the evolution of traditional capital's productivity, and not by ICT capital.

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 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, which was mainly driven by capital accumulation as we established in section 4.1.²³ Sector-specific routine labor augmenting technological

²²These 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.

²³To obtain these numbers we conducted an additional counterfactual, where we fixed everything at the 1960 level except for $\alpha_{r,J,t}$ 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 $\alpha_{r,J}$ (the 'routine' counterfactual of Figure 6).

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 to some extent, manual labor augmenting technologies hardly had any impact on the level of and on the differences in sectoral labor productivity growth. This is perhaps not surprising given the low growth rates of these technologies shown in Table 1.

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 sectoral differences. 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-augmenting technological change.

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

As we found a key 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 1, 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.²⁴ In particular we regress the change in log cell technologies between each consecutive period on a (time-varying) sector effect ($\gamma_{J,t}$), an

²⁴In 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).

occupation effect ($\delta_{o,t}$), and a time effect (β_t) in the following way

$$\Delta \ln \alpha_{oJ,t} \equiv \ln \alpha_{oJ,t} - \ln \alpha_{oJ,t-1} = \beta_t + \gamma_{J,t} + \delta_{o,t} + \varepsilon_{oJ,t}, \quad (14)$$

where we use weights $\omega_{oJ,t}$ to reflect the relative importance of the sector-occupation cell.²⁵ 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 .²⁶

Based on the results of (14), we compute counterfactual series for $\Delta \ln \alpha_{oJ,t}$, 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 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 $\Delta \ln \alpha_{oJ}$:

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

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 $1 - R^2$.²⁷

²⁵The 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.

²⁶To 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 .

²⁷The R^2 is defined as

$$R^2 = \frac{\sum_{o,J,t} \omega_{oJ,t} (\widehat{\Delta \ln \alpha_{oJ,t}} - \overline{\Delta \ln \alpha})^2}{\sum_{o,J,t} \omega_{oJ,t} (\Delta \ln \alpha_{oJ,t} - \overline{\Delta \ln \alpha})^2},$$

	neutral	full factor	sector	occupation
Distance measure	0.702	0.033	0.227	0.408

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 (29.8 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.²⁸

Finally, we evaluate the role of the different components of labor-augmenting technological change for sectoral labor productivity growth. For this we build counterfactual labor-augmenting technology series by adding the predicted changes to the initial 1960 levels of the cell-specific labor augmenting technologies. Figure 7 shows the results for counterfactuals, in which 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. Not surprisingly, labor-augmenting technological change that is neutral across sector-occupation cells can account neither for sectoral differences, nor for the level of labor productivity growth. The counterfactual based on the ‘full factor’ prediction, on the other hand, replicates the observed labor productivity growth rates well. This highlights that the growth of labor-augmenting technologies is well described as the sum of neutral, sector-specific and occupation-specific components. However, 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,

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

²⁸In appendix A.4 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 for the occupation component it is smaller.




Figure 7: 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 (14) based on the components listed below the bars, with all inputs as well as all capital augmenting technologies evolving as in the data.

whereas the sector component alone gets closer in terms of these aspects but quantitatively falls short. Overall this analysis reveals that both sector and occupation components are important drivers of labor productivity growth at the sectoral level. Despite the marked differences in labor-augmenting technological change across occupations, shown in Table 1, the sectoral differences within these occupation-augmenting technologies seem to be key.

4.4 Implications for aggregate labor productivity growth

We established that while capital accumulation was important for the level of labor productivity growth, especially in sector H , 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. In Figure 8 we




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.

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’ and the ‘all tech.’ 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’) played a larger role than labor reallocation across sector-occupation cells (‘labor inputs’). 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 1990-2017.²⁹ 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 substitution elasticities. We also show that the results of a model variant with sector-specific substitution elasticity between capital and the labor aggregate are very similar. Finally, we describe how to control for observable worker characteristics in our framework and demonstrate that our conclusions are robust to accounting for worker efficiency.

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, one at a time, to alternative values. The general 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. In the appendix we list in Table A2 the average annual growth rate of the various sector-specific factor-augmenting

²⁹That 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 1.

technologies for the different elasticities that we consider. This table shows that the general patterns described in section 4 for Table 1 remain the same. We also show figures analogous to Figures 6, 7 and 8 for the various elasticities.³⁰ Since this analysis establishes that our results are qualitatively unchanged – except for the role of sector- vs. occupation-components – and even quantitatively very similar, here we only point out for each elasticity the biggest differences compared to our baseline results.

