

ECONOMICS

Working Paper Series

No. 0 -2022

Occupational polarisation and endogenous task-biased technical change

Wenchao Jin^{*}

University of Sussex and Institute for Fiscal Studies

December 2, 2022

Abstract

Since the 90s many developed countries have experienced job polarisa-

1 Introduction

In many developed countries since the 90s, employment has shifted substantially away from middle-paying occupations towards both the top and the bottom (Goos et al., 2014). This phenomenon - employment polarisation - has been an important factor in rising income inequality. The dominant explanation of this in the literature is Routine-Biased Technological Change (RBTC)¹. Most of the polarisation literature interpret RBTC as a consequence of increasing availability or productivity of automation equipments, or their declining costs. In other words, it's an exogenous demand shift hitting the middle-paying routine-intensive jobs. This paper offers a complementary explanation: while incorporating exogenous technical change, the emphasis here is on the increasing supply of skilled labour and the consequent switch to routine-biased technology.

From a policy perspective, supply-side policies such as increasing education are important policy levers for addressing income inequality in the long run. Given the prevalence and the scale of employment polarisation and its adverse impact on inequality and social mobility, it is surprising that few papers have examined the role of supply-side shifts in the polarisation context². In theory, demand-side factors and supply-side factors could affect each other in endogenous ways: technological change may respond to supply-side shifts, while education choices may depend on expected demand shift. The UK provides a uniquely-suitable context to investigate this problem because its increasing supply of graduates was largely driven by policy.³

I build on the RBTC literature by allowing the adoption of technology to respond to skill supply shifts. The idea that firms' choice of production technology depends on the supply of skills is supported by a growing literature (Beaudry et al., 2010; Lewis, 2011; Akerman et al., 2015). Compared to standard theories of RBTC, incorporating endogenous adoption of technology gives different implications for the effects of supply shifts on wages. As we will see, the UK data

¹The general idea is that new technologies (embodied by computer software and automation equipment) have displaced workers in carrying out routine tasks, which are important in middle-paying occupations. Acemoglu and Autor (2011) and Autor (2022) provide a good summary.

²Some papers (Hardy et al., 2018; Salvatori, 2018) have argued for a major role of education increase in the growth of cognitive or high-paying jobs in Europe/UK, by decomposing over-time changes into between and within components. This paper uses an equilibrium model to provide a clearer conceptual distinction between supply-side shifts and demand-side shifts. A recent paper Patel (2022) also quantifies the contributions of supply shifts versus demand shifts, which will be discussed in more detail later.

³For about two decades since the early 90s, undergraduate student numbers in individual universities were capped by the government, and they were allowed to increase year on year.

⁴Typically, these studies use exogenous geographical variation in the supply of educated workers to prove the causality from skill supply to technology adoption.

support the latter. My model can explain not just employment polarisation, but three facts about occupations in the UK at the same time.

First, the pattern of employment polarisation in the UK is essentially a shift

larisation (Green and Sand, 2015). By contrast, the chain of events emphasised here starts with a policy-driven positive supply shift, which causes task-biased technical change, and therefore leads to the three aforementioned facts about oc-

technology for two reasons: the UK data pattern does not support the hypothesis of exogenous RBTC, and there is growing evidence elsewhere that skill supply affects the adoption of technology.

There is a rapidly growing literature on the endogenous adoption of specific technologies and its effects on employment or wages. They usually focus on a tangible technology, such as personal computers (Beaudry et al. (2010), Borghans and ter Weel (2008)), broadband internet (Akerman et al., 2015), software (Contractor and Taska, 2022), automation (Aghion et al., 2020), industrial robots (Graetz and Michaels (2018), Humlum (2019)), or computer numerical control (Boustan et al., 2022). They often find that the adoption of technology was indeed affected by the local supply of skills or local wages. Their research questions centre around the causal effects of adopting that technology on employment of different skill groups, wages, productivity and so on.⁹ By contrast, this paper aims to explain overall patterns in all parts of the economy in a unified framework. So I choose not to focus on one specific technology. In my model, technology boils down to the production function that combines tasks into output.¹⁰ In each industry, there will be an 'Old' technology and a 'New' technology. I believe technological changes take different forms in different industries. It could be robots in manufacturing, automated software in financial services, and some sort of organisational restructuring in another services firm. And all those kinds of technical changes may be complementary to each other (Bresnahan et al. (2002), Caroli and Van Reenen (2001)). Empirically, we will use a wide range of tangible and intangible measures to estimate the share of the 'New' technology at the industry-year level.

The paper is also closely related to Blundell et al. (2022). It noted that the rapid growth of graduate numbers in the UK had no noticeable impact on graduate wages, and explained it by an endogenous adoption of skill-biased technical change. This paper uses the same intuition but in a different context, because the aim here is to explain the facts about occupations and to allow counterfactual analysis. In addition, Carneiro et al. (2018) and Dustmann and Glitz (2015) also found that production technology responds to changes in the local supply of educated/uneducated workers. Like Blundell et al. (2022), they differentiate labor by education and have nothing to say about occupations.

⁹Most of these papers did not model general-equilibrium effects. To my knowledge, Humlum (2019) was the first to estimate a general equilibrium model of technology adoption. His model is rich in how manufacturing firms choose whether to adopt robots and parsimonious for the rest of the economy. Specifically, the production function outside manufacturing is Cobb-Douglas and contains no task-biased technical change.

¹⁰We do not model capital explicitly in this paper. We can think of the choice of capital equipment as a choice of the production function that combines occupational labor into output. For example, adopting robots in the production process means you would need more technicians and fewer production workers to produce one unit of output.

We fit the model to the UK data over 1997-2015 at the level of 9 occupations and 7 industries.¹¹ It can fit the observed trends pretty well. The good fit is not mechanically guaranteed by the model design, because most of the key parameters (such as preferences) do not vary over time. The estimates in most industries suggest that technological changes in the UK over 1997-2015 were bi-

2.1 Fact 1: employment polarisation

Employment polarisation refers to a 'hollowing out' along the occupation spectrum. This phenomenon has been documented extensively in the literature for the US (Acemoglu and Autor (2011), Autor and Dorn (2013), Hershbein and Kahn (2018)) as well as many other developed countries (Goos et al. (2014), Breemersch et al. (2017), Michaels et al. (2014)). It's been documented since the 1980s for the UK (Goos and Manning, 2007) and Germany (Kampelmann and Rycx, 2011) and even earlier for the US (Barany and Siegel (2018)). The phenomenon is robust to different ways of classifying and ranking occupations for both the US and the UK. When my model is brought to the UK data, occupation will be at the level of SOC2000 major occupation group¹³. So in this section I present occupational facts at this level, too. At this level of nine occupations, the three middle-paying occupations are normally considered 'routine': 'Administrative and Secretarial Occupations', 'Skilled Trades Occupations', and 'Process, Plant and Machine Operatives'. The three high-paying ones will be referred to as 'abstract', and the low-paying ones as 'manual'.


