ECONOMICS

Working Paper Series

No. 06-2022

Occupational polarisation and endogenous task-biased technical change

Wenchao Jin
University of Sussex and Institute for Fiscal Studies
Wenchao.Jin@sussex.ac.uk

Abstract: Since the 90s many developed countries have experienced job polarisation, whereby employment shifts away from middle-paying jobs and towards both higher-paid and lower-paid ones. The most popular explanation is that technological changes have been biased against routine tasks. This paper offers a complementary explanation that emphasises the increase in skill supply and the resulting adoption of technology. I exploit the large policy-driven expansion of higher education in the UK and argue that this supply-side shift has caused the adoption of routine-biased technology and thereby employment polarisation. This framework is supported by three facts observed in the UK. First, employment has shifted from the middle to the top, with not much change at the bottom of the occupation distribution. Second, there were relatively little movements in occupational wages and the pattern is not U-shaped. Third, over a period of rapidly increasing supply of graduates, occupational outcomes among graduates have been broadly stable. I build an equilibrium multi-sector model of occupational labor and fit it to 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, the shift in skills supply alone can account for between a third and two thirds of the actual decline in routine manual occupations.

JEL codes: J20, J24

Key words: job polarisation, Routine-Biased Technical Change

Occupational polarisation and endogenous task-biased technical change

Wenchao Jin*
University of Sussex and Institute for Fiscal Studies

December 2, 2022

Abstract

Since the 90s many developed countries have experienced job polarisation, whereby employment shifts away from middle-paying jobs and towards both higher-paid and lower-paid ones. The most popular explanation is that technological changes have been biased against routine tasks. This paper offers a complementary explanation that emphasises the increase in skill supply and the resulting adoption of technology. I exploit the large policy-driven expansion of higher education in the UK and argue that this supply-side shift has caused the adoption of routine-biased technology and thereby employment polarisation. This framework is supported by three facts observed in the UK. First, employment has shifted from the middle to the top, with not much change at the bottom of the occupation distribution. Second, there were relatively little movements in occupational wages and the pattern is not U-shaped. Third, over a period of rapidly increasing supply of graduates, occupational outcomes among graduates have been broadly stable. I build an equilibrium multi-sector model of occupational labor and fit it to 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, the shift in skills supply alone can account for between a third and two thirds of the actual decline in routine manual occupations.

*Email: wenchao.jin@sussex.ac.uk; 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, Lars Nesheim, Fabien Postel-Vinay, Uta Schönberg and John Van Reenen, 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 been an important factor in rising income inequality. The dominant explanation of this in the literature is Routine-Biased Technological Change (RBTC).[1] Most of the polarisation literature interpret RBTC as a consequence of increasing availability or productivity of automation equipments, or their declining costs. In other words, it's an exogenous demand shift hitting the middle-paying routine-intensive jobs. This paper offers a complementary explanation: while incorporating exogenous technical change, the emphasis here is on the increasing supply of skilled labour and the consequent switch to routine-biased technology.

From a policy perspective, supply-side policies such as increasing education are important policy levers for addressing income inequality in the long run. Given the prevalence and the scale of employment polarisation and its adverse impact on inequality and social mobility, it is surprising that few papers have examined the role of supply-side shifts in the polarisation context.[2] In theory, demand-side factors and supply-side factors could affect each other in endogenous ways: technological change may respond to supply-side shifts, while education choices may depend on expected demand shift. The UK provides a uniquely-suitable context to investigate this problem because its increasing supply of graduates was largely driven by policy.[3]

I build on the RBTC literature by allowing the *adoption* of technology to respond to skill supply shifts. The idea that firms' choice of production technology depends on the supply of skills is supported by a growing literature (Beaudry et al., 2010; Lewis, 2011; Akerman et al., 2015).[4] Compared to standard theories of RBTC, incorporating endogenous adoption of technology gives different implications for the effects of supply shifts on wages. As we will see, the UK data

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

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

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

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

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

First, the pattern of employment polarisation in the UK is essentially a shift from middle-paying occupations to high-paying ones, with very little change in low-paying ones. If we broadly classify occupations into 3 groups, the total employment share of the middle group fell by 10.5 percentage points between 1997 and 2015, while that of top group increased by 10.3 percentage points.

Second, there has been no wage polarisation in the UK since the mid-90s.[5] In fact, the movements in wages are uncorrelated with those in occupational employment in the UK. In my model, when the skills distribution shifts in a way that favours a certain task, firms will 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 the relative quantities of tasks rather than prices. Thus, it is consistent with job polarisation and a lack of wage polarisation at the same time.