Elasticity between capital and labor. 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. The most noticeable difference when changing the elasticity of substitution between capital and labor is in terms of aggregate labor productivity growth. In Figure A4 in the appendix we see that the counterfactual growth rates when shutting down technological change augmenting labor between 1960-1990, and augmenting capital between 1990-2017 do not overshoot the actual ones. With lower substitution elasticity between capital and labor, optimality requires a more similar growth in the effective capital input and the effective labor aggregate. Given our framework this leads to a change in the technologies that we infer from the data. Table A2 shows that a lower σ requires that within a sector the technologies of traditional capital and the labor aggregate (formed by all occupations and ICT-capital) have to grow at a more similar rate. This explains why with lower σ the contribution of capital- and labor-augmenting technologies are more likely to go in the same direction.

Elasticity between occupational labor inputs. Next we vary only the elasticity of substitution between the occupational labor inputs (incl. the routine aggregate), ρ .³¹ The only visible difference relative to our baseline results is in the role of sector- and occupation-specific components of labor-augmenting technologies. As Table A3 in the appendix shows, the larger is ρ , 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

³⁰We do not show the equivalent of Figure 5, as this figure looks virtually identical for all elasticities we consider.

³¹Changing the value of ρ does not affect the growth rate of the $\alpha_{k,J}$ s at all, see Table A2 in the appendix.

the occupation-only technologies. This is also reflected when considering the role of these components in observed sectoral labor productivity growth (see Figure A3 in the appendix). 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 to match the data holds for all elasticities.

Elasticity between routine labor and ICT capital. We consider two alternative values for σ_c : 1.5 in the midrange of values calibrated in Aum et al. (2018), and 2.5 in the midrange of the values implied by the estimation in Eden and Gaggl (2018). It is important to note that the value of σ_c has no effect on the growth rate of technologies except for routine labor and ICT capital, and quantitatively the effect is mostly on ICT-augmenting technologies (see Table A2 in the appendix). Given this it is not surprising that we see hardly any effect of σ_c on sectoral or aggregate labor productivity growth. The only discernible change is quantitative: the smaller this elasticity, the larger is the impact of technologies on aggregate labor productivity growth.

5.2 Sectoral heterogeneity in elasticities between capital and labor

We next 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. The last set of rows in Table A2 show the growth rates of α_s , and we can see that in low- and high-skilled services these values are the same

as when setting the common σ to 0.75. The values obtained for the goods sector are different, but overall the table mimics the patterns of our baseline quite closely. Figure A5 compares the effects of the various channels with those in the baseline, 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.

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, $\bar{e}_{oJ,t}$ and wages per efficiency units of labor, as we explain in appendix A.5.2.³²

To incorporate efficiency units of labor into the model, we assume that firms choose $n_{oJ,t} \equiv \bar{e}_{oJ,t} l_{oJ,t}$ 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 (8) to (13) to infer sector-factor technologies, whereas the measurement of all other variables remains unchanged.

Figure A6 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,³³ whether accounting for efficiency units or not, though their levels are somewhat different. Since we identify the within-sector ratios of occupational productiv-

³²We construct this in two different ways, by including/not-including the residuals from the Mincer wage regression in $\bar{e}_{oJ,t}$. Note that, even though we calculate sector-occupation wage rates from our accounting identity (see equation (21) in the appendix) as before, the relative wages within sectors are the same as those implied by the the Mincer wage regression.

³³From 2000 onwards, in high-skilled services there is somewhat of a divergence between relative average ('raw') wages and relative wages controlling for workers' characteristics.

ities precisely from these relative wages, the general conclusions about the inferred technological change are very similar, as shown in Table A5. 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 A7 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.

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 clearly demonstrate that technological change has not been neutral. 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-biased 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 of productivity growth, but neither contributed to the sectoral differences observed in 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 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 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, the job of a worker is not only described by the occupation, but also by the sector (or industry) of work. It is easy to see that the tasks performed in a very specific occupation, e.g. a cleaner, depend on whether the individual works in a car manufacturing plant or in the offices of a law firm. Thus 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.

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.
- 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.
- 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, and Richard Rogerson. Skill Biased Structural Change. Working Paper 21165, National Bureau of Economic Research, May 2015.
- 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 U.S. Structural Transformation and Regional Convergence: A Reinterpretation. *Journal of Political Economy*, 109(3): 584–616, 2001.
- 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.
- 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.