Figure 1 shows that each of the three routine occupations saw a very substantial decline in employment share. Over 1997-2015 (the period for which my model will be estimated), the total employment share of the 3 routine occupations fell from 39.1% to 28.5%: a decline of 10.6%. Meanwhile, each of the three abstract occupations grew substantially. In particular, professional occupations grew from 9.9% of aggregate employment to 15%. Together, the abstract employment share grew from 39.1% to 49.4% over the sample period: an increase of 10.3%. Among the manual occupations, there is some decline in elementary occupations¹⁴, which is more than compensated by the increase in personal services (such as care assistants). Overall, the pattern of employment polarisation in the UK is more of a shift of employment from the middle to the top, with very little change at the bottom.

At a similar level of aggregation, Figure 12 shows a V shape in employment growth across ISCO occupation groups in a number of European countries over 2002-14. This echoes the findings in Goos et al. (2014), which looked at 16 European countries and documented pervasive occupational polarisation over 1992-2010. On the other hand, some more recent studies looking at employment changes in European countries found no polarisation pattern but 'occupation upgrading' - meaning fastest growth in the 'best' jobs and weakest growth in the 'worst' jobs.

¹³There are 9 occupations in total : 1, managers and senior officials, 2 professional, 3 associate professional and technical, 4 administrative and secretarial, 5 skilled trades, 6 personal services, 7 customer services, 8 process, plant and machine operatives, and 9 elementary.

¹⁴which include labourers in agriculture, cleaners, waiters, kitchen assistants, labourers in construction, porters, postal workers and so on.

Figure 1: Employment shares by occupation



Note: the 9 occupations are major occupation groups under SOC2000. See section 4 for how we adjusted for discontinuities in SOC over 2000-01 and 2010-11.

For example, Ferrandez-Macas and Hurley (2017) looked at 23 European countries over 1995-2007 and found polarisation in a handful of countries but the most common pattern is occupational upgrading. Oesch and Piccitto (2019) looked at UK, France, Germany and Spain over 1992-2015 and found job growth was by far the weakest in the 'lowest-quality' jobs using a range of measures of job quality.¹⁵ Murphy and Oesch (2018) looked at Ireland and Switzerland over 1970-2010 and found 'occupational upgrading', and the patterns were consistent with changes coming from the supply side associated with female education and immigration. It's beyond the scope of this paper to investigate why those studies reach different conclusions. Notably, they all point to strong growth in high-paying occupations. We see in both Figure 1 and Figure 12 that the professional occupation stands out as having the strongest growth. This is an occupation in which university graduates are likely to have comparative advantage. In the framework proposed here, an increase in the supply of graduates will cause firms to adopt a technology that's more intensive in professional tasks, and therefore the professional employment share will increase. My model does not have a definitive prediction as to whether low-paying occupations should grow or decline relative to the middle. Both 'occupational polarisation' and 'occupational upgrading' could be the consequence of an increase in skills supply. The former follows if the new technology is biased against middle-skilled tasks and in favour of high-skilled tasks; while the latter follows if the new technology is biased in favour of high-skilled tasks and against low-skilled tasks.

2.2 Fact 2: no wage polarisation

Meanwhile, apart from the US, there is no such V shape in occupational wage growth in other developed countries that also saw employment polarisation.

Figure 2 ranks the 9 occupations from the lowest paid to the highest paid, and plots the occupational wage growth in red markers. The plotted wage changes are net of compositional shifts in education, age and gender.¹⁶ The three low-skilled occupations have slower wage growth than 5 of the other 6. Skilled trades and operatives have fairly strong wage growth, while admin had the slowest wage growth. The maximum difference between occupations in log wage changes over 1997-2015 is just under 0.08. This is small relative to the observed changes in employment shares.¹⁷

¹⁵The only exception, they found, is for the earnings-based indicator in the UK, which suggests a polarising pattern.

¹⁶In each year, I have regressed log wages on those demographics and occupation dummies. The coefficients on occupational dummies are interpreted as 'composition-adjusted' occupational wages.

¹⁷To give a sense of magnitude, if tasks are neither complements nor substitutes, the response


Figure 2: Changes in log occupational wages

One might have expected such a big supply-side shift to reduce the relative wage of graduates. In reality, that has not happened. Blundell et al. (2022) documents this and explains it in a model of endogenous technology adoption.

One might also expect the huge increase in graduate numbers to lead to 'occupational downgrading', that is, an adverse shift in occupational destinations of graduates over time. However, there has not been much occupational downgrading among graduates in the UK. The right subgraph in Figure 3 shows that among graduate workers, the proportion in abstract occupations has been stable over time, at around 80%. There seems to be a little fall after 2010, to around 75% by 2015, which is still very far above the level among high-school workers.

To give a sense of magnitude, I calculate how much the share of abstract occupations needs to fall within education group if the aggregate abstract share had been constant while the education composition improved²². These counterfactual trends are plotted as dashed lines in Figure 3: the proportion in abstract occupations conditional on education would need to fall by about a quarter. Thus, the UK story is one where the increase of graduates was quickly absorbed through employment growth in abstract occupations. The model in section 3 will formalise this intuition: increasing education means more workers now have comparative

Figure 3: Proportion of graduates and their occupation destination



Note: graduates are people with NVQ level 4 qualifications or above. High-school workers refer to those with NVQ level 2 or 3 qualifications. First degrees are NVQ level 4. A-levels and post-16 further education qualifications are NVQ level 3. O-levels and GCSEs (grade C+) are NVQ level 2. 'Abstract' refers to the first three occupations in SOC2000: managerial, professional and technicians.

with complete tertiary education was already 24% in the US by 1990, when the proportion in European countries was all below 15%. This supports the view that the US has been the leader of skill-intensive technologies in general, with other developed countries closely behind. This means when their workforce's education level catches up, the latter group (including the UK) are in a position to adopt newer technologies, and this choice would depend on prices and wages. Consistent with this view, Blundell et al. (2022) shows that in 11 OECD countries which experienced substantial increase in tertiary education, there was no significant decline in graduates' relative wages in 9 of them, like the UK. Finally, it has also been documented in Green and Henseke (2021) that in 24 European countries, the share of graduates in non-graduate occupations has increased 'only modestly' from 19 to 21 percent over 2005-13²³. All these similarities suggest that a model of endogenous adoption of technology (like the one proposed in section 3) might be more suitable for these non-US developed countries, whereas the US might need a model of endogenous innovations.

