Occupational polarisation and endogenous task-biased technical change

Wenchao Jin*
University College London and Institute for Fiscal Studies

February 3, 2022

Abstract

Since the 90s many developed countries have experienced job polarisation, which is defined as employment growth in both high-paid and low-paid occupations and a relative decline in middle-paying occupations. The most popular explanation is that recent technological change has been biased against routine tasks, which are more important in middle-paying occupations. This paper offers a new and complementary explanation that emphasises increasing skill supply and endogenous adoption of technology. I exploit the large policy-driven increase in education in the UK and argue that this supply shift has caused the adoption of routine-biased technology and thereby employment polarisation. This framework is consistent with two additional facts in the UK labor market. First, there were relatively little movements in occupational wages and the pattern is certainly not polarising. Second, over a period of rapidly increasing supply of graduates, occupational outcomes among graduates have not deteriorated much. I build and estimate a general equilibrium multi-sector model on UK data over 1997-2015. I find that in most industries, technical change over the period was biased against routine tasks and favoured managerial and professional tasks. Allowing endogenous technological change, I find that changes to the skills distribution and industry demand shifts can each explain about half of the decline in manual routine jobs.

*Email: wenchao_j@ifs.org.uk. I am extremely grateful to my PhD advisors Richard Blundell and Suphanit Piyapromdee for their guidance and encouragement on this project. I would like to thank David Autor, David Green, John Van Reenen, Lars Nesheim and Fabien Postel-Vinay for their insightful comments. I acknowledge financial support from the Economic and Social Sciences Research Council via the Institute for Fiscal Studies.

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 important implications for income inequality. In the polarisation literature, the leading explanation is Routine-Biased Technological Change (RBTC) or Task-biased technological change (TBTC). The basic idea is that technological change (embodied by IT and automation equipment) has displaced workers in carrying out routine tasks, which are important in middle-paying occupations. There is a large literature on this, both theoretical and empirical.[1] Most of it interpret RBTC as a consequence of increasing availability or productivity of capital equipments, or their declining prices. This paper offers a new perspective, which emphasises the role of increasing skills supply in the diffusion of routine-biased technology.

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).[2] The intuition is simple: when faced with two available technology options, if the local supply of skills is abundant, then economic incentives will be more likely to favour adopting the more skill-complementary technology.

This paper builds on the RBTC framework by allowing the adoption of technology to respond to skill supply shifts. This allows supply shifts to have different effects on the labour market than other theories of RBTC. The UK provides a uniquely-suitable empirical context to examine this, because it has experienced a huge policy-driven increase in higher education since the early 90s. By endogenizing the adoption of technology, my model can explain not only employment polarisation as a result of increasing skill supply, but also two additional phenomena observed in the UK.

First, there has been no wage polarisation in the UK. While the UK saw a significant shift in employment from middle-paying occupations to primarily high-paying ones and to a lesser extent low-paying ones, there were little relative movements in occupational wages since the mid-90s.[3] In fact, the movements in wages are uncorrelated with those in occupational employment in the UK. With endogenous adoption of Task-Biased Technical Change (TBTC), my model implies that skill supply shifts will primarily cause shifts in the occupational structure and have

[1] Acemoglu and Autor (2011) provides a good summary.

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

[3] Goos and Manning (2007) found substantial growth in both 'lousy and lovely jobs' over the 80s and 90s, but wages in lousy jobs were clearly falling relative to those in the middling jobs over their sample period.

less impact on wages.[4] Thus, it is consistent with job polarisation and a lack of wage polarisation at the same time.

The 2nd additional fact is more striking: the huge increase in education attainment since the early 90s has not led to much occupational downgrading for graduates in the UK. We will see in section 2, when the proportion of graduates in the workforce doubled from about 20% in the early 90s to over 40% in the mid-2010s, the share of graduates employed in abstract occupations has been stable around $75-80\%$. To the best of my knowledge, this fact has not been documented before.[5] This phenomenon is consistent with my model, and it would be harder to rationalise in models of exogenous technical change.[6] In fact, each of the three phenomena can be explained by alternative theories, but together they paint a picture consistent with the explanation proposed here. In section 5, I will discuss some regression results that reject the hypothesis of exogenous technical change and support my model.

The model proposed here is an equilibrium model with endogenous adoption of task-biased technologies in every industry. It is static because the aim here is to explain long-run trends. In each industry, firms choose between two technologies, which differ in task intensities. The choice depends on task prices. On the supply side, workers' productivity depends on two dimensions of observable skills and an unobserved general ability, and they choose their occupation based on their comparative advantage and task prices. When the skills distribution shifts in a way that favours a certain task, firms may switch towards the technology that's more intensive in that task, thus the resulting impact on prices and wages will be smaller than if technology is fixed. In other words, the endogenous adoption of technology helps absorb supply-side shocks, so the effects are seen in relative quantities of tasks rather than prices.

The idea that the adoption of technology responds to supply-side shifts is new to the polarisation literature (Autor et al., 2003; Goos et al., 2014; Goos and Manning, 2007; Acemoglu and Autor, 2011).The literature has interpreted the pervasiveness of employment polarisation across developed countries as a result of a global technology shock, while attributing the differences in wage trends to un-

[4]I call it TBTC because my model does not presume that technical change is biased against routine tasks per se; the direction of bias will be estimated from the data.

[5]There is an older and smaller literature on 'over-education', asking whether more graduates are now over-educated for their job (Dolton and Vignoles, 2000; Battu et al., 2000; Dolton and Silles, 2008). Sometimes the measure of 'over-education' is based on one's education relative to the mean or mode education level within the occupation, or the 'required' level as perceived by some professional or the workers themselves. While it is related, 'over-education' is more fluid concept.

[6]It would require the exogenous technical change to increase the demand for abstract tasks at the same time and by the same magnitude as the supply-side shift.

specified differences in institutions or differences on the supply side.[7] By contrast, the chain of events emphasised here starts with a positive supply shift (due to education policy in the UK), this causes task-biased technical change, and therefore leads to the three aforementioned facts about occupations. This is not a rejection of the hypothesis that technology shocks coming from cheaper machines are routine-biased. Such exogenous technical change is still allowed in my model; but the emphasis here is how the adoption of technology responds to supply-side shifts. This new feature is important because it yields different predictions for how a supply-side policy would affect the labor market. My model's ability to explain all three facts about occupations in the UK gives us confidence that it is a reasonable model for analysing potential policies in the UK, such as skill-based selection of immigrants.

The model also features flexible sectoral shifts. Because industries differ a lot in task intensities[8], differential productivity and demand trends in different industries may affect relative demand for different tasks. Exogenous factors such as population ageing and Chinese imports may lead to rising demand for personal services and falling demand for manufacturing goods. Counterfactual analysis using my model suggests between-industry demand shift played a big role in occupational polarisation over the period. When I bring my model to data, I allow 7 industries. This is a finer disaggregation than most papers in the polarisation literature. For example, Autor and Dorn (2013) distinguishes between low-skill services and the rest. Barany and Siegel (2018) built and calibrated a model of 3 sectors: low-skilled services, manufacturing and high-skilled services. They show that sectoral shift contributed significantly to changes in both occupational employment shares and occupational wages in the US since the 1950s.

The idea that technical change is endogenous is not new. Acemoglu (1998, 2002, 2003) argued that the extent of skill bias in the new technology is endogenous, which explained the apparent acceleration of skill-biased technical change after an initial increase in the supply of skills in the US. While such models of endogenous innovations are suitable for a big country on the technology frontier, like the US in the last 100 years; for countries that are followers, models of endogenous adoption of available technologies are more suitable. I believe overall the UK belongs to the latter group over my sample period 1997-2015. The UK had a much lower proportion of graduates than the US in the early 90s and have surpassed it by the end of my period. Moreover, models with endogenous innovations typically imply a downward-sloping short-run demand curve (just like the case with exogenous technology) and a flatter or upward-sloping long-run demand curve; whereas my

[7] As we'll discuss below, many developed countries experienced employment polarisation without wage polarisation (Green and Sand, 2015).

[8] For example, finance is intensive in professional task, and construction is intensive in skilled trades.

model with endogenous adoption imply a flatter short-run demand curve. The broad trends observed in the UK support the latter.

This paper also relates a growing literature on 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), automation (Aghion et al., 2020), or industrial robots (Graetz and Michaels (2018),Humlum (2019)). 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, 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. Instead, 'technology' is general-purpose within industry. 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 technical changes manifest differently in different firms. It could be automation equipment in manufacturing, some software in financial services, and some sort of organizational 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 proxies to measure the share of the 'New' technology at the industry-year level.

The paper is most closely related to Blundell et al. (2021). They 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 three facts about occupations and to allow policy analysis. First, the model in Blundell et al. (2021) has two labour inputs (graduates and others), whereas my model is about occupational tasks and it features multiple industries. Second, in this paper each worker has 2-dimensions of observable skills: analytical and social skills are what matters for productivity, not education per se.[11] This opens up the possibility of modelling changes in skills distribution within education groups over time, in a data-driven way. Third, they did not estimate

[9] Most 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 of 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.

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

[11] The model also allows an unobserved general ability, that can vary across individuals freely.

or calibrate their model, whereas I do. This means I can simulate the effects of counterfactual policies. Overall, my paper corroborates their story with a richer model and supportive empirical evidence. In addition to Blundell et al. (2021), Carneiro et al. (2018) and Dustmann and Glitz (2015) also found production technology responds to changes in the local supply of educated/uneducated workers. Like Blundell et al. (2021), they differentiate labor by education and have nothing to say about occupations.

Currently, the model is partly calibrated and partly estimated, at the level of 9 occupations and 7 industries. The 9 occupations are SOC2000 major occupation groups: 1 managerial 2 professional 3 technician 4 admin 5 skilled trades 6 personal services 7 sales and customer services 8 operatives 9 elementary. I will improve the estimation approach when time allows. As it stands, the model fits the UK trends pretty well. The good fit is not mechanically guaranteed by the model design, because most of the key parameters do not vary over time. The estimates suggest that technological change in the UK over 1997-2015 was biased against all three routine tasks, favoured managerial and professional tasks, and neither favoured nor biased against the remaining four (3 manual tasks and technicians).[12] Counterfactual analysis suggests that industry demand shift could explain about half of the decline of manual routine jobs, and so could the shift in the aggregate skills distribution.

The model offers a suitable framework to investigate some counterfactual policy questions. For example, how will further increase in higher education affect the wage structure? If the UK government select EU immigrants by skills[13], how will it affect the labour market? Such counterfactual scenarios are outside the range of historical observation and are difficult to answer by reduced-form methods. I plan to use the UK Skills for Life Survey to obtain the differences in literacy and numeracy skills between immigrants and natives, and to examine whether the skills distributions within education groups have deteriorated over birth cohorts.

For reasons that will be discussed in Section 2 and Conclusion, this paper also provides a promising framework that can be used to study occupational trends in other European countries. The paper is structured as follows. Section 2 documents three phenomena in the UK labour market, with comparison to other developed countries where possible. Section 3 develops the model and explains how to identify the unobserved technology share and the model parameters. Section 4 describes the data sources, including how we impute the unobservable technology share. Section 5 investigates various correlations in micro data, which are supportive

[12]This direction of biases are consistent with previous studies. For example, Humlum (2019) estimated that in Danish manufacturing, robot adoption reduced the productivity of production workers but increased that of tech workers (engineers, researchers and skilled technicians.

[13]It already does for non-EU immigrants, and now after Brexit it will be able to reject EU immigrants on skill-related criteria.

evidence of the model. Section 6 explains how various parameters are estimated, discusses some key estimates and the fit of the model, and conducts counterfactual analysis. The final section concludes.

2 Motivating facts

This section documents three phenomena in the UK labour market since the 90s. The first fact has been observed in many developed countries in the past couple of decades. The latter two pertain to the UK, which are clearly different to the US, and somewhat shared with other European countries.

1. Since the 90s, employment has shifted significantly away from middle-paying occupations towards both the high end and the low end.
2. There is no clear U-shaped pattern in occupational wage changes during the period of employment polarisation outside of the US.
3. The huge increase in education attainment in the UK has not led to much occupational downgrading, nor decline in the skilled wage premium.

The first point, of occupational polarisation, 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)). And it's been documented since the 1980s for the UK Goos and Manning (2007) and for Germany (Kampelmann and Rycx, 2011) and even earlier in 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 groups.[14] So in this section I present occupational facts at this level, too. At this level of nice 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 the '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%. Meanwhile, each of the three abstract occupations grew

[14]There are 9 occupations in total : 1, managerial, 2 professional, 3 technician, 4 admin, 5 skilled trade, 6 personal services, 7 customer services, 8 production and machine operatives, and 9 elementary.

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. Among the manual occupations, there is some decline in elementary occupations[15], which is more than compensated by the increase in personal services (such as care assistants).

At a similar level of aggregation, Figure 3 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. For example, Fernández-Macías 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.[16] Murphy and Oesch (2018) looked at Ireland and Switzerland over 1970-2010 and found 'occupational upgrading', and the patterns were consistent with changes on the supply side from women's education and immigration. It's beyond the scope of this paper to investigate why those studies reach different conclusions. But they all point to strong growth in high-paying occupations. And we see in both Figure 1 and Figure 3, 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 compared to the old technology, 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.

Meanwhile, apart from the US, there is no such V shape in occupational wage growth in other developed countries that also saw employment polarisation. This is our fact no.2.

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

[16] The only exception is for the earnings-based indicator in the UK, which suggests a polarising pattern.

Figure 1: Employment shares by occupation

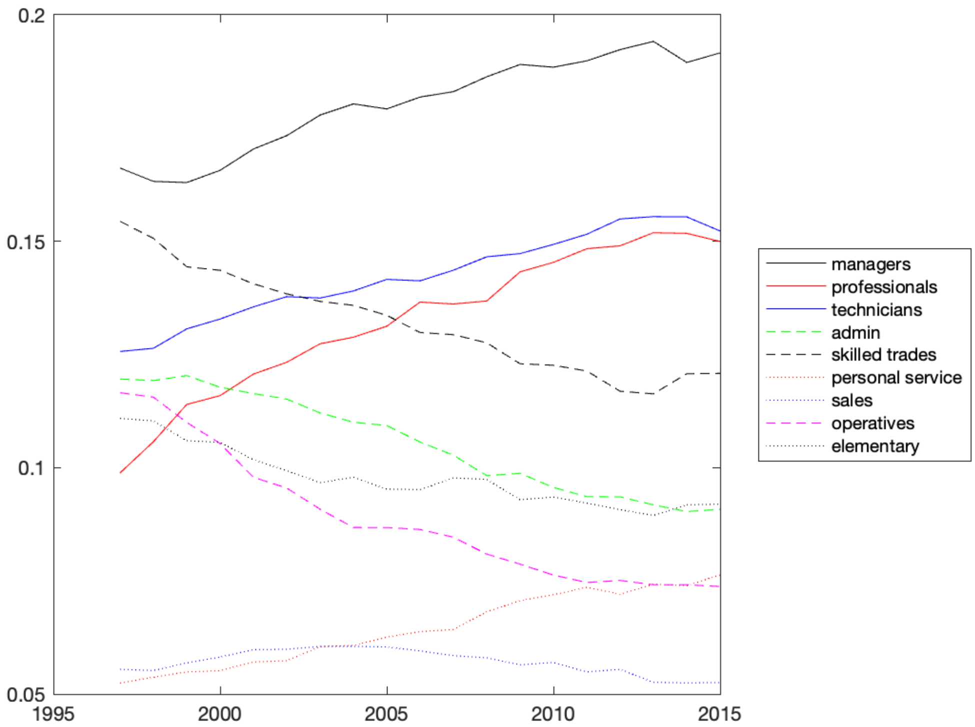

Note: the 9 occupations are major occupation groups under SOC2000. See section 4 for how we adjusted for discontinuities in SOC over 2000-01 and 2010-11.