The 3rd stylised fact is more striking: the huge increase in educational attainment since the early 90s has led to relatively little 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 $75 - 80\%$. 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 present some regression results that reject the hypothesis of exogenous task-biased technical change and support my model.

This paper's main contribution is to offer a different explanation of facts around polarisation, while incorporating three often-proposed mechanisms: 1) RBTC within industry, 2) between-industry demand shift, and 3) supply-side shifts. Most of 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 unspecified differences in institutions or differences on the supply side(Autor et al., 2003; Goos et al., 2014; Goos and Manning, 2007; Acemoglu and Autor, 2011). As we'll discuss below, many developed countries experienced employment polarisation without wage po-

[5]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.

[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.

larisation (Green and Sand, 2015). By contrast, the chain of events emphasised here starts with a policy-driven positive supply shift, which causes task-biased technical change, and therefore leads to the three aforementioned facts about 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 feature is new to models in the polarization literature. And it is important because it yields different predictions for how a supply-side policy would affect employment and wages. These predictions also differ from the implications of models with endogenous *innovations*Acemoglu (1998, 2002, 2003).[7] 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 second most popular explanation in the polarization literature is sectoral shifts. Intuitively, factors such as population ageing and rising Chinese imports may lead to rising demand for personal services and falling demand for manufacturing goods, which in turn reduces demand in routine occupations. To incorporate this channel, I allow 7 industries in my model. 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 shifts could explain a large part of the changes in occupational employment shares and in occupational wages in the US since the 1950s.

When it comes to supply-side shifts, this paper is among the first to quantify its contribution to job polarization. To the best of my knowledge, the only other paper that quantities the contributions of supply shifts versus demand shifts is contemporaneous work by Aseem Patel (Patel, 2022). Using French data between 1994 and 2015, Patel (2022) estimated that changes in education composition accounted for 3.1 ppts decline in the share of routine jobs, while RTBC accounted for 2.3 ppts, and the remaining 5.6 ppts decline was attributed to (unobserved) changes in preferences. There are several important differences in our models, which means we capture different channels. My model allows for exogenous between-industry demand shifts and it does not allow for preferences to change over time because they are unobserved.[8] Moreover, my model incorporates endogenous choice of

[7]Models with endogenous innovations 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 model with endogenous adoption implies a flatter short-run demand curve.

[8]Patel (2022) allows routine intensity to differ across firms (not just across industries) and this means RBTC will affect the firm-size distribution and offset the decline of routine jobs within firms. This between-firm channel is absent in my model.

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

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

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

[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 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 production function that combines occupational labor into output. For example, adopting robots in the production process means you would need more technicians and fewer production workers to produce one unit of output.

We fit the model to the UK data over 1997-2015 at the level of 9 occupations and 7 industries.[11] It can fit the observed trends pretty well. The good fit is not mechanically guaranteed by the model design, because most of the key parameters (such as preferences) do not vary over time. The estimates in most industries suggest that technological changes in the UK over 1997-2015 were 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 the shift in skills distribution alone can account for between a third and two thirds of the decline of manual routine occupations, and between a third and half of the increase in the three abstract occupations. The shift in the industry demand could account for similar magnitudes of occupational shifts.

The rest of 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 sets out the model and explains how to identify the technology trend and the model parameters. Section 4 describes the data sources, including the proxies used to impute the share of the 'New' technology. Section 5 presents correlations in UK data, which are further corroborative evidence for the model. Section 6 explains how various parameters are estimated, discusses some key estimates and the fit of the model, and conducts counterfactual analysis. Section 7 concludes.

2 Motivating facts

This section documents three phenomena in the UK labour market since the 90s.

1. Employment has shifted significantly away from middle-paying occupations towards primarily the higher end and to a lesser extent the lower end.

2. There is no clear U-shaped pattern in occupational wage changes during the period of employment polarisation.

3. The huge increase in education attainment in the UK has not led to much occupational downgrading among graduates, nor decline in the graduate wage premium.

[11] The 9 occupations are SOC2000 major occupation groups: 1, managers and senior officials, 2 professional, 3 associate professional and technical, 4 administrative and secretarial, 5 skilled trades, 6 personal services, 7 customer services, 8 process, plant and machine operatives, and 9 elementary.

[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).