²³See Figure 3 and the associated description in Green and Henseke (2021). They defined graduate occupations as the top three ISCO-08 major groups, which is very similar to the definition of abstract occupations in this paper. They looked at 10 Central and Eastern European countries and 14 old EU countries. Only 4 countries saw an increase in the share of graduates in non-graduate jobs by more than 5 percentage points, and they are all in Central and Eastern Europe. And in every one of the 24 countries, the share of graduates in the workforce increased over the period. The UK is around the middle in the distribution of the growth rate of the graduate share among the 24 countries over that decade.

3 Model of endogenous adoption of task-biased technology

This section develops an equilibrium model of occupational labour (called 'tasks' for brevity). The model is static because we are interested in long-run comparative statics. On the demand side, there are multiple industries and within each industry firms choose between two technologies that differ in task intensities. On the supply side, workers have two dimensions of observable skills and an unobservable general ability. They sort into occupations based on wages and preferences.

In this paper I will use 'occupations' and 'tasks' inter-changeably. In reality, the task content within occupations may change continuously as overall demand for tasks change. This is an interesting challenge for future research.²⁴ In this paper, 'tasks' should be interpreted as the output of specific occupations. For example, professional tasks are simply the output of workers in professional occupations, whether the actual activity carried out is writing reports or analysing data is not studied here.

Each industry produces one good. Denote the goods as $g \in \{1, 2, \dots, G\}$. The production of each good is a CES function of tasks $j \in \{1, 2, \dots, J\}$, given the technology choice.

To produce any given good g , there are two potential technologies, denoted by $T \in \{O, N\}$. Each firm can choose freely between the 'Old' technology and the 'New' technology. Firms are otherwise identical within the industry. The difference between two technologies is that they have different task intensities τ_{gj}^T . They also have their own TFP term A_{gt}^T , which is neutral with regard to tasks.

$$Y_{gt}^T = A_{gt}^T \left[\sum_j \tau_{gj}^T (y_{gjt}^T)^\frac{1}{\sigma} \right]^\sigma ; T \in \{O, N\} \quad (1)$$

Y_{gt}^T is the output produced in industry g at time t under technology T . y_{gjt}^T is the amount of task j employed in industry g , using technology

capital equipment as a choice of the production function that combines occupational labor into output. For example, adopting robots in the production process means you would need more technicians and fewer production workers to produce one unit of output. If the New technology uses robots, and the price of robots falls or the productivity of robots increases, then this would be reflected as an increase in A_{gt}^N . Both A_{gt}^O and A_{gt}^N are assumed to be exogenous.²⁵ If the New technology requires different amounts of capital, then by assuming that firms can switch to the new technology freely, I have also assumed a perfectly elastic supply of capital.²⁶

Each technology is assumed to have constant returns to scale. We normalise $\sum_j \pi_{gj}^T = 1$; π_{gj}^T ; T .

Consumers have CES preferences over G goods, with σ being the elasticity of substitution. Q_{gt} is output in industry g at time t . B_{gt} captures time-varying demand for good g . B_{gt} is assumed to be exogenous here.²⁷

$$U_t = \left[\sum_g B_{gt} Q_{gt}^{\frac{1}{\sigma}} \right]^{-\sigma} \quad (2)$$

$$Q_{gt} = Y_{gt}^O + Y_{gt}^N \quad (3)$$

A good produced by the Old technology is a perfect substitute for the same good produced by the New technology.

Because technology O and N differ in task intensities, we can think of a shift between technology O and N as task-biased technological change. This could be caused by changes in TFP in either technology option, industry demand shifts, or changes on the supply side. Ex ante, the model does not prescribe the New technology as routine-biased. It is left for the data to tell us how task intensities differ between the Old and New technologies.

The primary difference between my model and the RBTC literature is the presence of two technologies to choose from. If there's only one technology, then employment shares can only change due to changing task prices or changing parameters in the production function. The latter could be modelled as exogenously

²⁵This assumption rules out the possibility that the price of new capital equipment might respond to demand or supply shifts in the UK. Such an assumption would be questionable for a major innovator like the US.

²⁶If capital is not inelastically supplied, the impact of an education increase on relative wages would be different. ? developed a model with endogenous adoption of technology and where the input factors are skilled labor, unskilled labor and capital. They showed that holding aggregate capital constant, an increase in the skilled share of the workforce will increase the skilled to unskilled wage ratio by causing capital scarcity.

²⁷For future research, it would be interesting to allow income growth to differentially affect the demand for goods and services.

evolving share parameters in a CES production function, such as in Johnson and Keane (2013). The downsides are: 1) there are a lot more unobserved parameters (one will need α_{ijt} instead of α_{ij}), and 2) there is one less channel to absorb supply-side shocks, so the result of increasing skills supply will tend to be lower prices of high-skilled tasks. The reality is that the big increase in graduates did not reduce their relative wages, or the relative wage of abstract occupations. In my model, this happens through the endogenous shift towards the New technology, which is more intensive in the tasks that graduates have comparative advantage in. By contrast, in a model with exogenous technology, the technology's parameters would need to shift in favour of the tasks that graduates have comparative advantage in, and at a speed that happens to leave the task prices and the mapping from education to occupation relatively unchanged. In section 5, I will formally test the hypothesis of exogenous task-biased technical change and reject it in favour of my model.²⁸ It is worth noting that my model allows for exogenous technical change as well: the TFP trends A_{gt}^T are exogenous, and a sufficiently large increase in the New technology's TFP will induce all firms to switch to it.

The CES formulation is common to the task literature, and many papers make the more restricted assumption of Cobb-Douglas production.²⁹ One exception is Johnson and Keane (2013). Johnson and Keane (2013) differentiates labour by occupation, education, gender and age. Their production function is multi-level nested CES.³⁰ Their formulation is more detailed than my model. To fit the US data over 29 years of data, they found that it's necessary to allow the share parameters to follow 3rd or 4th order polynomials. By contrast, there is no time-variation in the share parameters in my model. Thus, ex ante, it's more challenging for my model to fit occupational trends.

That is the demand side. Now let's specify the supply side.

Suppose each person is endowed with two dimensions of observable skills and an unobserved general ability μ_i . The joint distribution of skills is assumed to be exogenous. Later on we will consider counterfactual policies that shift the skills distribution, through education or immigration. In reality, RBTC may induce workers to undertake more education or training in order to become more produc-

²⁸It's a rejection of the hypothesis that all technical change is exogenous. It does not reject the hypothesis that there is some exogenous shock to technology.