A Appendix

A.1 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.

Table A1: Classification of industries into sectors

	NAICS	SIC	ISIC	IND1990
Low-skilled services	<ul style="list-style-type: none"> Wholesale trade Retail trade Transportation & warehousing Arts, entertainment, recreation, accommodation & food serv. Other serv., except government 	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> 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) 	<ul style="list-style-type: none"> Wholesale & retail trade Transportation Entertainment Low-skilled business serv. Personal serv.
Goods	<ul style="list-style-type: none"> Agriculture, forestry, fishing & hunting Mining Construction Manufacturing 	<ul style="list-style-type: none"> Agriculture, forestry, & fishing Mining Construction Manufacturing 	<ul style="list-style-type: none"> Agriculture, forestry & fishing (A) Mining & quarrying (B) Construction (F) Total manufacturing (C) 	<ul style="list-style-type: none"> Agriculture, forestry & fishing Mining Construction Manufacturing
High-skilled services	<ul style="list-style-type: none"> Utilities Information Finance, insurance, real estate, rental & leasing Professional & business serv. Educational serv., health care & social assistance Government 	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> Electricity, gas & water supply (D-E) Information & communication (I) 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) 	<ul style="list-style-type: none"> Utilities Communications Finance, insurance & real estate Professional serv. High-skilled business serv. Public administration

A.2 Data Appendix

Capital targets. To back out all α s we need the rental rate of traditional, R_k , and of computer capital, R_c , the share of income going to both types of capital, Θ_J , and to computer capital alone, Θ_{cJ} , as well as the amount of traditional 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 traditional and computer 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 traditional and computer capital in 2009 is equal to the share of traditional and computer 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 traditional and computer 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 traditional 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:

$$R_k k + R_c c = Y^{nom} \cdot \sum_J VA_J \Theta_J, \quad (15)$$

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 traditional and computer 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 traditional capital in the next year. From these two equations we can calculate in each period the rental rates of traditional and of computer capital, R_k and R_c .

We calculate the allocation of computer 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 computer capital (per worker) in each sector is then obtained as $c_J = c \cdot \tilde{c}_J$. The share of income going to computer 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 traditional capital in each sector, k_J , can then be calculated from (15).

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

$$w_{oJ} l_{oJ} = Y^{nom} \cdot VA_J (1 - \Theta_J) \theta_{oJ}, \quad (16)$$

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 θ s and occupational employment shares, and therefore is equal to the relative wage observed in the Census/ACS data.

A.3 Derivations

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

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

$$\frac{\alpha_{cJ}l_{cJ}}{\alpha_{rJ}c_J} = \left(\frac{w_{rJ}\alpha_{cJ}}{R_c\alpha_{rJ}} \right)^{\sigma_c} = \left(\left[\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}} \right]^{\frac{1}{\sigma_c-1}} \right)^{\sigma_c} = \left(\frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}} \right)^{\frac{\sigma_c}{\sigma_c-1}}, \quad (17)$$

where the first equality comes from multiplying the relative optimal input use with the relative α s, and the second one comes from multiplying it with relative factor prices (and re-arranging). Using the above expression implies that at the optimum we can express the routine aggregate as:

$$RA = \left[(\alpha_{rJ}l_{rJ})^{\frac{\sigma_c-1}{\sigma_c}} + (\alpha_{cJ}c_J)^{\frac{\sigma_c-1}{\sigma_c}} \right] = (\alpha_{rJ}l_{rJ})^{\frac{\sigma_c-1}{\sigma_c}} \left[1 + \frac{\Theta_{cJ}}{(1-\Theta_J)\theta_{rJ}} \right]. \quad (18)$$

Plugging this into the first order condition on routine labor, (4), and dividing with the FOC on manual labor, (3), and re-arranging we get:

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

Multiplying the above with w_{rJ}/w_{mJ} and substituting in θ_{rJ}/θ_{mJ} we obtain (10):