2.1 Fact 1: employment polarisation

Employment polarisation refers to a 'hollowing out' along the occupation spectrum. This phenomenon has been documented extensively in the literature for the US (Acemoglu and Autor (2011), Autor and Dorn (2013), Hershbein and Kahn (2018)) as well as many other developed countries (Goos et al. (2014), Breemersch et al. (2017), Michaels et al. (2014)). It's been documented since the 1980s for the UK (Goos and Manning, 2007) and Germany (Kampelmann and Rycx, 2011) and even earlier for the US (Barany and Siegel (2018)). The phenomenon is robust to different ways of classifying and ranking occupations for both the US and the UK. When my model is brought to the UK data, occupation will be at the level of SOC2000 major occupation groups.[13] So in this section I present occupational facts at this level, too. At this level of nine occupations, the three middle-paying occupations are normally considered 'routine': 'Administrative and Secretarial Occupations', 'Skilled Trades Occupations', and 'Process, Plant and Machine Operatives'. The three high-paying ones will be referred to as 'abstract', and the low-paying ones as 'manual'.

Figure 1 shows that each of the three routine occupations saw a very substantial decline in employment share. Over 1997-2015 (the period for which my model will be estimated), the total employment share of the 3 routine occupations fell from 39.1% to 28.5%: a decline of 10.6%. Meanwhile, each of the three abstract occupations grew substantially. In particular, professional occupations grew from 9.9% of aggregate employment to 15%. Together, the abstract employment share grew from 39.1% to 49.4% over the sample period: an increase of 10.3%. Among the manual occupations, there is some decline in elementary occupations[14], which is more than compensated by the increase in personal services (such as care assistants). Overall, the pattern of employment polarisation in the UK is more of a shift of employment from the middle to the top, with very little change at the bottom.

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

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

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

Figure 1: Employment shares by occupation

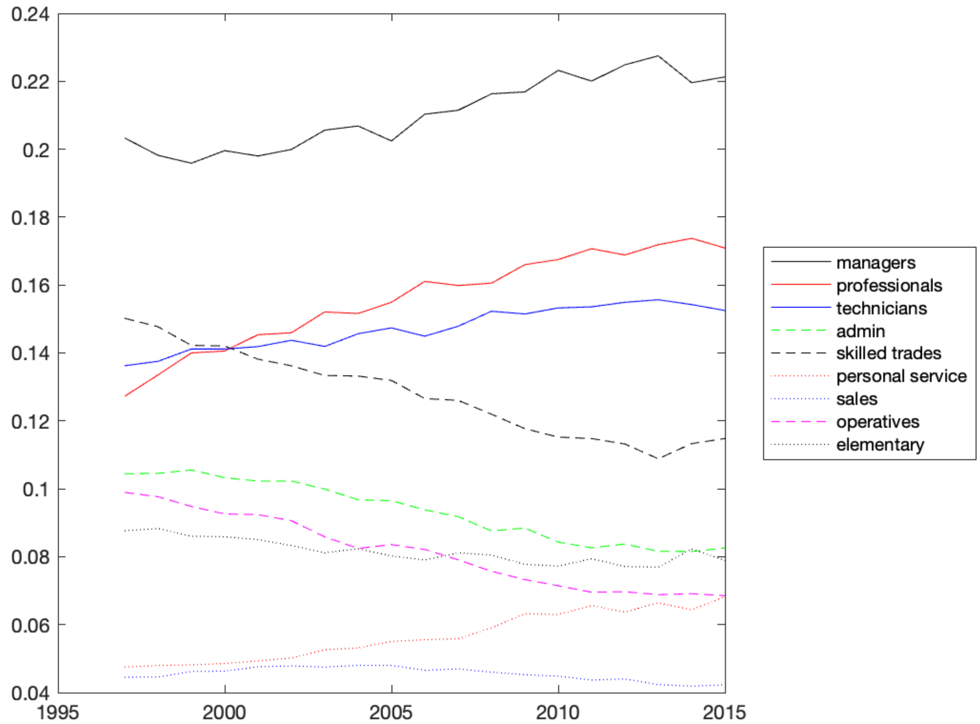

Note: the 9 occupations are major occupation groups under SOC2000. See section 4 for how we adjusted for discontinuities in SOC over 2000-01 and 2010-11.