²⁹For example, Autor (2013) defines output as CES over a continuum of tasks; Acemoglu and Autor (2011) models output as Cobb-Douglas over a continuum of tasks; Autor and Dorn (2013) models goods output as Cobb-Douglas over routine task and abstract task, and services output is simply manual labour times a scalar; Traiberman (2019) models output in each industry as a Cobb-Douglas function of capital, human capital in each occupation and intermediate inputs produced in other industries.

³⁰The bottom three levels are education, gender and age; at the top level, aggregate output is CES between unskilled task and skilled task; unskilled task is 2-level CES of 8 occupations, and skilled task is 2-level CES of capital and 2 occupations.

tive in abstract tasks (Battisti et al., 2017). Such an endogenous response on the skills distribution is left for future investigation.

In the workplace, only the individuals' skills matter for productivity, not their education per se. Each occupation produces one task. Occupation and task are both denoted by subscript j . The amount of task that worker i in occupation j produces is

$$y(i;j) = k_j e^{a_j a_i + s_j s_i + \alpha_i} \quad (4)$$

This formula follows from Autor and Handel (2013), where I specify observable skills to have 2 dimensions. a_i is analytical ability and s_i is social skill. α_i is worker's general ability which is unobserved. α_i can be correlated with observed skills freely. The coefficients a_j ; s_j are occupation-specific productivities of analytical and social skills. k_j is a j -specific scalar. The key assumption here is that comparative advantage is captured by 2 dimensions of skills, s_i ; and conditional on them, there is no omitted factor that makes a person more productive in one task rather than another.

The labour market is competitive. We assume workers do not directly care about the technology chosen by their employer or which industry they are in. Since a worker's task output is the same wherever they work, the task price must equalise

where p denotes the price vector of all tasks. Comparative advantage plays a role in the sorting into occupation: a worker with higher a_i is more likely to go to an occupation with higher a_j . A smaller σ means the preferences are less varied and so wages are more influential in occupation choices. Note that the unobserved heterogeneity term ϵ_i does not enter into occupational choice. Thus $\pi_k(i; p) = \pi_k(a_i; s_i; p)$.

Given task prices, the supply of task j in the economy is

$$LS_j(p) = \sum_{i \in Z} \pi_j(a_i; s_i; p) y(i; j) \quad (8)$$

$$= \int \pi_j(a; s; p) y(a; s; j) f(a; s) da ds \quad (9)$$

where $f(a; s)$ is the joint density function, and $y(a; s; j)$ is the expected output in task j conditional on observing $a; s$. The derivation of (9) is in Appendix A.2.

Thus, the only relevant unknowns on the supply side are π_j , $y(a; s; j)$ and $f(a; s)$. As long as we get $y(a; s; j)$

across occupation and industry $j(g)$. In the context of occupation-industry, a standard specification of exogenous technical change would use a j - g -specific time polynomial. I will test such a hypothesis in section 5, and show that the data cast doubt on it.

3.1 Equilibrium characteristics and effect of a supply-side shift

I define the equilibrium as log task prices ($\log p_t = f(\log p_{1t}, \dots, p_{Jt})$) and technology shares ($\theta_t = f$

of industries where the unit costs are equal. For a majority of years in our sample period (1997-2015), it has 7 dimensions.

The transformed case is observationally equivalent to the original one:

$$(1 - w_{gt})r_{gj}^O + w_{gt}r_{gj}^N = (1 - \tilde{w}_{gt})\tilde{r}_{gj}^O + \tilde{w}_{gt}\tilde{r}_{gj}^N, \forall j, t$$

Therefore, we will anchor the time series w_{gt} by assuming $w_{g0} = 0; w_{gT} = 1$. This 'normalisation' is not totally innocuous because it assumes that w_{gt} cannot go above w_{gT} or below w_{g0} . This seems true in the UK data, and it allows easy interpretation: we are effectively calling the production function at time 0 the Old technology and the one at time T the New technology.

Empirically, we will estimate w_{gt} from technology proxies. Suppose we have a proxy for new technology called z_t , such that $z_N > z_O$. The assumption here is that all firms with the New tech have the same level of z_t , which is higher than the level among old-tech adopters. There is no time variation within z_N or z_O . Thus, the observed change in z_{gt} at the industry level reveals the shift towards the New technology within this industry.

$$z_{gt} = (1 - w_{gt})z_O + w_{gt}z_N \quad (14)$$

where w_{gt} is the scale of new technology adopters relative to the entire industry.

In practice, we will use several measures of w_{gt} . We observe w_{gt} over time and at the industry level. If z_{gt} comes from employee survey, w_{gt} is the employment share of firms using the new technology in the industry-year. As we anchor w_{gt} to 0 at one point and 1 at another point, we would be setting $z_O = z_{g0}; z_N = z_{gT}$. Thus, we can impute w_{gt} as $\frac{z_{gt} - z_{g0}}{z_{gT} - z_{g0}}$. Thus, w_{gt} is just-identified by one proxy up to an affine transformation. If we have several measures of w_{gt} , we can allow errors in equation (14). In section 4.3, we will assume a latent factor model to impute w_{gt} .

3.3 Identification of model parameters

in employment will be attributed to unobservable preference shifts. Empirically, I search for β_j to match the observed employment shares in 2006 (the mid-point of my sample period).

The smaller β_j is, the more elastic task supply will be with regard to task prices. The identification of β_j relies on movements along the task supply curve. Had there been no changes to the skills distribution, small movements in task prices together with large movements in employment would imply that β_j is small.

The joint skill distribution comes from the numeracy score and the literacy score in the British Cohort Studies (BCS), measured at age 34. They are summarised to 7 points of support in each dimension²⁴. The skills distribution in the BCS data might be quite different from the aggregate skill distribution in the UK because the BCS only contains the 1970 birth cohort. The aggregate skill distribution might be changing over time due to increasing education as well as immigration. I assume the joint distribution of analytical and social skills is fixed

by industry and occupation will give usr_{gj}^O as the constant and r_{gj}^N r_{gj}^O as the

We regress each time series (in log terms) on a 5th order polynomial of time plus a dummy for $t < 2001$ and a dummy for $t \geq 2011$. In other words, we allow the occupation classification change to affect the level of the variable and nothing else. We deduct the estimated jump from the affected period. Figure 14 in the appendix plots the raw and adjusted p_{jt} for three example occupations. There are clearly jumps in some raw time series at 2001 and 2011, and the adjusted time series are smoother. We use the adjusted data in both descriptive graphs (Figure 1, Figure 2) and when estimating the model.

4.2 Skills distribution

We use numeracy and literacy skills in the British Cohort Study (BCS). The BCS is a longitudinal survey following around 17,000 people who were born in England in 1970. BCS contains many skill assessments at various ages, sometimes for a subset of the cohort. We are interested in skills measured after the completion of education, because education could have affected skills. We also prefer a larger sample. After age 16, there is only one wave (at age 34) when skills were assessed for the whole sample. Hence, in this paper we will use literacy and numeracy assessed at age 34. There are about 9500 observations with both skills measured at 34 in the BCS.