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.[17] I plan to use another dataset to estimate occupational wages with individual fixed effects in the near future.[18] 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 did have the slowest wage growth. The maximum difference between occupations in log wage change over 1997-2015 is just under 0.08. This is not big, given the difference in occupational log wage level can be more than 0.5 between the top and bottom groups. This is also small relative to the observed changes in employment

[17]In 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.

[18]This has been delayed due to data access issues during the pandemic.

shares.[19]

In other European countries, we have not seen wage polarisation since the 80s either. Before 2000, occupational wage growth was marked by increased inequality across the distribution, in the UK during the 80s and 90s (Goos and Manning, 2007), in Germany in the 90s (Dustmann et al., 2009), and in Canada (Green and Sand (2015)). In fact, Green and Sand (2015) summarizes that occupational wage polarisation was only observed in the US in the 90s, and not elsewhere or in other decades. After the turn of the century, there was less or no increase in inequality between occupations. In figure 3, we see in a number of European countries, wage growth over 2002-14 tends to be slightly slower in high-paying occupations such as the professionals. Naticchioni et al. (2014) looked at twelve European countries (subset of EU15) over 1995-2007 and found no evidence of wage polarisation, whether using industry level or individual level data.

The leading explanation for the employment polarisation is routine-biased technical change (RBTC thereafter) (Autor et al. (2003); Acemoglu and Autor (2011) ,Autor and Dorn (2013),Hershbein and Kahn (2018), Goos and Manning (2007),Goos et al. (2014) , Michaels et al. (2014) and many others). Broadly speaking, the hypothesis is that technological changes (such as automation and ICT) were biased against routine tasks, which are important in the semi-skilled occupations around the middle of the distribution. Such a technology-induced demand shock causes polarisation.[20] This has a lot of intuitive appeal, and it fits the polarising trends in employment and wage in the US in the 90s. Guided by the RBTC hypothesis, some papers have asked directly whether occupational wage change correlates negatively with its 'routineness', and the answer is no for Germany and Sweden(Kampelmann and Rycx, 2011; Adermon and Gustavsson, 2015), and yes for the US(Firpo et al., 2011; Böhm, 2020; Acemoglu and Restrepo, 2021).[21]

[19]To give a sense of magnitude, if tasks are neither complements nor substitutes, the response of the log wage ratio to log quantity ratio along the demand curve would be -1. That is assuming no demand shift, an increase in the log quantity of professional tasks by 0.5 (its employment share increased from 10% to 15%) would reduce its log wage by 0.5.

[20]A secondary explanation is the sectoral shift away from manufacturing towards the services. This is also found to contribute to polarisation because manufacturing is more intensive in middle-paying occupations(Autor and Dorn (2013), Barany and Siegel (2018)). But this story is also about a shift in the demand curve.

[21]Adermon and Gustavsson (2015) examined occupational employment and wages in Sweden over 1975-2005, and found that TBTC could explain changes in within-occupation wage differentials but not between-occupation wage differentials. Kampelmann and Rycx (2011) found in Germany, routine jobs have lost employment but there is "no consistent task bias in the evolution of pay rules". By contrast, for the US, Firpo et al. (2011) found that both changes in within- and between-occupation wage differentials in the 90s are consistent with predictions from TBTC. Acemoglu and Restrepo (2021) finds that demographic groups who specialize in tasks that were automated experienced relative wage falls.

Figure 2: Changes in log wages at SOC2000 occupation main group

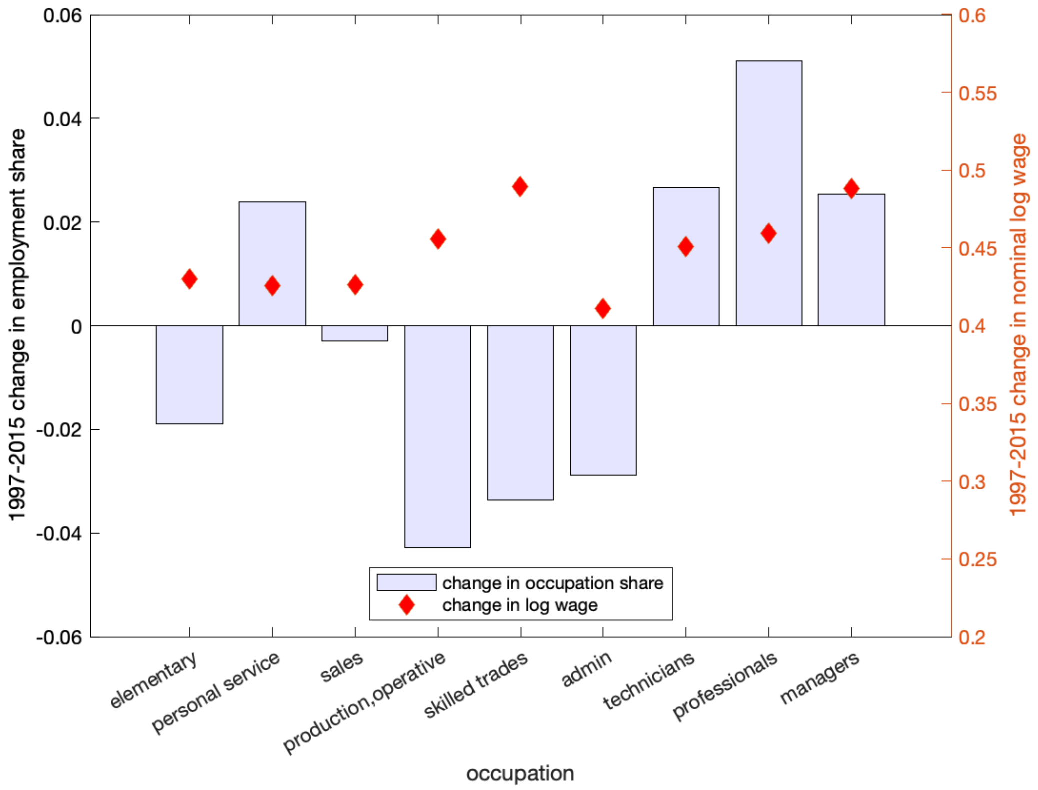

Note: in each year, we regress log wages on gender-age interactions, detailed education, and occupation dummies. This forms our 'composition-adjusted' occupational wage data log Pjt. Because the occupation classification changed in 2001 and 2011, we then fit each Pjt with a 5th-order polynomial with discrete jumps at 2001 and 2011, and subtract the estimated jump in both pre2001 and post2011 data. Here we show the change in the adjusted log Pjt between 1997 and 2015.

The lack of wage polarisation outside the US does not in itself reject the RBTC story. There are many reasons why RBTC could lead to substantial employment polarisation and no noticeable impact on observed wages: 1, the supply curve could have shifted at the same time in the same direction as the demand due to some exogenous reason. 2, supply could be highly elastic, which would be the case if wage is a key factor in people's selection of occupation and there isn't too high a barrier to switching occupations. 3, wages are sticky for institutional reasons. 4, observed wages are confounded by unobserved compositional changes. All these explanations could be true simultaneously.

It's beyond the scope of this paper to review the wage-setting institutions and supply-side changes in individual countries. What this paper offers is a unified

Figure 3: Employment and wage growth by ISCO major group

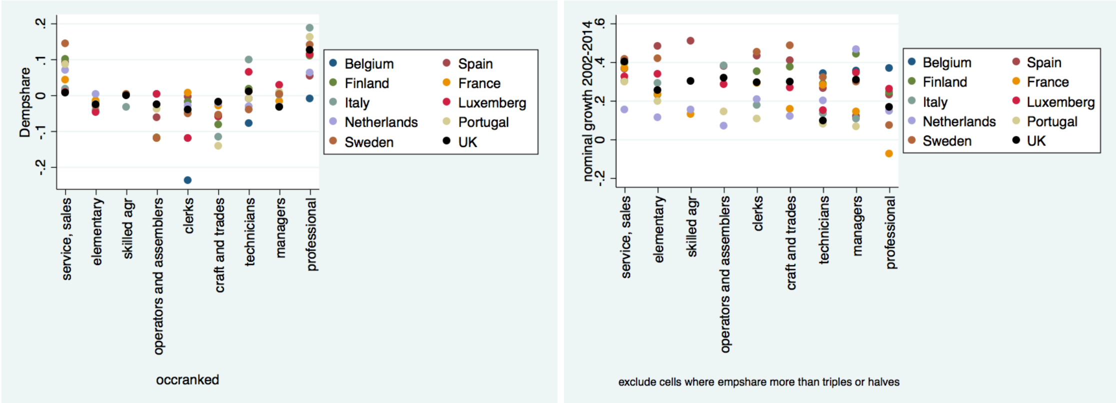

Source: SES 2002 and 2014

explanation of the facts without deviating from competitive labour markets. More importantly, the third fact (documented below for the UK) is a prediction from my model and is harder to explain otherwise.

It's worth stressing that the changes in occupational wages are not only small and dissimilar to employment changes, they are in fact uncorrelated with the movements in occupational employment in the UK. Section 5 further investigates this, comparing my model and the standard model with exogenous technology-driven demand shifts. The correlations between occupational wages and employment ratios do not support the latter.

Now let's turn to the third fact. The proportion of graduates has increased dramatically in the UK since the early 90s, with no significant deterioration in graduates' relative wage or occupation destinations.

This increase in education attainment was mostly driven by government policy. The vast majority of universities in the UK are publicly-funded: they receive direct grants from the government and tuition fees from students, who can take subsidised loans from the government. The Education Reform Act (ERA) of 1988 changed the funding formula of HE institutions and they responded by increasing their student intake dramatically. Then in 1994, the government introduced student number controls: the number of home students each university could admit every year were capped. This resulted in a steady increase in student numbers from then till the 2010s. In 2012, the cap was abolished for students whose grades are above a threshold. Since 2015, the cap was abolished for all. Throughout the period, entry to university was rationed by academic selection. Figure 4 shows that about 20% workers in the early 90s had higher-education qualifications, and this more than doubled over the next 2 decades. The pace of increase is much faster than the US.

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., 2021) documents this and explains it in a model of endogenous technology.

One might also expect the huge increase in graduates to lead to 'occupational downgrading', that is, an adverse shift in the occupation destinations of graduates over time. However, there has not been much occupational downgrading among graduates in the UK. To the best of my knowledge, this UK fact has not been documented before.

The right subgraph in Figure 4 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. This is striking when compared against the US. According to Beaudry et al. (2016), in the US the employment rate of cognitive occupations for college graduates fell by nearly 0.1 log point over 2000-2010. The UK trend is basically flat over the 2000s. And the UK saw a much faster increase of college graduates than the US.

To give a sense of magnitude, I calculate how much the abstract share needs to fall within education group if the aggregate abstract share had been constant while the education composition improves.[22] These counterfactual trends are plotted as dashed lines in Figure 4: 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 leads to an increase in the supply of skills; as those skills are relatively more important in abstract tasks; this would cause firms to switch to the abstract-task-intensive technology and create more abstract jobs.

Broadly speaking, most developed countries have seen some increase in tertiary education over the past couple of decades, and the UK is one of the countries with the fastest increase. The US, on the other hand, had the highest level to start with and a slower increase since the 90s compared to most European countries. According to Barro and Lee (2013), the proportion of 15-64 year olds 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 technology in general, with other developed countries closely behind. This means the latter group (including the UK) are in a position to choose between technologies that are not too different and so the choice should depend on prices and wages. Moreover, Blundell et al. (2021) shows that in 11 OECD

[22]In the counterfactual, the education-specific abstract share is proportional to its 1992 level, the aggregate abstract share is at the 1992 level, and the shares of education groups in the workforce are the actual values.

Figure 4: Proportion of graduates and their occupation destination

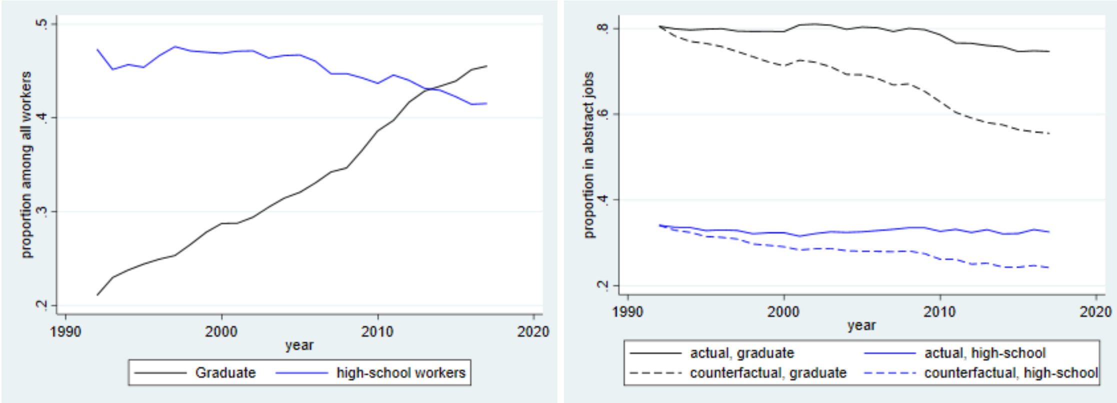

Note: graduates are people with NVQ level 4 qualifications or above. High-school workers refer to those with NVQ level 2 or 3 qualifications. First degrees are NVQ level 4. A-levels and post-16 further education qualifications are NVQ level 3. O-levels and GCSEs (grade C+) are NVQ level 2. 'Abstract'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

countries which experienced substantial increase in tertiary education, there was no significant decline in graduates' relative wages in 9 of them, like the UK. These similarities suggest that a model of endogenous adoption of technology might be more suitable for these non-US developed countries, whereas the US might need a model of endogenous innovations.