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

Next we express the labor aggregate as:

$$LA = \sum_{o=m,a} (\alpha_{oJ}l_{oJ})^{\frac{\rho-1}{\rho}} + RA^{\frac{\sigma_c}{\sigma_c-1} \frac{\rho-1}{\rho}} = (\alpha_{mJ}l_{mJ})^{\frac{\rho-1}{\rho}} \frac{1}{\theta_{mJ}} \left(1 + \frac{\Theta_{cJ}}{1-\Theta_J} \right), \quad (19)$$

using (18) and substituting in $\left(\frac{\alpha_{rJ}l_{rJ}}{\alpha_{mJ}l_{mJ}} \right)^{\frac{\rho-1}{\rho}} = \frac{\theta_{rJ}}{\theta_{mJ}} \left[1 + \frac{\Theta_{cJ}}{(1-\theta_J)\theta_{rJ}} \right]^{\frac{\sigma_c-\rho}{(\sigma_c-1)\rho}}$ and $\left(\frac{\alpha_{aJ}l_{aJ}}{\alpha_{mJ}l_{mJ}} \right)^{\frac{\rho-1}{\rho}} = \frac{\theta_{aJ}}{\theta_{mJ}}$ (obtained similarly to (17)), and that $\sum_o \theta_{oJ} = 1$. Plugging the expression for LA into the FOC for manual labor, (3), and dividing by the FOC on traditional capital, (6), and re-arranging we get:

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

Multiplying with R_k/w_{mJ} and re-arranging we get equation (11):

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

Finally we express sectoral output as a function of observables. Using the expression on LA (19) and substituting that $\left(\frac{\alpha_{kJ}k_J}{\alpha_{mJ}l_{mJ}} \right)^{\frac{\sigma-1}{\sigma}} = \frac{\Theta_J - \Theta_{cJ}}{(1 - \Theta_J)\theta_{mJ}} \left[\frac{1}{\theta_{mJ}} \left(1 + \frac{\Theta_{cJ}}{1 - \Theta_J} \right) \right]^{\frac{\sigma-\rho}{(\rho-1)\sigma}}$ (obtained similarly to (17)) we can express sectoral output as:

$$Y_J^{\frac{\sigma-1}{\sigma}} = LA^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma}} + (\alpha_{kJ}k_J)^{\frac{\sigma-1}{\sigma}} = (\alpha_{kJ}k_J)^{\frac{\sigma-1}{\sigma}} \frac{1}{\Theta_J - \Theta_{cJ}}.$$

Raising the above to the power of $\sigma/(\sigma - 1)$ we get the expression in the main text.

A.4 Decomposing labor-augmenting technological change




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.

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).

A.5 Robustness checks and extensions

A.5.1 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 A2 shows in the top rows the average annual growth rates of the factor augmenting technologies in each sector in our baseline. The subsequent segments show these growth rates for the various alternative calibrations (different σ , ρ , σ_c and heterogeneous σ^J across sectors). All display the same key features that we highlighted in the discussion of Table 1.


Similarly Figures A2, A3 and A4 show in the top row the baseline, and in subsequent rows the results from considering two alternative values for σ , ρ and σ_c , respectively. These figures demonstrate that all the results from the baseline are replicated for all alternative parametrizations. Figure A5 shows the robustness of the model to allowing for different σ^J across sectors. In this figure the column on the left shows the baseline results, and the one on the right the results with heterogeneity across sectors.

Additionally, Table A3 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 pre-


dictions 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 A2: Average annual growth rate of α s over 1960–2017 under alternative parameters


		occupations			capital	
		manual	routine	abstract	non-ICT	ICT
Baseline:						
$\sigma = 0.84,$	L	1.0025	1.0292	0.9933	1.0085	1.0200
$\rho = 0.6,$	G	1.0058	1.0559	1.0100	0.9839	1.0439
$\sigma_c = 2$	H	1.0067	1.0132	0.9763	1.0178	0.9803
Alternative σ:						
	L	1.0017	1.0283	0.9925	1.0097	1.0192
$\sigma = 0.75$	G	0.9979	1.0476	1.0021	0.9945	1.0357
	H	1.0121	1.0186	0.9815	1.0115	0.9856
	L	1.0013	1.0279	0.9921	1.0103	1.0187
$\sigma = 0.65$	G	0.9939	1.0434	0.9981	0.9999	1.0316
	H	1.0149	1.0214	0.9842	1.0083	0.9883
Alternative ρ:						
	L	1.0058	1.0262	0.9979	1.0085	1.0170
$\rho = 0.5$	G	1.0123	1.0527	1.0167	0.9839	1.0407
	H	1.0039	1.0079	0.9797	1.0178	0.9752
	L	0.9970	1.0342	0.9856	1.0085	1.0250
$\rho = 0.7$	G	0.9949	1.0614	0.9990	0.9839	1.0493
	H	1.0113	1.0221	0.9706	1.0178	0.9890
Alternative σ_c:						
	L	1.0025	1.0281	0.9933	1.0085	1.0640
$\sigma_c = 1.5$	G	1.0058	1.0550	1.0100	0.9839	1.0900
	H	1.0067	1.0095	0.9763	1.0178	0.9995
	L	1.0025	1.0295	0.9933	1.0085	1.0057
$\sigma_c = 2.5$	G	1.0058	1.0562	1.0100	0.9839	1.0290
	H	1.0067	1.0144	0.9763	1.0178	0.9740
Sector specific σ:						
$\sigma^L = 0.75$	L	1.0017	1.0283	0.9925	1.0097	1.0192
$\sigma^G = 0.9$	G	1.0191	1.0699	1.0234	0.9666	1.0578
$\sigma^H = 0.75$	H	1.0121	1.0186	0.9815	1.0115	0.9856