Figure 4 shows the distributions of two skills by education and gender. For each skill, the mean score clearly increases with education, while the distribution overlaps significantly between education groups. Both skills have raw scores with 20+ points but the lower range is very sparsely populated. I summarise them to 7 points of support in each dimension.

For obtaining wages conditional on skills and occupation, I pool all the waves together to increase sample size. I take age effects out of wages by simply regressing log wages on age dummies, and deducting the age effects from observed log wages. Then for each combination of skills and occupation, I use the mean wage excluding outliers as the data moment for $E[p_j y(i; j) | a; s; j]$.⁴¹ There are a number of empty $(a; s; j)$ cells (having no individual in the cell or no one reporting wages), and they all have rare combinations of skills where one skill is very high and the other skills is very low. In such cases,⁴² I use the observed average wage of that occupation.


4.3 Technology proxies

When setting out the model, I have not specified what the new technology is or means in practice. This is because I believe its practical manifestation would

⁴¹Within each (a, s, j) cell, I exclude the top and bottom 5% of wage observations in calculating mean wages.

⁴²Such pairs of $(a; s)$ constitute 0.9% of the BCS sample.

Figure 4: Distribution of literacy and numeracy scores in BCS



Note: from British Cohort Studies. The box edges correspond to the 25th percentile and the 75th percentile within the education and gender group. The line inside the box is the medium skill score. "HE" refers to higher education or above. "HS" refers to secondary school qualifications including A-levels, O-levels, GCSE C+ or equivalents. "None" refers to those without secondary school qualifications.

vary across industries and firms. It could be something tangible such as automation equipment in a manufacturing firm, or high-speed internet in a professional service firm; or it could be something intangible like a decentralized structure of management and decision-making. The different aspects of changes may be complementary to each other and skill-biased.(Bresnahan et al., 2002; Caroli and Van Reenen, 2001)

Guided by the literature (Michaels et al., 2014; Machin and Van Reenen, 1998), I consider measures of ICT capital and related tangible technology, as well as measures about intangibles, from two datasets: capital inputs in EU-KLEMS and the British Skills Survey (BSS). The former is available over 1997-2015. The BSS is available for 1986, 1992, 1997, 2001, 2006, 2012, 2017.

In EU-KLEMS, we observe various types of capital by year and across dozens of industries. At the industry-year level, I use the share of overall capital that is in each of the following four areas: Communication Technology, Information Technology, Software&database, and R&D. These variables about capital composition have increased over time. I have also verified that the graduate proportion is positively and significantly correlated with IT capital input at the industry-year level. Correlations with other capital inputs are mostly positive but insignificant, see table 1.

From the BSS, I obtain 5 proxies, which are responses to the following questions/statements: 'whether job involves use of computerised or automated equipment', 'my job requires that i keep learning new things', 'my job requires that i help my colleagues to learn new things', 'do you have a formal appraisal system at

Table 1: Capital input composition and the graduate proportion

	Comm. tech	Info. tech	Software&database	R&D
Graduate proportion	0.0047 (0.0054)	0.0280 (0.0143)	0.0309 (0.0244)	0.0809 (0.0141)
HS-Dropout proportion	-0.0045 (0.0050)	0.0139 (0.0133)	0.0189 (0.0228)	-0.0382 (0.0131)
Observations	133	133	133	133

Standard errors in parentheses

$p < 0.05$, $p < 0.01$, $p < 0.001$

Note: these regressions are at the level of industry-year, including industry dummies and year dummies. Each dependent variable is the share of overall capital in this type, with the industry-year. 'propBA' is the proportion of people with tertiary qualifications. 'propDO' is the proportion of people without GCSE grade C+ or equivalent.

your workplace', and 'In your workplace, what proportion of employees work with computerised or automated equipment?'.

Figure 5 shows the aggregate trend in these variables. They are mostly available for 5-6 waves in the BSS. They all increase strongly over time. Moreover, I summarise the data to the level of industry-region-year and regress each of the 5 proxies on the graduate proportion allowing for year dummies, industry dummies, region dummies. Table 2 shows that all these 5 proxies are very positively and significantly correlated with the local proportion of graduates, which is consistent with my model prediction.

Given a range of proxy measures z_{gt}^m , $1 \leq m \leq M$, we now impute w_{gt} in a latent variable model. Suppose each measure is a linear function of the latent variable w_{gt} plus some measurement error.


$$z_{gt}^m = \alpha_g^m + \beta_g^m w_{gt} + \epsilon_{gt}^m \quad (19)$$

The constant and the slope coefficient is specific to the measure m and the industry g . Because w_{gt} is unobserved, w_{gt} is only identified up to a linear transformation. I conduct a linear transformation of w_{gt} to equal 0 in 1997 and 1 in 2015. Figure 15 in Appendix B shows the resulting technology shares for all the industries.

5 Corroborative evidence

The key difference between my model and standard models in the RBTC literature is that the choice of technology in my model responds to supply shocks. This has different implications for how occupational wages respond to supply-side shocks.

Figure 5: Time trends in technology proxies in BSS



Note: the two learning measures take values between 0 to 3, 0 meaning 'strongly disagree' and 3 meaning 'strongly agree'. The other three are valued between 0 and 1.

Table 2: Proxies in BSS, correlation with graduate proportion

	own use PC	%PC at work	appraisal	learn new thing	help others
BA proportion	0.3276 (0.0707)	0.2733 (0.0443)	0.2000 (0.0702)	0.4234 (0.0929)	0.3081 (0.1244)
Observations	348	390	389	390	312

Standard errors in parentheses

p < 0:05, p < 0:01, p < 0:001

Note: all the outcomes are aggregated to the industry-year-region level. Each regression is at the industry-year-region level, including year dummies, industry dummies, region dummies. 'own use PC' is binary on 'whether job involves use of computerised or automated equipment'. '%PC in workplace' is 'In your workplace, what proportion of employees work with computerised or automated equipment?'. 'appraisal' is binary for 'do you have a formal appraisal system at your workplace'. 'learn new thing' is the reported agreeer8956928(orte28('learn)-282((k)-4282(new)-16F46 9.96)ne624e9 new thinse'.BApor is theroportigradufatis atf.527(the)-56[(industry-y)78(ear-region)-516(lev)81(e)-1leis the BSS1