For example, 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.[15] Murphy and Oesch (2018) looked at Ireland and Switzerland over 1970-2010 and found 'occupational upgrading', and the patterns were consistent with changes coming from the supply side associated with female education and immigration. It's beyond the scope of this paper to investigate why those studies reach different conclusions. Notably, they all point to strong growth in high-paying occupations. We see in both Figure 1 and Figure 12 that the professional occupation stands out as having the strongest growth. This is an occupation in which university graduates are likely to have comparative advantage. In the framework proposed here, an increase in the supply of graduates will cause firms to adopt a technology that's more intensive in professional tasks, and therefore the professional employment share will increase. My model does not have a definitive prediction as to whether low-paying occupations should grow or decline relative to the middle. Both 'occupational polarisation' and 'occupational upgrading' could be the consequence of an increase in skills supply. The former follows if the new technology is biased against middle-skilled tasks and in favour of high-skilled tasks; while the latter follows if the new technology is biased in favour of high-skilled tasks and against low-skilled tasks.

2.2 Fact 2: no wage polarisation

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

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

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

[16] 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.

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

In other European countries, we have not seen wage polarisation since the 80s either. Before 2000, wage inequality increased 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) summarised 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 wage inequality between occupations. In figure 12, 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 explanation is that new technologies (such as automation software and robots) can increasingly replace workers in carrying out routine tasks. Thus, the middle-paying routine occupations are hit by a negative technology-induced demand shock.[18] This has a lot of intuitive appeal, and it fits the polarising trends in both employment and wages in the US in the 90s. Inspired 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).[19] The absence of relative wage declines in routine occupations outside the US does not directly support the RBTC hypothesis, but it does not in itself reject it either.

of the log wage ratio to log quantity ratio along the demand curve would be -1. 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.

[18]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 is also about a shift in the demand curve.

[19]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 specialise in tasks that were automated experienced relative wage falls.

Figure 2: Changes in log occupational wages

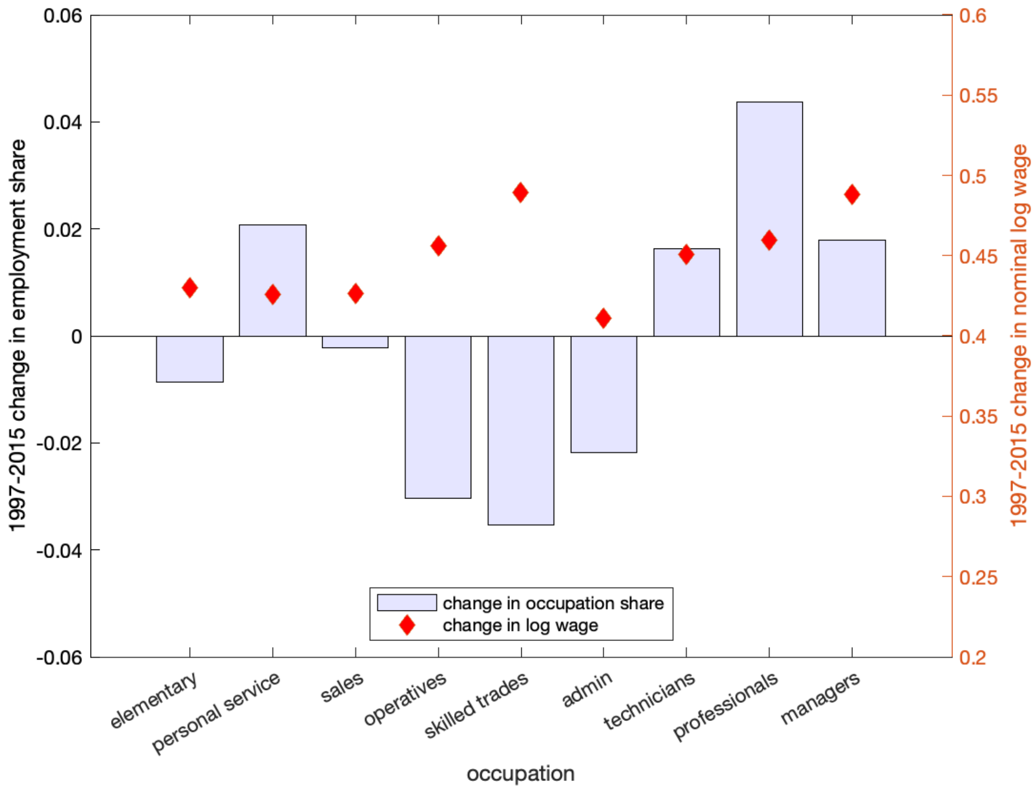

Note: in each year, we regress log wages on gender-age interactions, detailed education, and occupation dummies. This forms our 'composition-adjusted' occupational wage data log Pjt. Because the occupation classification changed in 2001 and 2011, we then fit each Pjt with a 5th-order polynomial with 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.