3 The model

This section develops a general-equilibrium model of tasks. It's 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. The endogenous adoption of technology means that one point of task prices is consistent with not one but a wide range of task ratios on the demand side. 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-changibly. In reality, the task content within occupations may change continuously, in response to changing demand for tasks. This is an interesting challenge for future research.[23] In this

[23]Conceptually, what matters in production is tasks, but what workers choose is occupation, The task content within occupation is a choice made by the firm, subject to potentially complex constraints (physical constraints, information constraints, supply constraints and so on). There's

paper, 'tasks' should be interpreted as the output of specific occupations. For example, professional task is simply the output of labour in professional occupations, whether the actual activity carried out is writing or analysing data is not studied here.

Each industry produces one good. Dente the goods as $g \in 1, 2, ..G$. The production of each good is a CES function of tasks $j \in 1, 2, ..J$, 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' tech and the 'New' tech. Firms are identical otherwise within the industry. The difference between two technologies is that they have different task intensities α_{gj}^T. They also have their own TFP term A_{gt}^T, which is neutral with regard to tasks.

$$Y_{gt}^T = A_{gt}^T [\sum_j \alpha_{gj}^T (y_{gjt}^T)^\rho]^{\frac{1}{\rho}} \quad , T \in \{O, N\} \tag{1}$$

where y_{gjt}^T is the amount of task j employed in industry g, using technology T at time t; α_{gj}^T is the share parameter of task j in technology T in industry g, note that it does not vary over time; ρ is 1 minus 1 over elasticity of substitution between tasks; ρ must be below 1. A negative ρ means tasks are complements. A_{gt}^T is Total Factor Productivity of technology T in industry g at time t. And Y_{gt}^T is the output produced by technology T in industry g at time t.

Consumers have CES preferences over G goods. A good produced by Old technology is a perfect substitute for the same good produced by New technology.

$$Q_{gt} = Y_{gt}^O + Y_{gt}^N \tag{2}$$

$$U_t = [\sum_g B_{gt} Q_{gt}^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \tag{3}$$

where Q_{gt} is output in industry g at time t, and B_{gt} captures time-varying demand for good g. B_{gt} is assumed to be exogenous here. For future research, it would be interesting to allow income growth to differentially affect the demand for goods and services.

Each technology has constant returns to scale. We normalize $\sum_j \alpha_{gj}^T = 1, \forall g, T$.

Because technology O and N differ in task intensities α. 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

also a question of how to organize all the tasks into bundles for each individual worker and then to combine them by management.

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 evolving share parameters in a CES production function, such as in Johnson and Keane (2013). The downsides are: 1) there is a lot more unobserved parameters to be estimated (one will need α_{gjt} instead of $\alpha_{gj}^{O}, \alpha_{gj}^{N}$), 2) 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.[24] 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 a way that happen to increase the demand for tasks that graduates have comparative advantage in, and at a speed that happen to leave the task prices and the mapping from education to occupation relatively unchanged. In subsection 5.1, I will formally test the hypothesis of exogenous task-biased technical change and reject it in favour of my model.[25] It is worth noting that my model allows for exogenous technical change as well, as the TFP trends A_{gt}^{T} are exogenous and a sudden 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 further assumes Cobb-Douglas production. For example, Autor (2013) define 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 is a Cobb-Douglas function of capital, human capital in each occupation and intermediate inputs produced in other industries.

One exception is Johnson and Keane (2013). Johnson and Keane (2013) differentiates labour by occupation, education, gender and age. Their production is multi-level nested CES: the 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. Their formulation is much more detailed than my model.

[24] The three abstract occupations clearly have much more graduates than other occupations, so I call them 'high-skilled'.

[25] It's a rejection of the hypothesis that all technical change is exogenous. That's not to say that there is no exogenous shock to technology.

To fit the US data over 29 years of data, they found 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. As we show later in this section, task-biased technological change in my model happens only through changing weights of New technology versus Old technology within industries. 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 i is endowed with two dimensions of observable skills: analytical ability a_i and social skill s_i; and a general ability μ_i. The joint distribution of a, s is assumed to be exogenous. Later on we will consider counterfactual policies that shift the skill distribution, through education or immigration. In reality, RBTC may induce workers to undertake more education or training in order to become more productive 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 worker i in occupation j produces is

$$y(i,j) = k_j e^{\beta_{aj} a_i + \beta_{sj} s_i + \mu_i} \tag{4}$$

This formula follows from Autor and Handel (2013), where I restrict the number of skills to 2-dimension rather than K-dimension. k_j is a j-specific scalar. μ_i is worker's general ability which is unobserved. It can be correlated with observed skills. The coefficients β_{aj}, β_{sj} determine which occupations reward which skills more. The key assumption here is that comparative advantage is captured by 2 dimensions of skills a_i, s_i; 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 equalize between firms that operate with different technologies and across industries. I denote the price of task j at time t as p_{jt}.

Because workers are perfectly substitutes in producing any given task (though individuals have different productivities), worker of ability a_i, s_i in occupation j in a firm adopting tech T gets paid the value of their output

$$W(i,j,t) = y(i,j) p_{j,t} \tag{5}$$

The utility worker i gets from occupation j is

$$U_{ij} = \ln(y(i,j)p_j) + \eta_j + e_{ij}, \quad j = 1, ...J \tag{6}$$

where η_j is occupation-specific amenities; e_{ij} follows iid Type-1 extreme value distribution, with location parameter at 0 and scale parameter ζ. In future, we may want to allow η_j to vary across demographic groups (e.g. gender, age, family type, and immigration status).

These preference shocks e_{ij} mean that for any given (a_i, s_i), there is positive probability of the worker going to any occupation j. The probability of worker i choosing occupation k is simply

$$\begin{aligned}
\pi_k(i, \mathbf{p}) &= (y(i,k)p_k e^{\eta_k})^{\frac{1}{\zeta}} / [\sum_j (y(i,j)p_j e^{\eta_j})^{\frac{1}{\zeta}}] && (7) \\
&= [e^{\beta_{ak}a_i + \beta_{sk}s_i + \mu_i + \eta_k} k_k p_k]^{\frac{1}{\zeta}} / \sum_j [e^{\beta_{aj}a_i + \beta_{sj}s_i + \mu_i + \eta_j} k_j p_j]^{\frac{1}{\zeta}} && (8) \\
&= [e^{\beta_{ak}a_i + \beta_{sk}s_i + \eta_k} k_k p_k]^{\frac{1}{\zeta}} / \sum_j [e^{\beta_{aj}a_i + \beta_{sj}s_i + \eta_j} k_j p_j]^{\frac{1}{\zeta}} && (9)
\end{aligned}$$

where $\mathbf{p}$ 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 β_{aj}. A smaller ζ means the preferences are less varied and so the wages are more important in determining the occupation choice. Note that the unobserved heterogeneity term μ_i does not enter into occupational choice. Thus $\pi_k(i, \mathbf{p}) = \pi_k(a_i, s_i, \mathbf{p})$.

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

$$\begin{aligned}
y(a, s, j) &= E[y(i,j)|a_i = a, s_i = s] && (10) \\
&= k_j e^{\beta_{aj}a + \beta_{sj}s} E[e^{\mu_i}|a_i = a, s_i = s] && (11)
\end{aligned}$$

Going back to (9) and substituting $y(a,s,j)/E[e^{\mu_i}|a_i = a, s_i = s]$ for $k_j e^{\beta_{aj}a + \beta_{sj}s}$, we get

$$\begin{aligned}
\pi_j(a, s, \mathbf{p}) &= [e^{\beta_{ak}a + \beta_{sk}s + \eta_k} k_k p_k]^{\frac{1}{\zeta}} / \sum_j [e^{\beta_{aj}a + \beta_{sj}s + \eta_j} k_j p_j]^{\frac{1}{\zeta}} \\
&= [e^{\eta_k} p_k y(a,s,j)/E[e^{\mu_i}|a_i = a, s_i = s]]^{\frac{1}{\zeta}} / \sum_j [e^{\eta_j} p_j y(a,s,j)/E[e^{\mu_i}|a_i = a, s_i = s]]^{\frac{1}{\zeta}} \\
&= [e^{\eta_k} p_k y(a,s,j)]^{\frac{1}{\zeta}} / \sum_j [e^{\eta_j} p_j y(a,s,j)]^{\frac{1}{\zeta}}
\end{aligned}$$

This last equation says occupation choice depends on task prices, parameters $\eta_j, \forall j$ and ζ, and $y(a, s, j)$.

Given task prices, the supply of task j is

$$LS_j(\mathbf{p}) = \sum_i \pi_j(a_i, s_i, \mathbf{p})y(i,j) \tag{12}$$

$$= \int\int \pi_j(a,s,\mathbf{p})y(a,s,j)f(a,s)dads \tag{13}$$

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

Thus, the only relevant unknowns on the supply side are η_j, ζ, $y(a,s,j)$ and $f(a,s)$. As long as we get $y(a,s,j)$, we don't need to estimate the distribution of unobserved heterogeneity μ_i or how it depends on (a,s), or the returns to skills β_{aj}, β_{sj}.

In the rest of this subsection, we derive a prediction about the relationship between task price ratio and task quantity ratio.

The FOC with regard to task j for a firm using technology T is:

$$p_{jt} = p_{gt}\frac{\partial Y_{gt}^T}{\partial y_{gjt}^T} = p_{gt}\alpha_{gj}^T(y_{gjt}^T/Y_{gt}^T)^{\rho-1} \quad \forall j,g,t,T \in \{O,N\} \tag{14}$$

Apply $j=1$ to (14) and take the ratio of the same equation between j and 1, we get

$$\frac{p_{jt}}{p_{1t}} = \frac{\alpha_{gj}^T}{\alpha_{g1}^T}(\frac{y_{gjt}^T}{y_{g1t}^T})^{\rho-1} \quad \forall j,g,t,T \in \{O,N\} \tag{15}$$

$$\frac{y_{gjt}^T}{y_{g1t}^T} = (\frac{p_{jt}\alpha_{g1}^T}{p_{1t}\alpha_{gj}^T})^{\frac{1}{\rho-1}} \quad \forall j,g,t,T \in \{O,N\} \tag{16}$$

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

$$\frac{EMP_{gjt}}{EMP_{g1t}} = \frac{y_{gjt}^O}{y_{g1t}^O + y_{g1t}^N} + \frac{y_{gjt}^N}{y_{g1t}^O + y_{g1t}^N} \tag{17}$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N}\frac{y_{gjt}^O}{y_{g1t}^O} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N}\frac{y_{gjt}^N}{y_{g1t}^N} \tag{18}$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N}(\frac{p_{jt}\alpha_{g1}^O}{p_{1t}\alpha_{gj}^O})^{\frac{1}{\rho-1}} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N}(\frac{p_{jt}\alpha_{g1}^N}{p_{1t}\alpha_{gj}^N})^{\frac{1}{\rho-1}} \tag{19}$$

Denote $w_{gt} = y_{g1t}^N/(y_{g1t}^O + y_{g1t}^N)$. We can interpret w_{gt} as the share of technology D in industry g at time t. Denote

$$r_{gj}^O = (\alpha_{gj}^O/\alpha_{g1}^O)^{1/(1-\rho)} \tag{20}$$

$$r_{gj}^N = (\alpha_{gj}^N/\alpha_{g1}^N)^{1/(1-\rho)} \tag{21}$$

Equation (19) simplifies to

$$\frac{EMP_{gjt}}{EMP_{g1t}} = (\frac{p_{jt}}{p_{1t}})^{\frac{1}{\rho-1}}[(1-w_{gt})r^O_{gj} + w_{gt}r^N_{gj}] \tag{22}$$

The last term is a weighted average between two technologies, where the weight w_{gt} is endogenous.

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

$$\ln(\frac{p_{jt}}{p_{1t}}) = (\rho-1)\ln\frac{EMP_{gjt}}{EMP_{g1t}} + (1-\rho)\ln[(1-w_{gt})r^O_{gj} + w_{gt}r^N_{gj}] \tag{23}$$

Equation (23) looks like a typical demand-side equation from the Skill-Biased Technical Change literature, where the term $\ln[w_{gt}r^O_{gj} + (1-w_{gt})r^N_{gj}]$ would represent technical changes. But it has a particular functional form: it's like a weighted average between two technologies, where the weight is at the industry-year level. The standard equation in the SBTC literature would have an exogenous time trend to represent technological progress (for example Katz and Murphy (1992) just had a linear time trend and Johnson and Keane (2013) had 3rd or 4th order polynomial). In the context of several j, g, the standard exogenous SBTC specification would be j-g-specific time polynomial. I will compare my model against the hypothesis of exogenous technical change in section 5, and show that the evidence supports mine.

3.1 equilibrium characteristics and effect of a supply-side shift

Denote $\omega_{gt} = Y^N_{gt}/(Y^N_{gt}+Y^O_{gt})$, the share of output produced by the new technology. So it's not the same as $w_{gt} = y^N_{g1t}/(y^N_{g1t} + y^O_{g1t})$, the share of new technology in terms of employment in the first occupation. But they are very strongly positively correlated. I define the equilibrium as log task prices ($\log \mathbf{p_{jt}} = \{\log P_{1t}, ...P_{Jt}\}$) and technology shares ($\omega_\mathbf{t} = \{\omega_1, ...\omega_{Gt}\}$) such that demand equals supply in each task, and that in each industry, the lower-cost technology is adopted. Both are adopted if their unit costs are equal.

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

$$unitcost^T_g = [\sum_j (\alpha^T_{gj})^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}]^{1-1/\rho}/A^T_g \tag{24}$$

The ratio of unit costs between the two technologies:

$$\frac{unitcost^N_g}{unitcost^O_g} = \frac{A^O_g}{A^N_g}[\frac{\sum_j (\alpha^N_{gj})^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha^O_{gj})^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}}]^{1-1/\rho} \tag{25}$$

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

$$[\frac{\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}}]^{1-1/\rho} = \frac{A_g^N}{A_g^O} \tag{26}$$

$$\Rightarrow \frac{\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}} = (\frac{A_g^N}{A_g^O})^{\frac{\rho}{\rho-1}} \tag{27}$$

$$\Rightarrow \sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}} - (\frac{A_g^N}{A_g^O})^{\frac{\rho}{\rho-1}} \sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}} = 0 \tag{28}$$

$$\Rightarrow \sum_j [(\alpha_{gj}^N)^{\frac{1}{1-\rho}} - (\frac{A_g^N}{A_g^O})^{\frac{\rho}{\rho-1}} (\alpha_{gj}^O)^{\frac{1}{1-\rho}}] p_j^{\frac{\rho}{\rho-1}} = 0 \tag{29}$$

The last equation is linear in $p_j^{\frac{\rho}{\rho-1}}$. Given alphas, this equation may have no solution in the positive domain if the TFP ratio is very far from 1. In that case, one technology will dominate in that industry. When the TFP ratio is not extreme, there are likely infinitely many points in the $(p_j > 0, 1 \leq j \leq 9)$ space that would equalize the unit costs between the two technologies in all 7 industries.