Table 3: Estimating wage response to supply-side shifts, by industry

Dependent var: $\log w_{gjt} = \log w_{g1t}$				
	natural resources	manufacturing	construction	trade
log emp ratio	0.2773 (0.4058)	0.0956 (0.1323)	-0.4225 (0.2629)	0.1665 (0.5868)
j-specific trend	yes	yes	yes	yes
Observations	200	200	200	200
	transport, information	finance, business serv	other services	
$\ln y_{gjt}/y_{g1t}$	-0.8401 (1.3341)	0.0048 (0.2665)	0.3706 (0.1446)	
j-specific trend	yes	yes	yes	
Observations	200	200	200	

Standard errors in parentheses

$p < 0.05$, $p < 0.01$, $p < 0.001$

Note: The dependent variable is log hourly wage ratio at the industry-occupation-year level. The definition of industry and occupation is the same as the rest of the paper. Occupation 1 is the reference occupation group. The key regressor is log occupational employment ratio $\ln emp_{gjt} = \ln emp_{g1t}$, where $\ln emp_{gjt}$ is the total hours in the g - j - t cell. The instruments for $\ln emp_{gjt} = \ln emp_{g1t}$ are $supply_{gjt}$; $supply_{g1t}$. $supply_{gjt}$ is a shift-share instrument at the g ; j ; t level, using contemporary shares of demographic groups and historical mapping from demographic groups to g ; j cells. Source: LFS 1993-2017.

New technology: w_{gt} will increase. If instead, we have $\frac{N_{gj}}{O_{gj}} < \frac{O_{gj}}{N_{gj}}$, then w_{gt} will fall. In either case, the term $(1 - w_{gt})r_{gj}^O + w_{gt}r_{gj}^N$ will increase. This will partially offset the negative effect through the first term.

Now let's see how wages have responded to supply-side shifts in the UK data. Specifically, we will regress the log occupational wage ratio on the log occupational employment ratio and a j - g -specific time trend:

$$\ln\left(\frac{p_{gjt}}{p_{g1t}}\right) = (\quad - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + \sum_{k=0}^5 \frac{X_{gj}^k}{k!} t^k + u_{gjt} \quad (21)$$

The log employment ratio will be instrumented by supply-side shifts. The instruments are of shift-share style, using the shift in the demographic composition of the population (defined by education-gender-age) and historical mappings from each demographic group to tasks. Thus, it captures variation that comes from aggregate changes in the demographic composition. The coefficient on the log employment ratio ($\quad - 1$) is interpreted as the slope of the demand curve. The specification of time trend is a 5th order polynomial of year, plus two dummies to capture classification discontinuities over 2000-1 and 2010-11. The regression is run separately by industry.

The results are reported in table 3. I find that the key estimate ($\quad - 1$) is small and not significantly different from zero in most industries. It is negative in only

two out of seven industries, and it is significantly positive in one industry. The instruments are reasonably strong: the standard errors are small enough to rule out $(\beta - 1) < -1$ in most industries. Overall, the estimates suggest the demand curve is not as downward-sloping as would be expected from standard models. My framework with endogenous technical change offers an explanation as to why it may be flat.

The finding that occupational wages do not respond negatively to supply-side shifts in the above regression analysis is not surprising, given that the canonical SBTC model with two education groups has been shown to provide a poor fit of UK data (Blundell et al., 2022).⁴⁴

6 Empirical results

I calibrate two of the structural parameters and estimate the rest. I calibrate $\sigma = 0.1$ and $\alpha = 0.1$. $\sigma = 0.1$ corresponds to Goos et al. (2014)'s 0.9 estimate of the substitution elasticity between tasks. I have experimented with several values of σ and found $\sigma = 0.1$ yields a good fit of the data overall.

Given the calibrated σ and α , I estimate all the other structural parameters according to the methods discussed in section 3.3. Given all the parameters, I solve for the equilibrium $(p_t; w_t)$ in each year. I search for the equilibrium that is closest to the observed and satisfies

tasks. This is what we expect. And this is driven by the data: within manufacturing employment has shifted substantially away from manual routine to managerial. Meanwhile, in non- financial services, the new technology is less intensive in admin

Figure 6: Estimated task intensities in each industry






Figure 7: Fit of occupation employment share



Note: The actual time trends of occupational employment shares are solid lines. The corresponding baseline predictions are dashed lines of the same colour.


Figure 8: Fit of log task prices P_{jt}



Note: The actual time trends of task prices are solid lines. The corresponding baseline predictions are dashed lines of the same colour.


6.2 Counterfactuals

Figure 9: Counterfactual: only skills distribution shifted



Note: using the lagged $\log P_{jt} ; w_{gt}$ as the benchmark. For log task prices, we normalise the average change across 9 occupations to 0.

Figure 10: Counterfactual: only industry demand shifts



Note: using lagged $\log P_{jt} ; w_{gt}$ as the benchmark. For log task prices, we normalise the average change across 9 occupations to 0.

stract tasks. This technology shift helps to absorb the impact of the supply shock on wages. As a result, we get substantial movements in employment shares, little changes in occupational wages, and little change in the mapping from skills to occupation. To the extent that the skills distribution within graduates are stable, the model predicts little occupational downgrading within graduates.

The calibrated model can't UK data well over 1997-2015. While the estimated direction of technical change varies across industries, the overall pattern is that the New technology is less intensive in all three routine tasks and more intensive in managerial and professional tasks, with less difference in other tasks. The shift in skills distribution alone can account for between a third and two thirds of the actual decline in routine manual occupations, and between a third and half of the increase in each of the three abstract occupations. The shift in industry demand can account for similar magnitudes of employment declines in routine manual occupations and increases in professionals and technicians.

While this paper focuses on the UK, it provides a promising framework to study issues around occupations and education in other advanced economies other than the US. Many of these countries share some of the key facts observed in the UK since the 90s. First, like the UK, employment growth has been strongest in high-paid occupations in most European countries. This is consistent with the New technology being more intensive in abstract tasks. Second, occupational wages did not polarise outside the US. And third, the US had the highest proportion of graduates in 1990 and a slower increase afterwards than many European countries.

of skill-supply-induced adoption of technology might be much smaller than other factors in the determination of occupational trends.

Finally, the proposed framework offers a data-driven approach to answer several policy questions about the labour market. By having analytical and social skills (instead of education) as determinants of worker productivity, it allows a lot of heterogeneity within education groups and opens up the possibility of modelling changes in the group-specific skills distribution over time. The approach also makes clear that for analysing any policies that shift labor supply, it is important to model

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics* 130(4), 1781-1824.

Autor, D. (2022). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. Technical report, National Bureau of Economic Research.

Autor, D. H. (2013). The 'task approach' to labor markets: An overview. Technical Report 13-04, MIT Department of Economics Working Paper.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market.