There are at least four reasons why exogenous RBTC could lead to substantial employment polarisation and no noticeable impact on observed wages, and they could be true simultaneously. First, the supply curve could have shifted at the same time in the same direction as the demand. Second, 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. Third, wages are sticky for institutional reasons. Fourth, observed wages are confounded by unobserved compositional changes. In particular, the unobserved compositional shift is likely to be negative in expanding occupations and positive in shrinking occupations, and therefore the observed occupational wage changes are attenuated.[20] In the

[20]Using German longitudinal data to address the composition shifts, Böhm et al. (2019) found

UK case, Cavaglia and Etheridge (2020) has used longitudinal data to estimate changes in task prices for four broad groups of occupations over 1991-2008, and found that the price of abstract tasks has increased by a statistically significant 0.126 log points, while changes in the other task prices are not significantly different from zero.

It's beyond the scope of this paper to investigate all the plausible explanations of the absence of wage polarisation. What this paper offers is a unified explanation of the facts without deviating from competitive labour markets. 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 a model with exogenous demand shifts. The correlations between occupational wages and employment ratios do not support the latter.

2.3 Fact 3: little change in graduates' occupational destinations

The proportion of graduates has more than doubled in the UK since the early 90s, but there has been no significant deterioration in graduates' relative wage or occupation destinations.

This increase in educational 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. The student number controls were raised by the government by a little every year. This resulted in a steady increase in student numbers until the 2010s. In 2012, the cap was abolished for students whose exam grades were above a certain threshold. Since 2015, the cap was abolished for all. Throughout the period, university admission was rationed by prior academic achievement. Figure 3 shows that about 20% of workers in the early 90s had higher-education qualifications, and this more than doubled over the next two decades.[21] The pace of increase is much faster than the US.

that the movement in skill price was actually positively correlated with the employment change at detailed occupation level.

[21]Here 'higher-education' qualifications include both Bachelor's degrees and other tertiary-level qualifications like nursing qualifications. This is a slightly broader definition than 'BA's in Blundell et al. (2022), because in this paper we will also need the British Cohort Studies, which only allows this broader definition of graduates.

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

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

To give a sense of magnitude, I calculate how much the share of abstract occupations needs to fall within education group if the aggregate abstract share had been constant while the education composition improves.[22] These counterfactual trends are plotted as dashed lines in Figure 3: the proportion in abstract occupations conditional on education would need to fall by about a quarter. Thus, the UK story is one where the increase of graduates was quickly absorbed through employment growth in abstract occupations. The model in section 3 will formalise this intuition: increasing education means more workers now have comparative advantage in abstract tasks; this would cause firms to switch to the abstract-task-intensive technology and create more abstract jobs.

The relatively flat trend documented in Figure 3 is robust to the classification of abstract or higher-paid occupations. Consistent with what's shown here, Salvatori (2018) found a 'very small' shift in graduate employment away from top occupations; and Green and Henseke (2016) defined 'graduate jobs' based on skills requirements of detailed occupations and found no increase in over-education following the huge increase of graduates in the UK labour force.

This lack of occupational downgrading among UK graduates 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 was basically flat over the 2000s, even though the UK saw a much faster increase of college graduates than the US.

Broadly speaking, the UK experience is more similar to many other developed countries except the US. Many developed countries have seen large increases 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

[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 3: Proportion of graduates and their occupation destination