Given all task prices, the demand for task j is

$$\sum_g [\delta_{jg}^N(\mathbf{p_{jt}}) Q_g(\mathbf{p_{jt}}) \omega_g + \delta_{jg}^O(\mathbf{p_{jt}}) Q_g(\mathbf{p_{jt}})(1-\omega_g)] \tag{30}$$

$$= \sum_g (\delta_{jg}^N(\mathbf{p_{jt}}) - \delta_{jg}^O(\mathbf{p_{jt}})) Q_g(\mathbf{p_{jt}}) \omega_g) + \sum_g \delta_{jg}^O(\mathbf{p_{jt}}) Q_g(\mathbf{p_{jt}}) \tag{31}$$

where

- δ_{jg}^T is the amount of task j required by tech T to produce one unit of output in industry g, and it's a function of all task prices $\mathbf{p_{jt}}$;
- industry output Q_g is a function of $\mathbf{p_{jt}}$ through industry prices $\mathbf{p_{gt}}$. It does not depend on $\mathbf{w_t}$.

Task demand is not one single point given task prices $\{p_1, ...p_J\}$. Instead, movements in $0 \leq \omega_g \leq 1$ allows task demand to move within the cone of diversification. The cone of diversification has as many dimensions as the number of industries where the unit costs are equal. For a majority of years in our sample period (1997-2015), it has 7 dimensions.

Recall (13), the supply of task j takes this form:

$$LS_j(\mathbf{p_{jt}}) = \int\int \pi_j(a, s, \mathbf{p_{jt}}) y(a, s, j) f(a, s) dads \tag{32}$$

where $\mathbf{p_{jt}}$ is the vector of all task prices, $f(a,s)$ is the joint density function, and $y(a,s,j)$ is the amount of task j that workers with skills (a,s) will produce. The latter two do not depend on $\mathbf{p_{jt}}$.

Market clearing requires:

$$\sum_g (\delta^N_{jg}(\mathbf{p_{jt}}) - \delta^O_{jg}(\mathbf{p_{jt}}))Q_g(\mathbf{p_{jt}})\omega_g + \sum_g \delta^O_{jg}(\mathbf{p_{jt}})Q_g(\mathbf{p_{jt}}) - LS_j(\mathbf{p_{jt}}) = 0 \tag{33}$$

Given all task prices, these market-clearing constraints are a system of 9 linear equations, linear in the 7-element vector w_g.

When some supply-side shock shifts the supply curve (shifting the function $LS_j(.)$ as in (32)), it's possible that a change in ω_g will clear the markets without any change in $\{p_1, ...p_J\}$. This requires the shift in LS_j to be in the cone of diversification. In other words, the shift between technologies may absorb supply-side shocks and leave the equilibrium task prices unchanged.[26] Recall that occupational choice probabilities are functions of two skills and task prices. When task prices do not change, the occupational employment shares conditional on skills will not change. This is consistent with the UK fact that during a period of rapid increases in higher education, the occupational destinations among graduates did not change very much (Figure 4). The small amount of occupational downgrading observed within education groups could be interpreted as the education-specific distribution of skills having deteriorated slightly.[27] In short, through technology shifts, an increase in the supply of skills can leave the task prices unchanged, the occupation destinations conditional on skills unchanged, and increase the aggregate share of abstract jobs.

Finally, let's discuss the possibility of there being multiple equilibria of $\omega_\mathbf{t}$ with the same task prices. $\omega_\mathbf{t}$ only enters the system of linear equations that equals demand to supply (33). The coefficient on ω_g in the system is

$$(\delta^N_{jg} - \delta^O_{jg})Q_g \tag{34}$$

$$= [(\frac{p_j}{\alpha^N_{gj}p_g})^{\frac{1}{\rho-1}} - (\frac{p_j}{\alpha^O_{gj}p_g})^{\frac{1}{\rho-1}}](cons/p_g)^\sigma \tag{35}$$

$$= [(\frac{1}{\alpha^N_{gj}})^{\frac{1}{\rho-1}} - (\frac{1}{\alpha^O_{gj}})^{\frac{1}{\rho-1}}](p_j)^{\frac{1}{\rho-1}} cons^\sigma p_g^{-\frac{1}{\rho-1}-\sigma} \tag{36}$$

where $p_g = \frac{1}{A^T_{gt}}[\sum_j (\alpha^T_{gj})^{\frac{1}{1-\rho}} p_j^{\frac{\rho}{\rho-1}}]^{\frac{\rho-1}{\rho}}$ and $cons$ is a scalar that does not vary by j or g. Suppose the two technologies are sufficiently different ($\alpha^N_{gj} \neq \alpha^O_{gj}$), then there

[26] If the supply-side shocks are outside the cone of diversification, then price changes will be necessary to return the economy to equilibrium.

[27] In the empirical estimation part of the paper, I will assume that the skills distribution is fixed within education-gender. I hope to relax this assumption in future: when I find skills data for cohorts with different education composition, the data will tell whether the skills distribution of graduates worsens when a larger share of the cohort are graduates.

is no reason to expect the vectors of coefficients to be linearly dependent. Thus, it's likely that given $\{p_1, ...p_J\}$, there won't be multiple solutions of $\mathbf{w_t}$ clearing all the task markets.

3.2 identification of w_{gt}

We don't observe w_{gt} directly, nor do we observe y^O_{gjt}, y^N_{gjt} separately as opposed to $y^O_{gjt} + y^N_{gjt}$. If we knew ρ, we could use observed $y^O_{gjt} + y^N_{gjt}$ to obtain $(1 - w_{gt})r^O_{gj} + w_{gt}r^N_{gj}$ through (23),. However, r^O_{gj}, r^N_{gj} are also unknown. In fact, the level of w_{gt} is not identified even if we directly observe $(1-w_{gt})r^O_{gj} + w_{gt}r^N_{gj}$. To see why, consider an affine transformation of w_{gt} :

$$\begin{aligned}
\hat{w}_{gt} &= kw_{gt} + c, \forall t \\
\hat{r}^N_{gj} &= r^O_{gj} + \frac{1-c}{k}(r^N_{gj} - r^O_{gj}), \forall j \\
\hat{r}^O_{gj} &= r^O_{gj} - \frac{c}{k}(r^N_{gj} - r^O_{gj}), \forall j
\end{aligned}$$

The transformed case is observationally equivalent to the original one:

$$(1 - w_{gt})r^O_{gj} + w_{gt}r^N_{gj} = (1 - \hat{w}_{gt})\hat{r}^O_{gj} + \hat{w}_{gt}\hat{r}^N_{gj}, \forall j, t$$

Therefore, we will anchor the time series $\{w_{gt}\}$ by assuming $w_{g0} = 0, w_{gT} = 1, \forall g$. 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, such that $z_N > z_O$. The assumption here is that all firms with the New tech share the same level of z, which is higher than the level among old-tech adopters. There is no time variation within z_N or z_O. In practice, we will try several measures of z. We observe z over time and at the industry level.

$$z_{gt} = (1 - \tilde{w}_{gt})z_O + \tilde{w}_{gt}z_N \tag{37}$$

where $\tilde{w}_{gt}$ is the scale of new technology adoptors relative to the entire industry. If z_{gt} comes from employee survey, $\tilde{w}_{gt}$ is the employment share of firms using the new technology in the industry-year. As we anchor $\tilde{w}_{gt}$ to 0 at one point and 1 at another point, we would be setting $z_O = \tilde{z}_{g0}, z_N = \tilde{z}_{gT}$. Thus, we can impute w_{gt} as $\frac{\tilde{z}_{gt} - \tilde{z}_{g0}}{\tilde{z}_{gT} - \tilde{z}_{g0}}$. Thus, w_{gt} is just-identified by one proxy up to an affine transformation. If we have several measures of z, we can allow errors in equation (37). In section 4.3, we will assume a latent factor model to impute w_{gt}.

3.3 identification of model parameters

The parameters fall into two broad categories: supply-side and demand-side.

On the supply side, the unknowns are: η_j, the preference for working in task j; ζ, the scale of preference shocks; $f(a, s)$, the joint distribution of analytical and social skills ; and $y(a, s, j)$, the expected task output conditional on skills (a, s). Note that we don't need the returns to skills per se, just the four listed here. The reason was explained around equation (13).

η_j is the preference for working in task j, and we normalize $\eta_1 = 0$. The higher η_j, the more people will select into task j, all else equal. Therefore, η_j can be identified from the occupational employment shares in any given year. If we allow η_j to vary over time without any restriction, we could fit employment shares in every year perfectly. For now, I choose to have fixed η_j, so that no changes in employment will be attributed to preference shifts. Empirically, I search for η_j to match the observed employment shares in LFS 2006 (the mid-point of my sample period).

The smaller ζ is, the more elastic task supply will be with regard to task prices. The identification of ζ 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 ζ is small. Currently we search over a grid of ζ to minimize a loss function. In the next revision, I will use instruments for demand shocks to estimate the elasticity of task supply.

The joint skill distribution comes from the numeracy score and the literacy score the British Cohort Studies, measured at age 36. They are summarized to 7 points of support in each dimension.[28] 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 conditional on gender and education.[29] We obtain the distribution from the BCS for each gender-education, get gender-education weights from the LFS for each year, and aggregate up to the aggregate distribution.

Because this is a competitive labour market, workers are paid their task output times task price. We can get wages conditional on skills directly from the BCS, and dividing it by p_{jt} gives us the expected task output conditional on skills.

[28] Currently 7 is selected so that each group has at least 10% density. In future, I will experiment with having more or fewer points of support.

[29] In future, I will use other data to test this assumption, by comparing between generations who have very different education composition. This cannot be tested in the BCS because it contains only one birth cohort.

On the demand side, the unknowns are: ρ, which governs the substitution elasticity between tasks; tasks intensities $\alpha^T_{gj}, T \in \{O, N\}, 1 \leq g \leq G, 1 \leq j \leq J$; TFP trends $A^T_{gt}, T \in \{O, N\}, 1 \leq g \leq G, \forall t$; industry demand B_{gt}; and σ, which governs consumers' substitution elasticity.

In principal, ρ is identified from (23): the correlation between log wage ratio and log quantity ratio conditional on the technology share w_{gt}. Because we don't observe w_{gt} directly, non-linear estimation of (23) is tricky. I hope to estimate ρ directly from (23) in the near future. Currently, I calibrate $\rho = -0.1$, which corresponds to Goos et al. (2014)'s estimate of the substitution elasticity between tasks at 0.9.

Given ρ, all the other production parameters are well-identified.

Recall equation (23):

$$\ln(\frac{p_{gjt}}{p_{g1t}}) - (\rho - 1)\ln\frac{y^O_{gjt} + y^N_{gjt}}{y^O_{g1t} + y^N_{g1t}} = (1 - w_{gt})r^O_{gj} + w_{gt}r^N_{gj} \tag{38}$$

$$= r^O_{gj} + (r^N_{gj} - r^O_{gj})w_{gt} \tag{39}$$

Given ρ, we can calculate the LHS directly for all g, j, t. The RHS is a linear function of w_{gt} with unknown parameters. So, regressing the term (39) on w_{gt} by industry and occupation will give us r^O_{gj} as the constant and $r^N_{gj} - r^O_{gj}$ as the slope. Given $r^T_{gj} = \alpha^T_{gj}/\alpha^T_{gj}$, and the fact that $\sum_j \alpha^T_{j,g} = 1$, we can back out all α^T_{gj} from r^T_{gj}.

We can get A^T_{gt} as an analytical function of $(\alpha^T_{gj}, P_{jt}, P_{gt}, \rho)$, assuming $\rho \neq 0$. This is because the profit maximization gives a FOC:

$$\begin{aligned}
&P_{gt}A^T_{gt}[\sum_j \alpha^T_{gj}(y^T_{gjt})^\rho]^{\frac{1-\rho}{\rho}}\alpha^T_{gj}(y^T_{gjt})^{\rho-1} = P_{jt}\\
\Rightarrow &P_{gt}A^T_{gt}[\sum_j \alpha^T_{gj}(y^T_{gjt})^\rho]^{\frac{1-\rho}{\rho}}[\alpha^T_{gj}(y^T_{gjt})^\rho]^{\frac{\rho-1}{\rho}}(\alpha^T_{gj})^{1/\rho} = P_{jt}\\
\Rightarrow &[\frac{1}{[\sum_j \alpha^T_{gj}(y^T_{gjt})^\rho]}]^{\frac{\rho-1}{\rho}}[\alpha^T_{gj}(y^T_{gjt})^\rho]^{\frac{\rho-1}{\rho}} = \frac{P_{jt}}{P_{gt}A^T_{gt}(\alpha^T_{gj})^{1/\rho}}\\
\Rightarrow &\frac{\alpha^T_{gj}(y^T_{gjt})^\rho}{[\sum_j \alpha^T_{gj}(y^T_{gjt})^\rho]} = [\frac{P_{jt}}{P_{gt}A^T_{gt}(\alpha^T_{gj})^{1/\rho}}]^{\frac{\rho}{\rho-1}}\\
\Rightarrow &1 = \sum_j [\frac{P_{jt}}{P_{gt}A^T_{gt}(\alpha^T_{gj})^{1/\rho}}]^{\frac{\rho}{\rho-1}}\\
\Rightarrow &(P_{gt}A^T_{gt})^{\frac{\rho}{\rho-1}} = \sum_j [\frac{P_{jt}}{(\alpha^T_{gj})^{1/\rho}}]^{\frac{\rho}{\rho-1}} \qquad (40)
\end{aligned}$$

This last step gives A^T_{gt} as a function of $(\alpha^T_{gj}, p_{jt}, p_{gt})$, once we have identified all the alphas and over-time changes in p_{jt}, we identify the over-time changes in each A^T_{gt} that $y^T_{gjt} > 0$. For industry g where tech T was not adopted at time t, (40) gives the upper bound of A^T_{gt}. We can get the size of A^O_{gt} relative to A^N_{gt}. The absolute scale of A^T_t is meaningless because it's just the inverse of the scale of y^T_{gj}.