- Borghans, L. and B. ter Weel (2008). Understanding the technology of computer technology diffusion: Explaining computer adoption patterns and implications for the wage structure. *Journal of Income Distribution* 17(3), 37{70.
- Boustan, L. P., J. Choi, and D. Clingingsmith (2022). Automation after the assembly line: Computerized machine tools, employment and productivity in the united states.
- Breemersch, K., J. P. Damijan, and J. Konings (2017). Labour market polarization in advanced countries. Technical Report 197, OECD.
- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002, 02). Information technology, workplace organization and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics* 117(1), 339{376.
- Card, D. and T. Lemieux (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *The Quarterly Journal of Economics* 116(2), 705{746.
- Carneiro, P. M., K. Liu, and K. G. Salvanes (2018). The supply of skill and endogenous technical change: evidence from a college expansion reform. *NHH Dept. of Economics Discussion Paper* (16).

- Goos, M. and A. Manning (2007, 02). Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics* 89(1), 118{133.
- Goos, M., A. Manning, and A. Salomons (2014, August). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review* 104(8), 2509{26.
- Graetz, G. and G. Michaels (2018, 12). Robots at Work. *The Review of Economics and Statistics* 100(5), 753{768.
- Green, D. A. and B. M. Sand (2015). Has the canadian labour market polarized? *Canadian Journal of Economics* 48(2), 612{646.
- Green, F. and G. Henseke (2016). The changing graduate labour market: analysis using a new indicator of graduate jobs. *IZA Journal of Labor Policy* 5(1), 1{25.
- Green, F. and G. Henseke (2021). Europe's evolving graduate labour markets: supply, demand, underemployment and pay. *Journal for Labour Market Research* 55(1), 1{13.
- Hardy, W., R. Keister, and P. Lewandowski (2018). Educational upgrading, structural change and the task composition of jobs in europe. *Economics of Transition* 26(2), 201{231.
- Hershbein, B. and L. B. Kahn (2018, July). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review* 108(7), 1737{72.
- Humlum, A. (2019). Robot adoption and labor market dynamics. Technical report, Princeton University.
- Johnson, M. and M. P. Keane (2013). A dynamic equilibrium model of the us wage structure, 1968-1996. *Journal of Labor Economics* 31(1), 1{49.
- Kampelmann, S. and F. Rycx (2011). Task-biased changes of employment and remuneration: The case of occupations.
- Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics* 107(1), 35{78.

- Machin, S. and J. Van Reenen (1998). Technology and changes in skill structure: evidence from seven oecd countries. *Quarterly journal of economics* 113(5), 1215{1244.
- Michaels, G., A. Natraj, and J. Van Reenen (2014, 03). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics* 96(1), 60{77.
- Murphy, E. C. and D. Oesch (2018). Is employment polarisation inevitable? occupational change in ireland and switzerland, 1970{2010. *Work, employment and society* 32(6), 1099{1117.
- Naticchioni, P., G. Ragusa, and R. Massari (2014). Unconditional and conditional wage polarization in europe.
- Oesch, D. and G. Piccitto (2019). The polarization myth: Occupational upgrading in germany, spain, sweden, and the uk, 1992{2010. *Work and Occupations* 46(4), 441{469.
- Patel, A. (2022). the role of firms in shaping job polarization. Technical report.
- Salvatori, A. (2018). The anatomy of job polarisation in the uk. *Journal for labour market research* 52(1), 1{15.
- Traiberman, S. (2019, December). Occupations and import competition: Evidence from denmark. *American Economic Review* 109(12), 4260{4301.

A Appendix

A.1 Derivation of a demand-side equation

In this section, we will derive a prediction about the relationship between task price ratio and task quantity ratio. That is equation (10) in the paper.

The F.O.C. with regard to task j for a firm using technology T is:

$$p_{jt} = p_{gt} \frac{\partial Y_{gt}^T}{\partial y_{jt}^T} = p_{gt} \frac{Y_{gt}^T}{y_{jt}^T} (y_{gt}^T = Y_{gt}^T)^{-1} \delta_j; g; t; T \in \{O, N\} \quad (22)$$

Apply $j = 1$ to (22) and take the ratio of the same equation between j and 1, we get

$$\frac{p_{jt}}{p_{1t}} = \frac{Y_{gt}^T}{Y_{g1t}^T} \left(\frac{y_{gjt}^T}{y_{g1t}^T} \right)^{-1} \delta_j; g; t; T \in \{O, N\} \quad (23)$$

$$\frac{y_{gjt}^T}{Y_{g1t}^T} = \left(\frac{p_{jt}}{p_{1t}} \frac{Y_{g1t}^T}{Y_{gt}^T} \right)^{-\frac{1}{1-\delta_j}} \delta_j; g; t; T \in \{O, N\} \quad (24)$$

Because we don't directly observe technology, we don't observe y_{gjt}^T . What we can observe is industry-level occupational employment $EMP_{gjt} = y_{gjt}^O + y_{gjt}^N$.

$$\frac{EMP_{gjt}}{EMP_{g1t}} = \frac{y_{gjt}^O}{y_{g1t}^O + y_{g1t}^N} + \frac{y_{gjt}^N}{y_{g1t}^O + y_{g1t}^N} \quad (25)$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N} \frac{y_{gjt}^O}{y_{g1t}^O} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N} \frac{y_{gjt}^N}{y_{g1t}^N} \quad (26)$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N} \left(\frac{p_{jt}}{p_{1t}} \frac{Y_{g1t}^O}{Y_{gt}^O} \right)^{-\frac{1}{1-\delta_j}} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N} \left(\frac{p_{jt}}{p_{1t}} \frac{Y_{g1t}^N}{Y_{gt}^N} \right)^{-\frac{1}{1-\delta_j}} \quad (27)$$

Denote $w_{gt} = \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N}$. We can interpret w_{gt} as the share of 'New' technology in industry g at time t . Denote

$$r_{gj}^O = \left(\frac{Y_{gt}^O}{Y_{g1t}^O} = \frac{Y_{gt}^O}{Y_{g1t}^O} \right)^{1-\delta_j} \quad (28)$$

$$r_{gj}^N = \left(\frac{Y_{gt}^N}{Y_{g1t}^N} = \frac{Y_{gt}^N}{Y_{g1t}^N} \right)^{1-\delta_j} \quad (29)$$

Equation (27) simplifies to

$$\frac{EMP_{gjt}}{EMP_{g1t}} = \left(\frac{p_{jt}}{p_{1t}} \right)^{-\frac{1}{1-\delta_j}} [(1 - w_{gt}) r_{gj}^O + w_{gt} r_{gj}^N]$$

Flipping the task price ratio to the left hand side, we get

$$\ln\left(\frac{p_{jt}}{p_{1t}}\right) = (1 - \alpha) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + \alpha \ln[(1 - w_{gt})r_{gj}^O + w_{gt}r_{gj}^N] \quad (31)$$

A.2 Derivation of task supply equation

Let's denote expected task output conditional on observed skills as

$$y(a; s; j) = E[y(i; j) | a_i = a; s_i = s] \quad (32)$$

$$= k_j e^{a_j \alpha + s_j \beta} E[e^{i_j} | a_i = a; s_i = s] \quad (33)$$

Note that $y(a; s; j)$ does not condition on the actual occupational choices, which would be endogenous.