(a) Baseline




(b) $\sigma = 0.65$




(c) $\sigma = 0.75$




(d) $\rho = 0.5$



(e) $\rho = 0.7$




(f) $\sigma_c = 1.5$




(g) $\sigma_c = 2.5$

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. The sets of bars are exactly the same as in Figure 6, which is also reproduced in graph (a) above.




(a) Baseline




(b) $\sigma = 0.65$




(c) $\sigma = 0.75$




(d) $\rho = 0.5$



(e) $\rho = 0.7$




(f) $\sigma_c = 1.5$




(g) $\sigma_c = 2.5$

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 different elasticities. The sets of bars are exactly the same as in Figure 7, which is also reproduced in graph (a) above.




(a) Baseline




(b) $\sigma = 0.65$




(c) $\sigma = 0.75$




(d) $\rho = 0.5$



(e) $\rho = 0.7$




(f) $\sigma_c = 1.5$




(g) $\sigma_c = 2.5$

Figure A4: 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 different elasticities. The sets of bars are exactly the same as in Figure 8, which is also reproduced in graph (a) above.




(a) Role of inputs, baseline




(b) σ^J




(c) Role of technologies, baseline




(d) σ^J




(e) Role of components, baseline



(f) σ^J



(g) Aggregate labor prod., baseline



(h) Aggregate labor prod., σ^J

Figure A5: Baseline vs sector specific σ^J

This figure shows the differences between the magnitude of the various channels when considering sector-specific σ s 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.

Table A3: Distance measure of the different predictions

ρ	neutral	full factor	sector	occupation
0.5	0.674	0.024	0.150	0.455
0.6	0.702	0.033	0.227	0.408
0.7	0.751	0.049	0.359	0.339
0.8	0.832	0.076	0.587	0.247
0.9	0.942	0.111	0.912	0.167

A.5.2 Allowing for efficiency units of labor

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

$$\log w_{ioJt} = \delta_{oJt} + \beta' X_{it} + \varepsilon_{ioJt}, \quad (20)$$

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 ($\bar{w}_{oJ,t}$), we need to assign the cell-year average of the exponent of the residuals from (20) 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, $\bar{e}_{oJ,t}^1$, as the average of $\hat{e}_{ioJt}^1 = \exp(\beta' X_{it})$ within the sector-occupation-year cell. In this case the implied sector-occupation-year unit wages are given as $\hat{w}_{oJ,t}^1 = \bar{w}_{oJ,t} / \bar{e}_{oJ,t}^1$. Alternatively we construct sector-occupation cell efficiency wages per hour, $\hat{w}_{oJ,t}^2 = \exp(\delta_{oJt})$. The implied average sector-occupation-year efficiency units per hour worked are then $\bar{e}_{oJ,t}^2 \equiv \bar{w}_{oJ,t} / \hat{w}_{oJ,t}^2$.

We use the equivalent of (16) to get sector-occupation wages per efficiency unit

$(\tilde{w}_{oJ,t})$:

$$\tilde{w}_{oJ,t}^M l_{oJ,t} \bar{e}_{oJ,t}^M = Y_t^{nom} \cdot VA_{J,t} (1 - \Theta_{J,t}) \theta_{oJ,t}, \quad (21)$$

where $\bar{e}_{oJ,t}^M$ 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:

$$\frac{\tilde{w}_{oJ,t}^M}{\tilde{w}_{rJ,t}^M} = \frac{\frac{\theta_{oJ,t}}{l_{oJ,t} \bar{e}_{oJ,t}^M}}{\frac{\theta_{rJ,t}}{l_{rJ,t} \bar{e}_{rJ,t}^M}} = \frac{\frac{\bar{w}_{oJ,t}}{\bar{e}_{oJ,t}^M}}{\frac{\bar{w}_{rJ,t}}{\bar{e}_{rJ,t}^M}} = \frac{\hat{w}_{oJ,t}^M}{\hat{w}_{rJ,t}^M},$$