Finally, industry demand trends can be identified from observed quantities and prices of all the goods. It doesn't rely on ρ or w_{gt}. The CES utility function means that

$$\frac{1}{\sigma-1}\ln[\sum_g B_g t Q_{gt}^{\frac{\sigma-1}{\sigma}}] + \ln B_{gt} = \frac{1}{\sigma}\ln Q_{gt} + \ln P_{gt} + \ln\lambda_t \tag{41}$$

where λ_t is a Lagrangian multiplier in utility maximization. Taking differences between industry g and industry 1 within t, we get the relative trends of B_{gt}:

$$\ln B_{gt} - \ln B_{1t} = \frac{1}{\sigma}(\ln Q_{gt} - \ln Q_{1t}) + \ln(P_{gt} - P_{1t}) \tag{42}$$

σ is unknown. Industry-level prices and outputs can be obtained from the ONS.[30] We can estimate σ by assuming $\ln B_{gt} - \ln B_{2t}$ follows a time polynomial and regressing relative outputs on relative prices. We get $\hat{\sigma} = 0.16$. The absolute level of all B_{gt} is not identified, nor is it necessary because the model features Constant Returns to Scale. To impute $\ln B_{gt}$, we use our own production function to impute industry output rather than directly use the ONS measures. This is because my model does not include capital explicitly, the industry output based on observed employment and estimated production parameters in my model will be lower than actual output in capital-intensive industries. To be internally consistent, we calculate industry output from the production function, then combined with observed industry prices and $\hat{\sigma}$, equation (42) gives the relative demand trends.

4 Sources of moments of data

4.1 Occupational employment and wages

The main data source for occupational employment and wages is the UK Labour Force Survey. This a representative quarterly survey of households in the UK, focusing on education and work-related topics. It is similar in nature to the US Current Population Survey (CPS). I have used the UK LFS data from the first quarter of 1993 to the last quarter of 2017. The main estimation is restricted to the period 1997-2015, because a key dataset for technology proxy is only available over that period.

Occupation in the LFS is based on the Standard Occupational Classification of that decade: SOC1990 up to 2000, SOC2000 for LFS2001-2010, and SOC2010

[30]Source: GDP output approach low-level aggregates from the ONS website .

from 2011 onwards. There are 300+ occupations within each SOC classification. When I bring the model to data, occupation will be based on the 9 major groups under SOC2000. The occupations are : 1, managerial, 2 professional, 3 technician, 4 admin, 5 skilled trade, 6 personal services, 7 customer services, 8 production and machine operatives, and 9 elementary.[31] I construct a probabilistic mapping from SOC1990 to SOC2000 on the basis of a subsample of LFS observations linked between LFS2000Q4 and LFS2001Q2, who were in the same job and hence reported SOC1990 and SOC2000 in those two quarters. The mapping takes into account 3digit SOC1990 (300+ occupations) and individual's gender and education. On the other hand, SOC2010 is mapped to SOC2000 using the transition matrix from the Office for National Statistics.

Industry is a slight aggregation from SIC80 divisions (in the LFS until 2008) and SIC92 sections (since 2009). To ensure consistency over time and across datasets, I group industries to 7 categories: 1) agriculture, mining, energy and water supply, let's call it natural resources thereafter; 2) manufacturing; 3) construction; 4) wholesale, retail, hotel and catering; 5) transport, storage, and communication; 6) finance, real estate and business activities; 7) all other services including government administration, health, education, social and other services.

For occupational employment y_{gjt}, we add up all the actual weekly hours in the relevant cell.

For task price p_{jt}, we run a log wage regression every year on occupation dummies, gender-age interactions and detailed education dummies. We add the observed mean log wage in the reference occupation to the coefficient estimates on occupation dummies. It's possible that the resulting p_{jt} is still contaminated by unobserved compositional changes. I could estimate task prices from the New Earnings Survey Panel Data by allowing individual fixed effect. This however will have to wait till the next revision of the paper because doing analysis in the secure data lab and taking results out takes time.

The change in occupational classification causes discontinuities in the observed y_{gjt} and p_{jt}. We get rid of discontinuities in the time series by the following method. 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 5 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.

[31]Elementary includes cleaners, waiters, kitchen assistants, labourers in agriculture and in construction, security guards, postal workers and so on.

Figure 5: adjusting occupational wage for classification changes

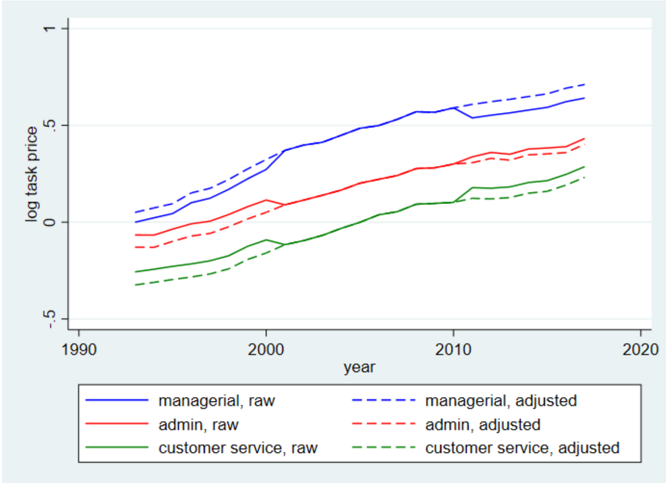

Source: UK Labour Force Survey 1993-2017.

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 the 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 BCS. [32] Figure 6 shows the two distributions of 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 scores with 20+ points but the lower range is very sparsely populated. I will summarize them to 7 points of support in each dimension during estimation. The BCS tracks people over time. We pool all the waves together and increase sample size. I have tested the hypothesis that wage returns to skills do not vary by age, and that is not rejected in most cases. So I take age effects out of wages by simply regressing log wages on age dummies; and I then get the mean log wage conditional on skills and occupation.

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

[32]At age 16, I find in the BCS data fewer than 4000 observations with arithmetic scores, and fewer than 5000 observation with raw vocabulary scores. Even the whole sample at age 16 is over 11,000.

Figure 6: Distribution of literacy and numeracy scores in BCS

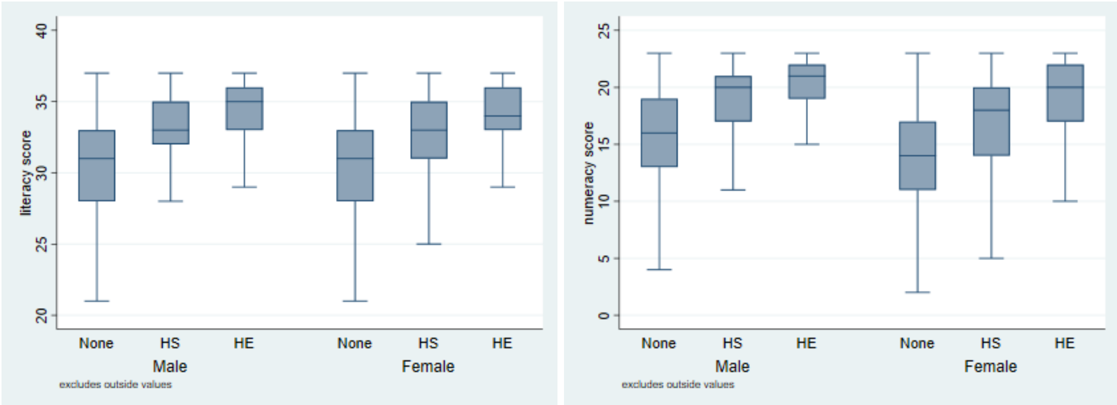

score. "HE" refers to higher education or above. "HS" refers to secondary school qualifications including A-levels, O-levels, GCSE C+ or equivalents.

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 should 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 literature (Bresnahan et al. (2002), Caroli and Van Reenen (2001)) suggests that the different aspects of changes may be complementary to each other and skill-biased.

In theory, a good proxy for the New technology should: 1) have a strong increasing trend; 2) be available at the industry-year level for a good number of years; 3) It should be positively correlated with the local proportion of graduates in the cross-section.[33]

Guided by the literature, I have considered measures of ICT capital and related tangible technology, measures about worker autonomy, organizational structure and managerial practice. Based on three criteria listed above, I have chosen proxies from two datasets: capital inputs in EUKLEM 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. It contains many variables about the organization of work.

In EUKLEMS, we observe investment and capital stock at the level of year and dozens of industries. I consider capital in 4 areas: Communication Technology, Information Technology, Software&database, and R&D. Most of the variables about investment composition and capital composition in those areas do show an increasing time trend. I have also verified that the BA proportion is positively and significantly correlated with IT capital input. Correlations with other capital inputs are mostly positive but insignificant, see table 1.

[33]This is based on the model predictions. Suppose skilled workers have comparative advantage in abstract tasks and the new technology is more intensive in abstract tasks, then an increase in skill supply will cause firms to adopt the new technology.

Table 1: Capital input proxies and the BA proportion

	cap_CT	cap_IT	cap_Soft_DB	cap_pca
propBA	0.0043	0.0301**	-0.0172	4.4629*
	(0.0036)	(0.0093)	(0.0164)	(2.1850)
propDO	-0.0057	0.0069	0.0095	-0.2704
	(0.0048)	(0.0123)	(0.0217)	(2.8928)
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, controlling for industry dummies and year dummies. 'propBA' is the proportion of people with Bachelor degree or more. 'propDO' is the proportion of people without GCSE grade C+ or equivalent.

In the BSS, there are dozens of questions about technology and organizational/managerial practices. I summarize the data to the level of industry-region-year. I run two regressions to determine which variables are reasonable proxies. First, I regress each variable on the BA proportion allowing for year dummies, industry dummies, region dummies. Second, I regress each variable on year (linearly) during the period 1992+, allowing for industry dummies and region dummies. I select variables for which the BA proportion is significant in the first regression and for which year is significant in the second regression, and the two signs must be the same. There are 8 such variables: 'whether job involves use of computerised or automated equipment', 'which is important in determining how hard works -clients/customers', 'which is important in determining how hard works -fellow workers or colleagues[34], '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 your workplace', 'i am willing to work harder in order to help this organisation succeed', and 'In your workplace, what proportion of employees work with computerised or automated equipment?'.

Figure 7 shows the aggregate trend in these 8 variables. They are mostly available for 5-6 waves. Not all of them increase steadily over time. Given the trends, I keep the following 5: two PC measures, appraisal, and two learning measures. Table 4.3 shows that all these 5 proxies are very positively and significantly correlated the local BA proportion.

Given a range of proxy measures $z^m_{gt}, 1 \leq m \leq M$, 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.

$$z^m_{gt} = \zeta^m_g + \psi^m_g w_{gt} + \epsilon^m_{gt} \tag{43}$$

The constant and the slope coefficient is specific to measure m and industry g. Because w_{gt} is unobserved, w_{gt} is only identified up to affine transformation. We have 4 measures of capital composition from 1997 to 2015 annually and 5 measures from 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 minimize the sum of squared ϵ^m_{gt}, they are equally important. Then, I do an affine transformation of w_{gt} to equal 0 in 1997 and 1 in 2015. Finally, I smooth each time series with a cubic spline and constrain the value to be in the [0,1] range. Figure 8 shows

[34]other answers include machine, boss, own discretion, pay, appraisals, and none of these

Figure 7: technology proxies in BSS, time trend

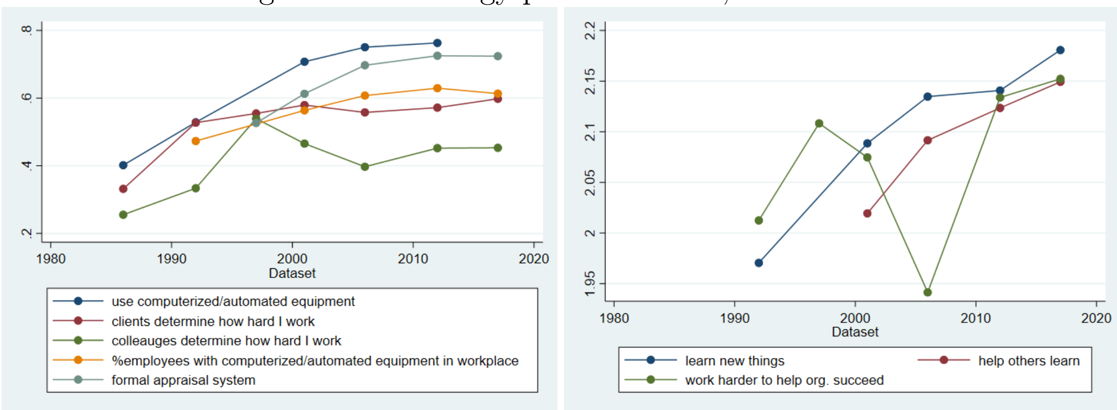

Note: I will use the following 5 to estimate w_{gt}: two PC measures, appraisal, and two learning measures.

Table 2: proxies in BSS, correlation with BA proportion

	useauto	bnewthin	bhelpoth	E_propcom	E_eapprais
BAprop	0.3276***	0.4234***	0.3081*	0.2733***	0.2000**
	(0.0707)	(0.0929)	(0.1244)	(0.0443)	(0.0702)
Observations	348	390	312	390	389

Standard errors in parentheses

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

Note: 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. 'useauto' is binary on 'whether job involves use of computerised or automated equipment'. 'bnewthin' is the reported agreement with the statement 'my job requires that i keep learning new things', it has range 0-3, higher value for more agreement with the statement. 'bhelpoth' is similar, for the statement 'my job requires that i help my colleagues to learn new things'. 'E_propcom' is the mean of 'In your workplace, what proportion of employees work with computerised or automated equipment?'. 'E_eapprais' is binary for 'do you have a formal appraisal system at your workplace'.

the resulting time series for all the industries. Note that the levels of w_{gt} are not comparable between industries, because the affine transformation is industry-specific.

Figure 8: estimated w_{gt} from 9 proxies measures

Note: estimated w_{gt}under the assumption that it follows 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 within industry.

5 Corroborative evidence

In this section, we empirically assess 2 implications of the model. First, increasing supply of skills may have little impact on occupational wages. Second, changes in the skill supply predict occupational shifts at the local level.

5.1 do occupational wages respond to supply shifts

The key difference between my model and standard models in the SBTC and RBTC literature is that technical change in my model responds to supply shocks. This has different implications for the relationship between wages and employment.

Recall equation (23):

$$\ln(\frac{p_{jt}}{p_{1t}}) = (\rho - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + (1 - \rho) \ln[(1 - w_{gt}) r^O_{gj} + w_{gt} r^N_{gj}]$$

This is similar to the SBTC formulation (44), in that both specify log price ratio equals $\rho-1$ times log quantity ratio plus a term for technical change. It represents a demand-side relationship, where the last term captures technology-driven shifts of the demand curve.

$$\ln(\frac{p_{jt}}{p_{1t}}) = (\rho - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + \theta_{gjt} \tag{44}$$

In my model, the last term is a weighted average between the old and new technologies where the weight w_{gt} is endogenous. In the exogenous SBTC literature, the last term would be an exogenous trend and it is usually approximated by some polynomial of time. The trend is linear in Katz and Murphy (1992) and Card and Lemieux (2001).