Going back to (7) and using (33) to substitute for $k_j e^{a_j \alpha + s_j \beta}$, we get

$$\begin{aligned} p_j(a; s; p) &= [e^{a_k \alpha + s_k \beta} k_k p_k]^{-1} = \frac{X}{j} [e^{a_j \alpha + s_j \beta} k_j p_j]^{-1} \\ &= [e^{a_k \beta} p_k y(a; s; k) = E[e^{i_j} | a_i = a; s_i = s]]^{-1} = \frac{X}{j} [e^{i_j} p_j y(a; s; j) = E[e^{i_j} | a_i = a; s_i = s]]^{-1} \\ &= [e^{a_k \beta} p_k y(a; s; k)]^{-1} = \frac{X}{j} [e^{i_j} p_j y(a; s; j)]^{-1} \end{aligned}$$

This last equation says occupation choice depends on task prices, occupation amenities p_j , and $y(a; s; j)$ for all j .

Given task prices, the supply of task j is

$$LS_j(p) = \frac{X}{Z^j Z} p_j(a; s; p) y(a; s; j) \quad (34)$$

$$= \int p_j(a; s; p) y(a; s; j) f(a; s) da ds \quad (35)$$

where $f(a; s)$ is the joint density function.

A.3 Derivation of when will firms be indifferent between two technologies

This section derives equation (11).

Given the CES production function, the cost of using technology T to produce one unit of output in industry g is

$$\text{unitcost}_{gt}^T = \left[\sum_j \left(\frac{p_{jt}}{p_{gt}} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}} \right]^{1-\alpha} = A_{gt}^T \quad (36)$$

The ratio of unit costs between the two technologies is:

$$\frac{\text{unitcost}_{gt}^N}{\text{unitcost}_{gt}^O} = \frac{A_{gt}^O}{A_{gt}^N} \left[\frac{P_j \left(\frac{N}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}}{P_j \left(\frac{O}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}} \right]^{1-\alpha} \quad (37)$$

When the two technologies in industry g have exactly the same unit cost, we have

$$\begin{aligned} & \left[\frac{P_j \left(\frac{N}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}}{P_j \left(\frac{O}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}} \right]^{1-\alpha} = \frac{A_{gt}^N}{A_{gt}^O} \\ &) \frac{P_j \left(\frac{N}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}}{P_j \left(\frac{O}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}}} = \left(\frac{A_{gt}^N}{A_{gt}^O} \right)^{\frac{1}{1-\alpha}} \\ &) \sum_j \left(\frac{N}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}} \left(\frac{A_{gt}^N}{A_{gt}^O} \right)^{\frac{1}{1-\alpha}} \sum_j \left(\frac{O}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}} = 0 \\ &) \sum_j \left[\left(\frac{N}{g_j} \right)^{\frac{1}{1-\alpha}} \left(\frac{A_{gt}^N}{A_{gt}^O} \right)^{\frac{1}{1-\alpha}} \left(\frac{O}{g_j} \right)^{\frac{1}{1-\alpha}} p_{jt}^{-\frac{1}{1-\alpha}} \right] = 0 \end{aligned}$$

A.4 Derivation of an equation to identify TFP terms

We can get A_{gt}^T as an analytical function of $(\frac{T}{g_j}; p_{jt}; p_{gt})$, assuming $\alpha \neq 0$. This is because the profit maximisation gives a FOC:

$$\begin{aligned} & p_{gt} A_{gt}^T \left[\sum_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]^{1-\alpha} \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} = p_{jt} \\ &) p_{gt} A_{gt}^T \left[\sum_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]^{1-\alpha} \left[\left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]^{-1} \left(\frac{T}{g_j} \right)^{\frac{1}{1-\alpha}} = p_{jt} \\ &) \left[\frac{P_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}}}{\left[\sum_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]} \right]^{-1} \left[\left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]^{-1} = \frac{p_{jt}}{p_{gt} A_{gt}^T \left(\frac{T}{g_j} \right)^{\frac{1}{1-\alpha}}} \\ &) \frac{P_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}}}{\left[\sum_j \left(\frac{T}{g_j} (y_{gjt}^T) \right)^{\frac{1}{1-\alpha}} \right]} = \left[\frac{p_{jt}}{p_{gt} A_{gt}^T \left(\frac{T}{g_j} \right)^{\frac{1}{1-\alpha}}} \right]^{-1} \\ &) 1 = \sum_j \left[\frac{p_{jt}}{p_{gt} A_{gt}^T \left(\frac{T}{g_j} \right)^{\frac{1}{1-\alpha}}} \right]^{-1} \\ &) (p_{gt} A_{gt}^T)^{-1} = \sum_j \left[\frac{p_{jt}}{\left(\frac{T}{g_j} \right)^{\frac{1}{1-\alpha}}} \right]^{-1} \quad (38) \end{aligned}$$

Figure 12: Employment and wage growth by ISCO major group




Source: SES 2002 and 2014. To compute the change in hourly wages, we exclude cells where the occupation's employment share has more than tripled or halved because those cases may involve large compositional changes.

B Additional figures


The flip side of the flat proportion of abstract occupations within graduates is that

Figure 13: Within-between decomposition of the change in occupational employment shares




Source: UK Labour Force Survey

Figure 14: Adjusting occupational wage for classification changes



Source: UK Labour Force Survey 1993-2017.

Figure 15: Estimated w_{gt} from 9 proxies measures



Note: We have 4 measures of capital composition from 1997 to 2015 annually and 5 measures from the BSS available at 4-5 points between 1992 and 2017. Because the different measures have different scales, I standardise each measure within industry so that when I minimise the sum of squared $\hat{w}_{gt}^m$, they are equally important. Finally, I smooth each time series with a cubic spline and constrain the value to be in the $[0,1]$ range: w_{gt} is assumed to follow a cubic spline in between each pair of nodes, nodes are 3 years apart from 1997 to 2015, the value in 1997 is constrained to be 0 and the 2015 value is constrained to be 1. Note that w_{gt} is not really comparable between industries, because of the affine transformation is within industry.

Figure 16: Fit of New technology's share w_{gt}