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 A4: Sector-occupation efficiency units of labor 1960–2017

J	o	1960	1970	1980	1990	2000	2010	2017
L	m	1.685	1.642	1.498	1.548	1.598	1.622	1.626
L	r	1.848	1.794	1.690	1.743	1.800	1.834	1.825
L	a	2.034	1.991	1.890	1.930	2.005	2.044	2.032
G	m	1.975	1.903	1.809	1.795	1.831	1.860	1.874
G	r	1.878	1.840	1.756	1.820	1.874	1.939	1.930
G	a	2.148	2.185	2.156	2.199	2.279	2.356	2.329
H	m	1.844	1.801	1.732	1.812	1.853	1.871	1.877
H	r	1.748	1.700	1.655	1.721	1.488	1.424	1.441
H	a	2.125	2.121	2.069	2.139	2.200	2.243	2.248

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

J	o	1960	1970	1980	1990	2000	2010	2017
L	m	1.975	1.924	1.707	1.715	1.775	1.762	1.751
L	r	2.060	2.003	1.922	1.985	2.038	2.093	2.082
L	a	2.450	2.360	2.249	2.270	2.332	2.380	2.383
G	m	2.130	2.099	2.056	2.023	2.103	2.102	2.087
G	r	2.098	2.044	1.982	2.032	2.091	2.192	2.153
G	a	2.447	2.462	2.420	2.482	2.588	2.685	2.670
H	m	2.009	1.984	1.928	2.017	2.075	2.105	2.113
H	r	1.858	1.855	1.847	1.929	1.682	1.612	1.658
H	a	2.416	2.409	2.324	2.414	2.501	2.556	2.580

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

Table A4 shows efficiency units by sector-occupation over time for the two meth-

ods. 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 $n_{o,J,t} \equiv \bar{e}_{o,J,t} l_{o,J,t}$ in each period, instead of just hours worked ($l_{o,J,t}$). This implies that we need to use wages per efficiency unit of labor in equations (8) to (13), but the measurement of all other variables remains the same as in the baseline model. Figure A6 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 (21), and the solid lines show our baseline (of wages per hour worked from (16)). Note that all alternative lines qualitatively show the same patterns, some are also quantitatively very close. The only larger difference is for manual and abstract wages relative to routine in sector H between 2000 and 2017, which can be traced back to a fall of routine workers' efficiency units in this sector for this period as shown in Table A4. Given that the relative wage path




Figure A6: 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 (16), solid lines), (ii) from fitted efficiency units ((21) based on method 1, dotted lines), (iii) from fitted efficiency wages ((21) based on method 2, dashed lines).

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, $\tilde{w}_{o,J,t}^M$ 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 A5 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 (8) to (12) 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 α_{cJ} or α_{kJ} .

Table A5: Average annual growth rate of α s over 1960–2017 accounting for efficiency units of labor


	occupations				occupations		
	manual	routine	abstract		manual	routine	abstract
<i>L</i>	1.0031	1.0294	0.9933	<i>L</i>	1.0046	1.0290	0.9938
<i>G</i>	1.0067	1.0554	1.0086	<i>G</i>	1.0061	1.0554	1.0085
<i>H</i>	1.0064	1.0166	0.9753	<i>H</i>	1.0058	1.0152	0.9752

(a) based on fitted efficiency units, \bar{e}^1 (b) based on residual efficiency units, \bar{e}^2


Notes: The change in the capital inputs' technologies (the α_{cJS} and α_{kJJS}) is exactly the same as in Table 1 and not shown here.

Comparing Table A5 to Table 1 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 α_{oJ} but also quantitatively. This is perhaps not that surprising given that we established in Figure A6 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 A7 shows this for the model variant based on fitted efficiency units, \bar{e}^1 .




(a) Role of inputs, baseline




(b) Role of inputs, efficiency units




(c) Role of technologies, baseline




(d) Role of technologies, efficiency units




(e) Role of components, baseline



(f) Role of components, efficiency units



(g) Aggregate labor prod., baseline



(h) Aggregate labor prod., efficiency units

Figure A7: 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.