In both (23) and (44), the coefficient on log quantity ratio is -1 over the elasticity of substitution.[35] Goos et al. (2014) estimated the substitution elasticity between tasks to be 0.9, which would mean a coefficient of -1.1 in front of log quantity ratio.[36] In the case of exogenous TBTC, if one attempts to regress log price ratio on log task quantity ratio, the last term would be approximated by task-specific time polynomials. I do exactly this later, and instrument the log task quantity ratio with a shift-share IV.

My model implies that w_{gt} will respond to changes in skill supplies and task prices. As explained in section 3.1, if the supply-side shift happens to fall into the cone of diversification, the task prices will stay constant while w_{gt} adjusts to equalize demand and supply. More generally, the endogenous technological shift will tend to offset exogenous shocks on the supply side, so the estimate will be biased towards zero. Intuitively, imagine a supply shock tends to increase the professional employment and reduce professional wage. This is a case of a positive change in the employment ratio. As the new technology is more intensive in professional task $r^N > r^O$, the lower professional wage will cause a shift to the new technology: w_{gt} increases. The whole term $(1 - \rho) \log[(1 - w_{gt})r^O + w_{gt}r^N]$ increases. In other words, the correlation between the employment ratio and the omitted variable is positive, which leads to a positive bias in the coefficient estimate. As the true $\rho - 1$ must be negative, this means the empirical estimate will be biased towards zero.

Now let's put the hypothesis of exogenous technical change to test.

Firstly, we can approximate the tech term in (23) by j-g-specific time polynomials and run OLS. Estimating it separately by industry, I find the coefficient on log quantity ratio to be almost always near zero and insignificant. Table 5.1 reports the results based on LFS1993-2016. The near-zero coefficient estimate comes from a lack of correlation between occupational wage change and employment change in the data. Figure 9 plots annualized changes in wages and employment at the level of 9 occupations and 7 industries. Over each period (1993-2000,2001-08,2009-2017)[37], there is substantial movements in the employment dimension and much smaller movements in wages. And there is no obvious correlation between the two dimensions.

Second, we can instrument for the log quantity ratio. One possibility is that omitted demand shock correlates positively with both quantity and wage, so my estimate of $(\rho - 1)$ includes a positive bias. To address this, I construct shift-share style IVs for the task ratio, using shifting demographic composition of the population (defined by education-gender-age) and historical mappings from each demographic group to tasks. Using LFS 1993-2016, I find that the 2SLS

[35]In Katz and Murphy (1992) the coefficient is estimated to be -0.7 (implying an elasticity of 1.4). In Card and Lemieux (2001), the substitution elasticity between college and high-school labour equivalents is estimated to be in the 2-2.5 range. But those estimates are not really comparable to mine because they differentiate labour by education, whereas I do by occupation.

[36]Their estimate did not come from such a regression.

[37]The periods are cut to avoid classification changes in occupation and industry.

Table 3: testing the exogenous TBTC case by OLS

	natural resources	manufacturing	construction	trade
ln y_gjt/y_g1t	0.0233	0.2532***	0.0305	0.2552***
	(0.0384)	(0.0451)	(0.0473)	(0.0529)
Observations	200	200	200	200
	transport, infomation	finance, business serv	other services	
ln y_gjt/y_g1t	0.0283	0.0189	-0.0986**	
	(0.0416)	(0.0514)	(0.0322)	
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 column headings refer to industry division (SIC80). The "natural" category groups together agriculture, mining, energy and water. In each regression (separately by industry), we approximate the tech term by j-specific 5th-order polynomial of year. Source: LFS 1993-2016

Figure 9: annualized change in wage and total hours, at g, j level

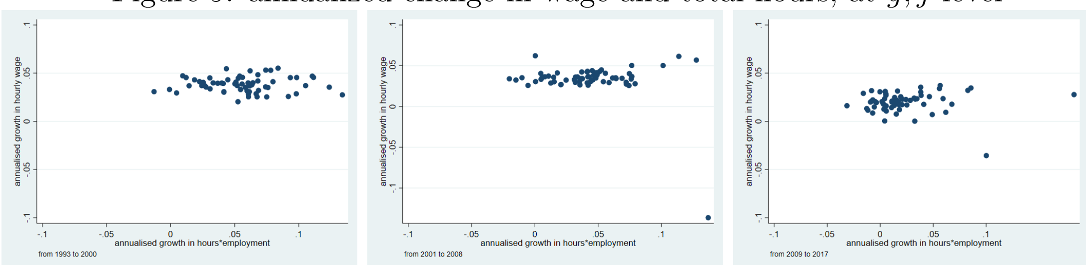

Note: based on UK Labour Force Survey 1993-2017. Each dot is an occupation-industry observation. There are 9 occupations and 7 industries.

Table 4: testing the exogenous TBTC case with 2SLS

	natural resources	manufacturing	construction	trade
ln y_gjt/y_g1t	0.8936	0.2141	-0.9201	0.5166**
	(1.3214)	(0.1275)	(0.9624)	(0.1714)
Observations	200	200	200	200
	transport, information	finance, business serv	other services	
ln y_gjt/y_g1t	0.1200	0.5330**	-0.5050	
	(0.1115)	(0.1859)	(0.2673)	
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 column headings refer to industry division (SIC80). The "natural" category groups together agriculture, mining, energy and water. In each regression (separately by industry), we approximate the tech term by j-specific 5th-order polynomial of year. The occupations are : 1, managerial, 2 professional, 3 technician, 4 admin, 5 skilled trade;6 personal services;7 customer services; 8 production and machine operatives, 9 elementary. We use a shift-share instrument at the g, j, t level, using contemporary share of demographic group and historical mapping from demographics to g, j cells. When occupation 1 is the reference occupation group, the instrument for $\ln y_{gjt}/y_{g1t}$ is $supply_{gjt}, supply_{g1t}$ Source: LFS 1993-2016.

estimates are still small and sometimes positive (See table 4). I found that the IV is reasonably strong, the standard errors are small enough to rule out $(\rho - 1) < -1$ in most industries.

In short, we find the coefficient on log quantity ratio to be near zero in the framework of exogenous technical change. This would mean tasks are perfect substitutes, which is implausible. By contrast, my model with endogenous technical change offers an explanation as to why the coefficient estimate is biased towards zero.

5.2 local skill supply change predict occupational shifts

Intuitively, there are two possible outcomes of an exogenous increase in skills supply: 1) the occupational destination conditional on skills may deteriorate ('occupational downgrading'), 2) the occupational structure shifts favourably to absorb the extra supply. At the aggregate level, the UK experienced more of 2) than 1). While there is no geography in the model, we can think of applying it to separate local labor markets, and thus obtaining the prediction that an increase in local skill supply should cause firms to switch to the technology that's more intensive in abstract tasks and therefore create more abstract jobs. Below we verify this at the regional level.

At the level of region and 3-year periods, I use the measure of skilled labour supply as in Blundell et al. (2021) (which is log total hours of degree-educated workers relative to other workers) and examine how it correlates with occupational employment shares. In terms of levels, the correlations are very strong and have the expected signs. What's more interesting to us is their correlation in changes. In table 5, I take the long difference in both skilled labour supply and occupational employment shares between 1993-95 and 2014-16, and regress the long difference on the long difference. The resulting coefficient is positive and large for managers and professional, and significantly negative for admin and personal services.

Table 5: local skill supply change and occupational shift

	managerial	professional	technician	admin	skilled trades
change in skill supply	0.0582*	0.0378	0.0105	-0.0792*	0.0233
	(0.0244)	(0.0193)	(0.0166)	(0.0351)	(0.0292)
Observations	19	19	19	19	19

	personal	sales	production	elementary
change in skill supply	-0.0520*	-0.0257	0.0325	-0.0055
	(0.0212)	(0.0227)	(0.0314)	(0.0225)
Observations	19	19	19	19

Note: the dependent variable is the change in employment share of the respective occupation between 1993-95 and 2014-16. There are 19 regions in each regression. Source: LFS 1993-2016.

At the aggregate level, we also see the increase in education accounting for a large part of the increase in abstract employment. Using the LFS (1997-2015), I decompose the change in occupational employment shares into within-gender-education-group component and between-group component. Figure 10 suggests that all of the increase in abstract employment is due to between-group, and almost all of the decline in skilled trades and operative employment is due to between-group.

6 Estimation approach and results

Currently, the model partly calibrated and partly estimated. The process is the following:

1. I guess some value for ρ and ζ.
2. Given ζ, I estimate all other supply-side parameters according to methods discussed in subsection 3.3.
3. Given any value of ρ, I estimate all other demand-side parameters $\alpha^O_{gj}, \alpha^N_{gj}, A^O_{gt}, A^N_{gt}, B_{gt}$. Their identification and estimation (conditional on ρ) are detailed in subsection 3.3
4. Given all the parameters, we solve for the equilibrium $(\mathbf{p_t}, \mathbf{w_t})$ in each year. I search for the equilibrium that is closest to the observed and satisfies all the equilibrium constraints within tolerance[38].
5. We compute a loss function as the sum of the following deviations in all years and tasks/industries: $\log P_{jt}, w_{gt}, \log P_{gt}, \log EMP_{jt}$. The last term is log employment share of occupation j in year t; we use log terms because a difference of say 1 ppt can be rather substantial relative to the observed employment share (e.g. the smallest occupation share averages 6 ppts in the period).
6. We select (ρ, ζ) based on the loss function.

[38]Demand minus supply in any occupational employment share is at most 1e-4 in absolute value. The unit costs of two technologies can differ up to 1%. This is not very big relative to the uncertainties in our parameters estimates. For example, a change of one standard deviation in r^O_{g2} (while holding other r^O_{gj} the same) would change the log unit cost by 0.01-0.04, depending on the industry g.

Figure 10: within-between decomposition of change in occupation employment shares

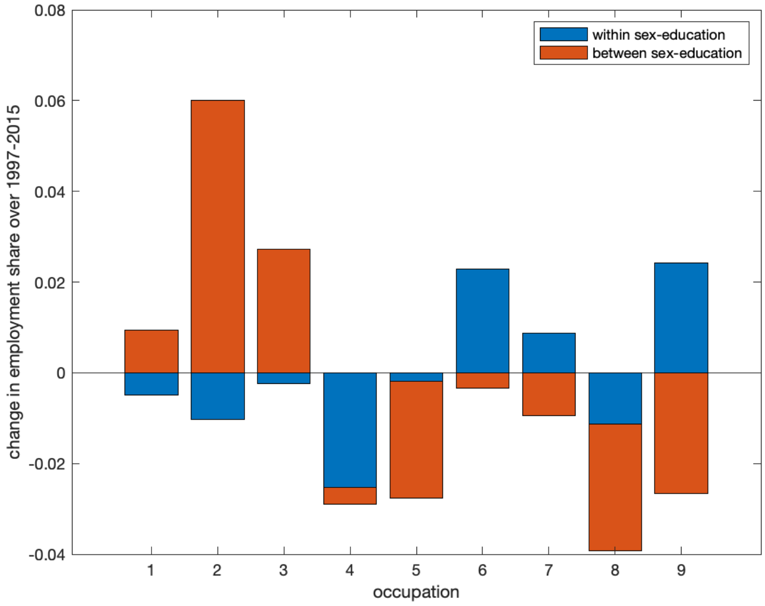

Note: occupation = 1 "managerial" 2 "professional" 3 "technician" 4 "admin" 5 "trades" 6 "personal" 7 "sales" 8 "operative" 9 "elementary".
Source: LFS 1997-2015.

In principle, we can search through the space of all parameters simultaneously to minimize the loss function; or adopt some other iterative estimation procedure. This is left for the next revision of the paper.[39]

I found $\zeta = 0.1$ minimizes the loss function at most values of ρ. Conditional on $\zeta = 0.1$, the loss function does not vary a lot across ρ if $\rho < 0.5$, and it gets very big if ρ exceeds 0.5. This means the tasks are likely complements. In what follows, we will set $\rho = -0.1$, which corresponds to Goos et al. (2014)'s 0.9 estimate of the substitution elasticity between tasks.

[39] I intend to improve the algorithm so as to find equilibrium with smaller tolerances. Currently, numerically there are sometimes multiple equilbria which can be qualitatively different to each other, so it's an issue for counterfactual analysis.

.

6.1 estimation results

First, let's compare the estimated task intensities between the two technologies. Figure 11 shows the task intensities $\alpha^O_{gj}, \alpha^N_{gj}$ in all 7 industries. Manufacturing is intensive in three tasks: managerial, skilled trades and machine operatives. The new technology is more intensive in managerial task, and less intensive in the other two manual routine tasks. This is what we expect. And this is driven by the data: within manufacturing employment has shifted away from manual routine to managerial. Meanwhile, in non-financial services, the new technology is less intensive in admin and elementary and more intensive in all 3 abstract tasks and personal service task. Again, this is true qualitatively regardless of ρ. Some patterns are common across industries. In all industries except natural resource, the new technology is more intensive in professional task and less intensive in admin task. In 5 out of 7 industries, the new technology is more intensive in managerial task. In the natural resources, the new technology compared to the old technology mainly involves a shift from operatives to skilled trades. For skilled trades and operatives, the new technology is less intensive in them than the old technology in most industries where the two tasks are sizable. Among the lower-skill task (personal service, sales and elementary), there is little evidence of the new technology being more or less intensive. While the direction of bias of technological change varies across industries, the overall pattern is that the New technology is biased against the routine tasks and towards managerial and professional tasks.

Next, we examine how the key endogenous variables in the model fit the actual trends. Recall that the only time-varying exogenous factors in the model are TFP of both technologies, industry demand, and aggregate skills distribution. And the last is based on education-gender-specific distribution and the evolving composition, so there's no free parameter in that aspect. The parameters particularly important for employment shares such as the task intensities α^T_{gj} and the occupational amenity η_j are assumed to be constant. Therefore, the design of the model does not mechanically guarantee a good fit of time trends.

Figure 12 shows the observed and predicted trends in occupational employment shares. For all of the 9 occupations, the model fit is quite good. Figure 13 shows the same for log task prices; Figure 15 shows the same for log industry prices; and Figure 14 shows the fit for technology shares w_{gt}. Note that given the partly-calibrated-partly-estimated parameters, the endogenous variables are obtained through a search for $\log \mathbf{p_t}, \mathbf{w_t}$ that is closest to the observed and subject to satisfying the equilibrium constraints.[40] This means it is expected that we get a good fit for $\log P_{jt}, w_{gt}, \forall j, g, t$. The fact that the model can capture the trends in occupational employment share movements means that the calibrated/estimated parameters are not too bad.

[40]This is because there are multiple points of $\log \mathbf{p_t}, \mathbf{w_t}$ that satisfy the equilibrium constraints within tolerance.

Figure 11: estimated task intensities in each industry

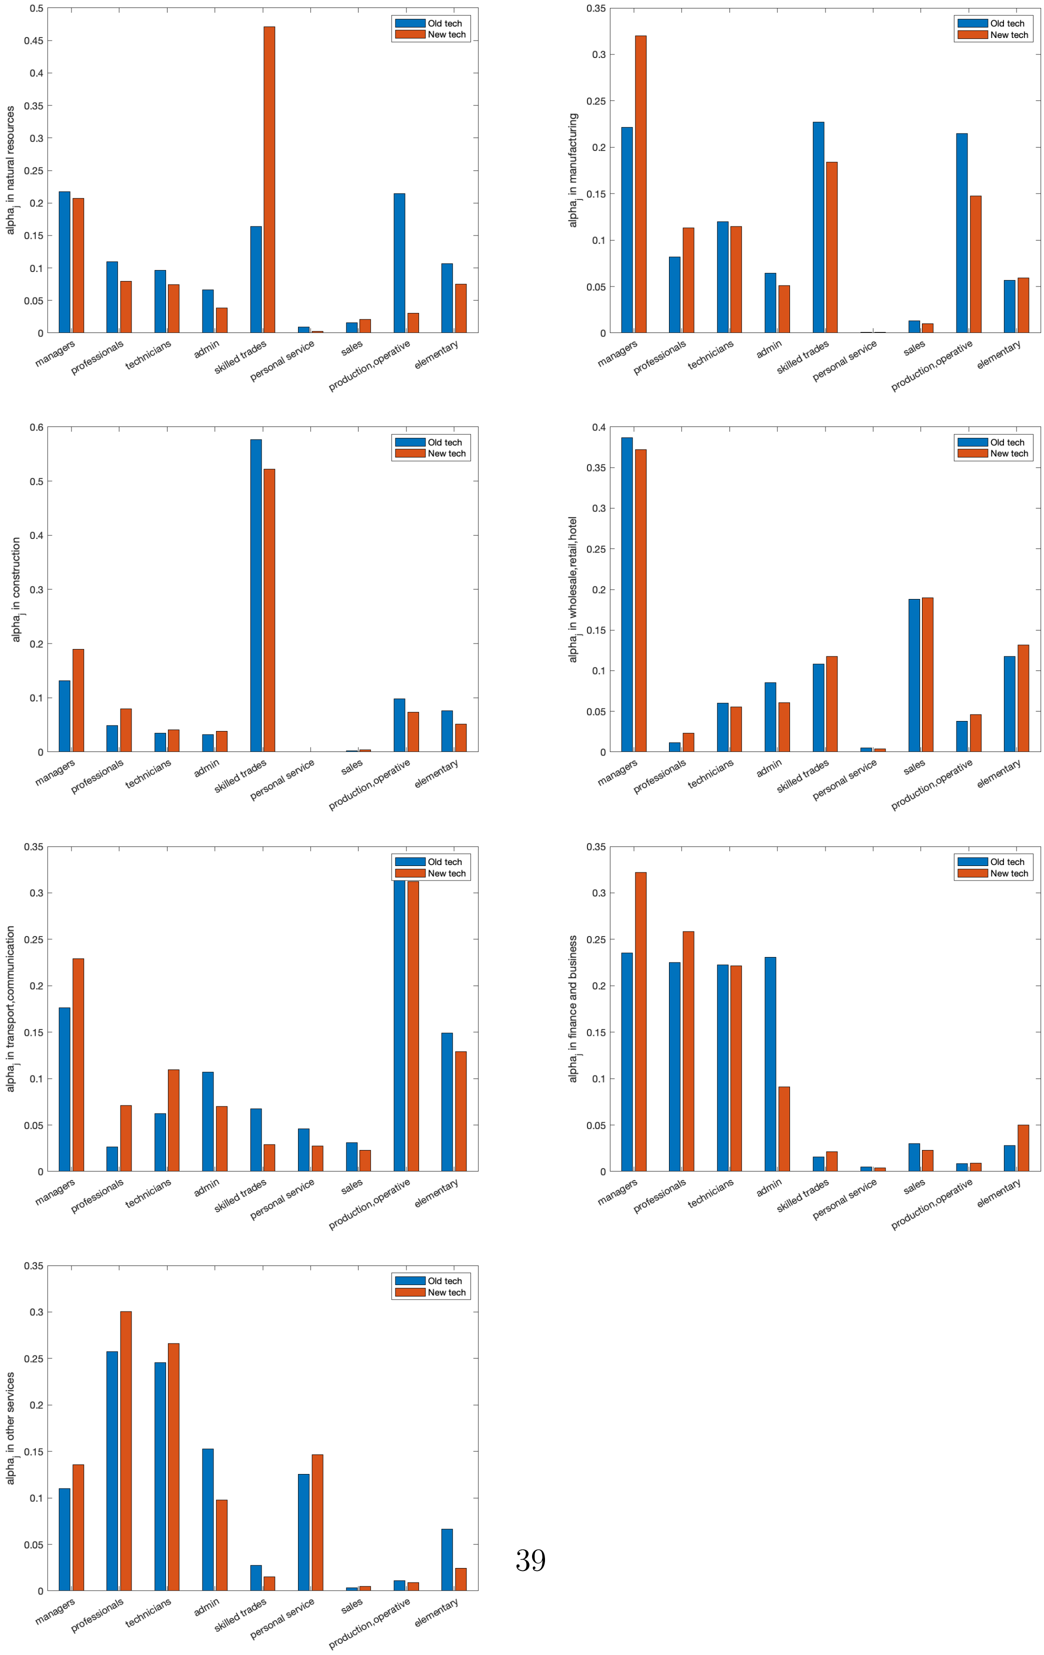

Figure 12: fit of occupation employment share

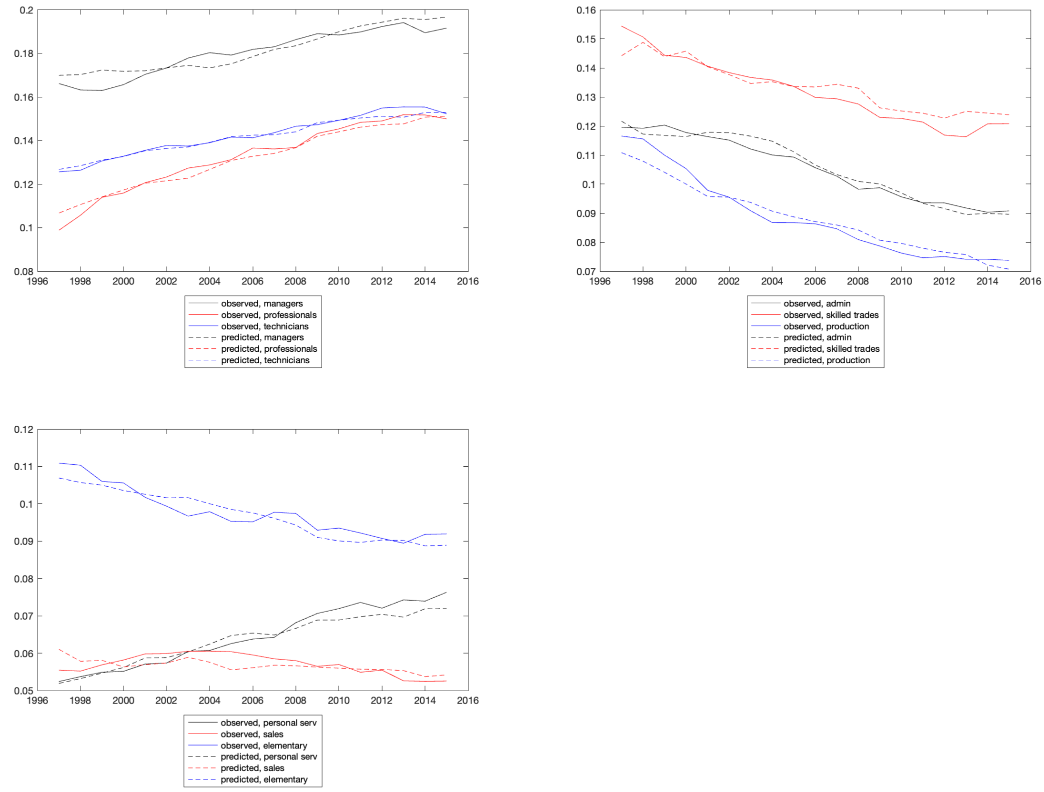

Note: The actual time trends of occupational employment shares are solid lines. The corresponding baseline predictions are dashed lines of the same colour.

Figure 13: fit of log task prices P_{jt}

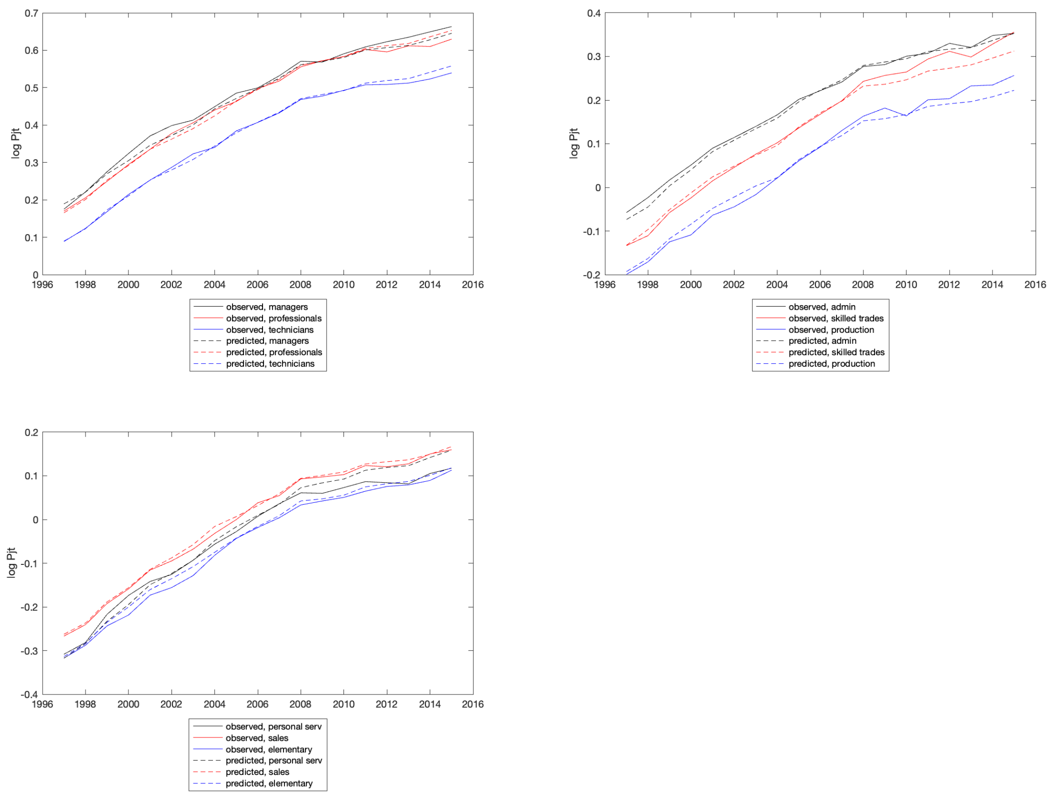

Note: The actual time trends of task prices are solid lines. The corresponding baseline predictions are dashed lines of the same colour.

Figure 14: fit of New technology's share w_{gt}

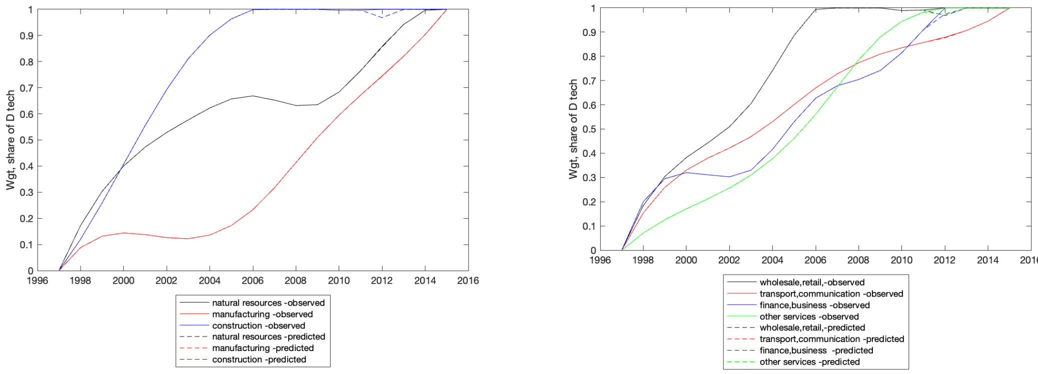

Note: The actual time trends of technology shares are solid lines. The corresponding baseline predictions are dashed lines of the same colours.

Figure 15: fit of log industry prices P_{gt}

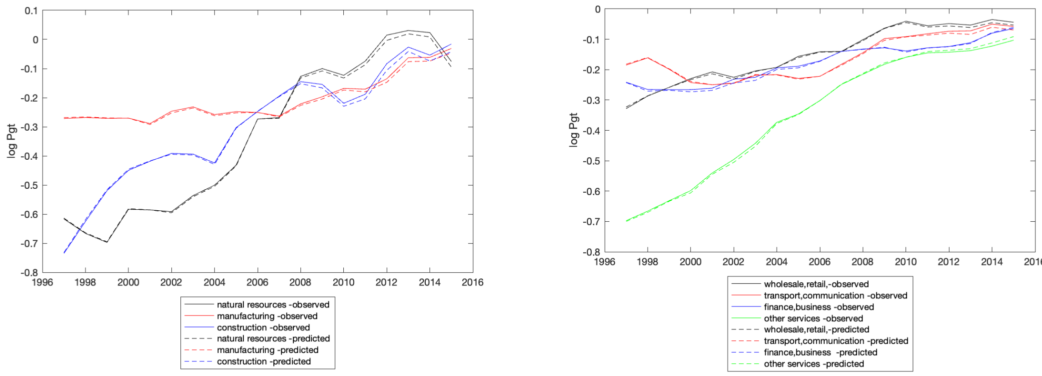

Note: The actual time trends of industry prices are solid lines. The corresponding baseline predictions are dashed lines of the same colour.

6.2 counterfactuals

The model contains three sources of exogenous factors: TFP of two technologies, industry demand, and the skills distribution. In this section, we will examine how each of them affect occupational prices and employment in the past. In future, I would also like to examine counterfactuals about Brexit-induced shift in the supply of skills and future education increases.

In each counterfactual, only one exogenous factor changes over time while others stay the same as 1997. Because numerically there are multiple equilibria, I search for $\log \mathbf{p_t}, \mathbf{w_t}$ that is closest to some benchmark subject to equilibrium constraints. I set the benchmark to the 1997 observations and interpret the result as a lower bound on the effect of shifting that factor.

Figure 16: Counterfactual: only industry demand shifts

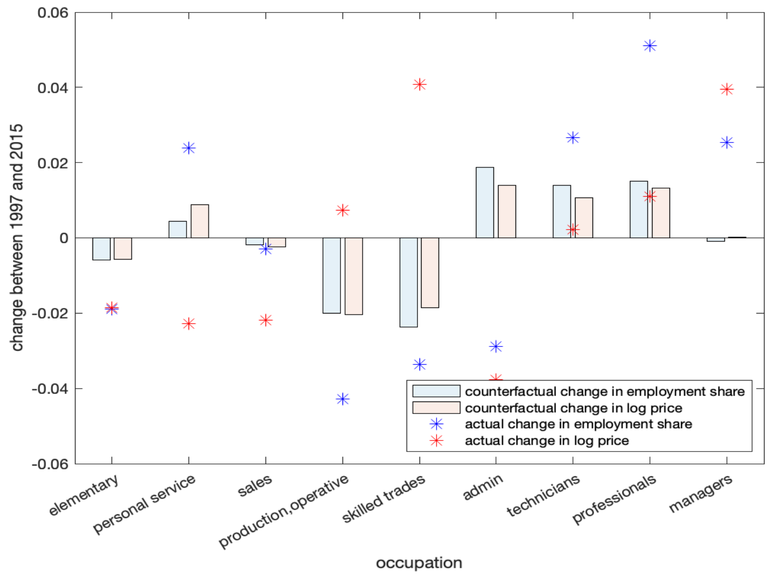

Note: using initial log P_{jt}, w_{gt} as the benchmark.

Figure 16 examines the effect of industry demand shifts. We hold TFP and skills distribution constant, and let B_{gt} follow the actual trend (imputed using estimated parameters). This counterfactual represents a shift in the demand curve and movement along the supply curve. Unsurprisingly, we see that occupational employment and prices move in the same direction.

Figure 16 also shows the actual changes as markers, so we can see that industry demand shift alone can account for half of the employemnt decline in skilled trades, machine operatives.

Figure 6.2 considers the counterfactual where the skills distribution shifts according to the baseline over 1997-2015 while TFP and industry demand are constant. In this counterfactual scenario, the equilibrium task employment would shift significantly. As the 1st subgraph shows, the supply shift alone could account for more than half of the decline in manual routine occupations and a third to a half of the increase in professional and technician occupations over the 18 year period. However, the supply shift did not contribute to the decline of admin employment. In this scenario, occupational wages change by a smaller amount. For example, the skilled trades employment share falls from 14.5% to 12.0% in the counterfactual, that is -0.17 in log terms. A demand elasticity of, say 5, would mean a wage increase of 0.03. In fact, skilled trades' wage falls relative to other occupations in this counterfactual (2nd graph in Figure 6.2). It's worth noting that occupational wages and employment do not move in the opposite direction for 7 occupations in this scenario. This counterfactual scenario is akin to a shift of the supply curve and a move along the demand curve. This illustrates the unusual implication of my model: that occupational wages do not respond to supply-side shocks along a downward-sloping demand curve.

Figure 17: Counterfactual: only skills distribution shift

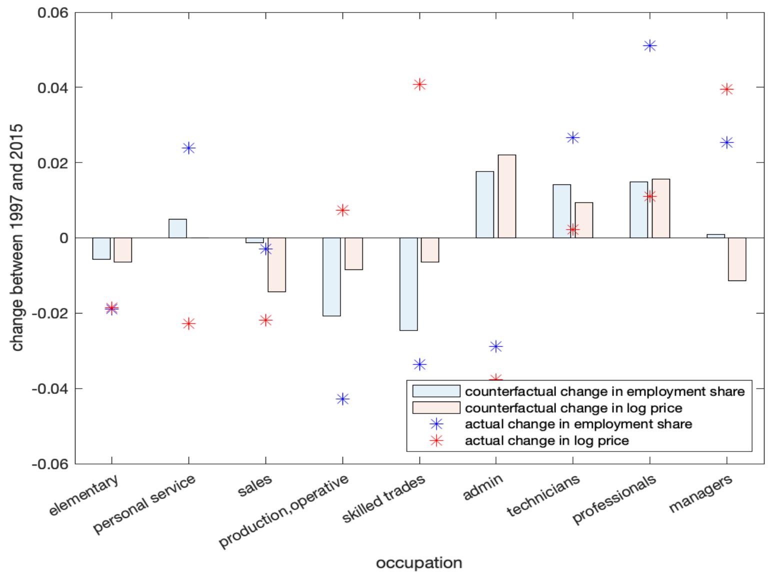

Note: using initial log P_{jt}, w_{gt} as the benchmark. For log task prices, we normalize the average change across 9 occupations to 0.

Finally, we consider the scenario with just the old technology available. In Figure 18, we let skills distribution and industry demand follow their baseline trends and suppose only the old technology was available.[41] The counterfactual incorporates industry demand shift, skills supply shift, and evolving TFP in technology C; so the only difference between the counterfactual and the baseline case is whether technology D is an option. Since the baseline prediction fits the actual changes well, the difference in Figure 18 between the counterfactual and the actual changes can be interpreted as the effects of having technology D. It suggests that having a 2nd technology option was responsible for all of the increase in managerial employment and all of the decline in admin employment.

7 Conclusion

This paper develops a multi-sector general equilibrium model of endogenous task-biased technological change that simultaneously explains three notable phenomena in the UK labour market. First, the UK has seen very strong employment growth in high-paying occupations and an even bigger decline in middle-paying occupations since the 90s. Second, changes in occupational wages are small and uncorrelated with employment changes. Third, I document the striking fact that

[41]If we held them constant and get rid of the new technology, then the baseline predictions in the 1st year would fit, because only technology C is adopted in that year.

Figure 18: Counterfactual: only the old technology was available

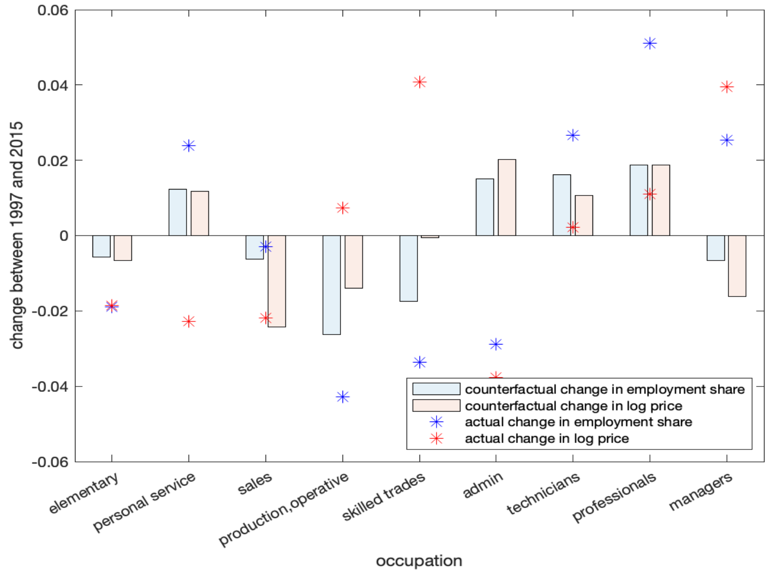

Note: in the counterfactual, we shift skill supply from its 1997 distribution to 2015 distribution, allow industry demand to follow the baseline trend, allow TFP in the old technology to follow the baseline trend, and remove the option of the new technology. For log task prices, we normalize the average change across 9 occupations to 0. We have used the initial log P_{jt}, w_{gt} as the benchmark. The results would be very similar had we used the observed log P_{jt}, w_{gt} as the benchmark.

there was relatively little occupational downgrading within education groups during a period of rapid increases in education. This is consistent with my model, and harder to explain in models with exogenous technical change. In addition to these three macro facts, I have provided regression analysis to corroborate my story and to reject the hypothesis of exogenous technical change.

This paper contributes to the polarisation literature by emphasising the endogenous nature of technology 'adoption'. Instead of an exogenous technology shock reducing the demand for routine labour, the key driving force in my explanation is a large positive shift in the supply of skills. This supply shift causes firms to adopt a new technology that's biased against routine tasks and in favour of abstract 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. The second result implies the third because any given individual's choice of occupation depends only on wages. To the extent that the skills distribution within graduates are stable, the last outcome means little occupational downgrading within graduates.

The model is estimated on UK data. While the direction of technical change varies across industries, the overall pattern is that compared to the old technology, 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 about half of the actual decline in routine manual occupations, and so can the shift in industry demand.

While the paper focuses on the UK, it provides a promising framework to study issues around occupations and education in other advanced economies except the US. Many of them share some of the key facts observed in the UK since the 90s, and are more different to the US. First, like the UK, employment growth has been strongest in high-paid occupations, compared to both the middle and the bottom. This is consistent with the new technology being more intensive in abstract tasks. Second, occupational wages did not polarise, except for the 90s in the US. And third, the US had the highest proportion of graduates in 1990 and a slower increase afterwards than many European countries. Among the European countries that saw large increases in higher education, the majority did not see a significant impact on graduates relative wage. These empirical differences between the US and the other advanced economies are intriguing, and worth further investigations with cross-country data.

Conceptually, the main point of my proposed framework is that the adoption of technology depends on current prices and skill supply. This is fundamentally different from the scenario where a new technology suddenly becomes available and it's unambiguously better than the existing one that all firms should adopt the new technology in the absence of fixed costs or frictions. That scenario might be a good enough approximation of reality in some historical episodes of technological revolution. In general, incremental changes of the technology frontier mean that there is often a meaningful choice to be made between relevant technology options. I believe many European countries are close enough to the technology frontier that their firms are in a position to choose between recent technologies, and that decision depends on prices and skill supply. In principle, the same argument of endogenous adoption should apply to the US as well; but because it's a major innovator and has experienced a smaller increase in education in the past three decades, the role of skill-supply-induced adoption of technology might be much smaller than other factors in determining the observed occupational trends.

Finally, the paper offers a data-driven approach to answer a few policy questions about the labour market. By having analytical and social skills (rather than education) as determinants of worker productivity, it allows a lot of heterogeneity within groups of labour and opens up the possibility of modelling changes in the group-specific skills distribution over time. For example, I plan to use the UK Life for Skills Survey to check whether the education-specific skill distribution has deteriorated between generations, as higher education becomes less selective. Then, the relevant data moments can be fed into the full model to investigate the effects of education expansion. The approach also makes clear that for analysing any policies about labor supply, it's important to model potential changes in the distribution of skills that matter for productivity, rather than labels like education.

Another interesting question to investigate is the effects of immigration on the aggregate labour market. Currently, immigrants in the UK are over-represented in both high-paying occupations and low-paying occupations. My next step is to use the UK Life for Skills Survey to obtain skill differences between British workers and immigrants, and examine to what extent the differences in occupation destinations are explained by skills (not just reported education level), as opposed to preferences or discrimination. Then, we can examine differences in skills distribution between European immigrants and non-EU immigrants, and ask what's the effect of forcing European immigrants to be as skilled as the non-EU ones on the UK labour market. This is a policy that the UK government can now pursue after Brexit.

References

Acemoglu, D. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *Quarterly journal of economics 113*(4), 1055–1089.

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies 69*(4), 781–809.

Acemoglu, D. (2003). Patterns of skill premia. *The Review of Economic Studies 70*(2), 199–230.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics 4*, 1043–1171.

Acemoglu, D. and P. Restrepo (2021). Tasks, automation, and the rise in us wage inequality. Technical report, National Bureau of Economic Research.

Adermon, A. and M. Gustavsson (2015). Job polarization and task-biased technological change: Evidence from sweden, 1975–2005. *The Scandinavian Journal of Economics 117*(3), 878–917.

Aghion, P., C. Antonin, S. Bunel, and X. Jaravel (2020). What are the labor and product market effects of automation? new evidence from france. Technical report.

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. 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. *American economic review 103*(5), 1553–97.

Autor, D. H. and M. J. Handel (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics 31*(S1), S59–S96.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics 118*(4), 1279–1333.

Barany, Z. L. and C. Siegel (2018, January). Job polarization and structural change. *American Economic Journal: Macroeconomics 10*(1), 57–89.

Barro, R. and J.-W. Lee (2013). A new data set of educational attainment in the world, 1950-2010. *Journal of Development Economics 104*, 184–198.

Battisti, M., C. Dustmann, and U. Schonberg (2017). Technological and organizational change and the careers of workers. Technical report, Working Paper.

Battu, H., C. R. Belfield, and P. J. Sloane (2000). How well can we measure graduate over-education and its effects? *National Institute Economic Review 171*(1), 82–93.

Beaudry, P., M. Doms, and E. Lewis (2010). Should the personal computer be considered a technological revolution? evidence from us metropolitan areas. *Journal of Political Economy 118*(5), 988–1036.

Beaudry, P., D. A. Green, and B. M. Sand (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics 34*(S1), 199–247.

Blundell, R., D. A. Green, and W. Jin (2021, 07). The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies 0*(0), 1–39. rdab034.

Böhm, M. J. (2020). The price of polarization: Estimating task prices under routine-biased technical change. *Quantitative Economics 11*(2), 761–799.

Borghans, L. and B. ter Weel (2008). Understanding the technology of computer technology diffusion: Explaining computer adoption patterns and implicaitons for the wage structure. *Journal of Income Distribution 17*(3), 37–70.

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).

Caroli, E. and J. Van Reenen (2001). Skill-biased organizational change? evidence from a panel of british and french establishments. *Quarterly journal of economics*, 1449–1492.

Dolton, P. and A. Vignoles (2000). The incidence and effects of overeducation in the uk graduate labour market. *Economics of education review 19*(2), 179–198.

Dolton, P. J. and M. A. Silles (2008). The effects of over-education on earnings in the graduate labour market. *Economics of Education Review 27*(2), 125–139.

Dustmann, C. and A. Glitz (2015). How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics 33*(3 Part 1), 711–750.

Dustmann, C., J. Ludsteck, and U. Schönberg (2009, 05). Revisiting the German Wage Structure*. *The Quarterly Journal of Economics 124*(2), 843–881.

Fernández-Macías, E. and J. Hurley (2017). Routine-biased technical change and job polarization in europe. *Socio-Economic Review 15*(3), 563–585.

Firpo, S., N. M. Fortin, and T. Lemieux (2011). Occupational tasks and changes in the wage structure. Technical report.

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.

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.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics 126*(2), 1029–1069.

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, 1992–2015. *Work and Occupations 46*(4), 441–469.

Traiberman, S. (2019, December). Occupations and import competition: Evidence from denmark. *American Economic Review 109*(12), 4260–4301.