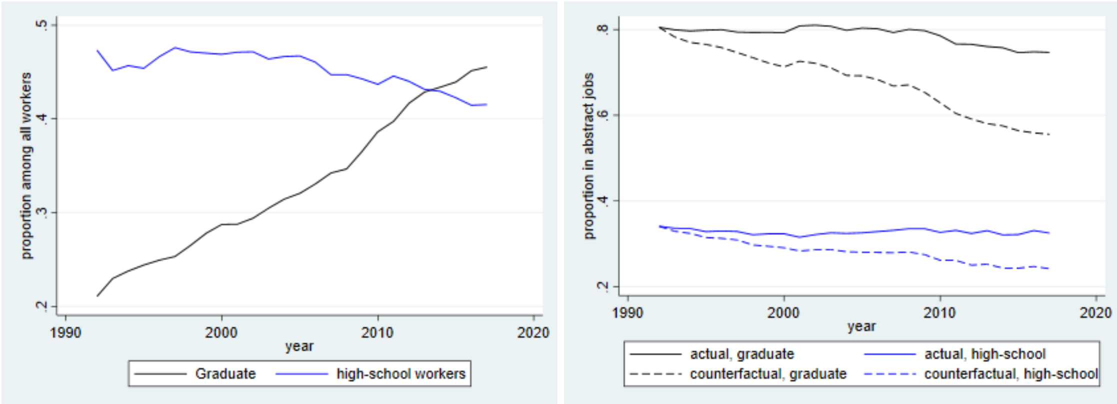

Note: graduates are people with NVQ level 4 qualifications or above. High-school workers refer to those with NVQ level 2 or 3 qualifications. First degrees are NVQ level 4. A-levels and post-16 further education qualifications are NVQ level 3. O-levels and GCSEs (grade C+) are NVQ level 2. 'Abstract'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

with complete tertiary education was already 24% in the US by 1990, when the proportion in European countries was all below 15%. This supports the view that the US has been the leader of skill-intensive technologies in general, with other developed countries closely behind. This means when their workforce's education level catches up, the latter group (including the UK) are in a position to adopt newer technologies, and this choice would depend on prices and wages. Consistent with this view, Blundell et al. (2022) shows that in 11 OECD countries which experienced substantial increase in tertiary education, there was no significant decline in graduates' relative wages in 9 of them, like the UK. Finally, it has also been documented in Green and Henseke (2021) that in 24 European countries, the share of graduates in non-graduate occupations has increased 'only modestly' from 19 to 21 percent over 2005-15.[23] All these similarities suggest that a model of endogenous adoption of technology (like the one proposed in section 3) might be more suitable for these non-US developed countries, whereas the US might need a model of endogenous innovations.

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

3 Model of endogenous adoption of task-biased technology

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

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

Each industry produces one good. Denote the goods as $g \in \{1, 2, ..G\}$. 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' technology and the 'New' technology. Firms are otherwise identical 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.

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

Y_{gt}^{T} is the output produced in industry g at time t under technology T. y_{gjt}^{T} is the amount of task j employed in industry g, using technology 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 the elasticity of substitution between tasks. ρ must be below 1. ρ is negative iff tasks are complements. A_{gt}^{T} is Total Factor Productivity of technology T in industry g at time t.

The model does not contain capital explicitly. We can think of the choice of

[24]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 also a question of how to organise all the tasks into bundles across individual workers and then to combine them by management.

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

Each technology is assumed to have constant returns to scale. We normalise $\sum_j \alpha^T_{gj} = 1, \forall g, T$.

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

$$U_t = [\sum_g B_{gt} Q_{gt}^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}} \tag{2}$$

$$Q_{gt} = Y^O_{gt} + Y^N_{gt} \tag{3}$$

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

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

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

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

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

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

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

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

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

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

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

[29]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 as a Cobb-Douglas function of capital, human capital in each occupation and intermediate inputs produced in other industries.

[30]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.

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

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

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

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

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

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

The utility that worker i gets from occupation j at time t is

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

where η_j is occupation-specific amenities; e_{ijt} follows iid Type-1 extreme value distribution, with location parameter at 0 and scale parameter ζ. Idiosyncratic preference shocks e_{ijt} mean that for any given (a_i, s_i), there is positive probability of the worker going to any occupation j. Omitting the time subscript t for simplicity, the probability of worker i choosing occupation k is simply

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

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 wages are more influential in occupation choices. Note that the unobserved heterogeneity term μ_i does not enter into occupational choice. Thus $\pi_k(i, \mathbf{p}) = \pi_k(a_i, s_i, \mathbf{p})$.

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

$$LS_j(\mathbf{p}) = \sum_i \pi_j(a_i, s_i, \mathbf{p}) y(i, j) \tag{8}$$

$$= \int\int \pi_j(a, s, \mathbf{p}) y(a, s, j) f(a, s) da ds \tag{9}$$

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

Thus, the only relevant unknowns on the supply side are 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_i, s_i), or the returns to skills β_{aj}, β_{sj}.[31]

On the demand side, we do not observe y^O_{gjt}, y^N_{gjt} separately as opposed to $EMP_{gjt} = y^O_{gjt} + y^N_{gjt}$. After some algebraic manipulation, we obtain a demand-side prediction about the relationship between task price ratio and the observable task quantity ratio.

$$\ln(\frac{p_{jt}}{p_{1t}}) = (\rho - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + (1 - \rho) \ln[(1 - w_{gt}) r^O_{gj} + w_{gt} r^N_{gj}] \tag{10}$$

where w_{gt} is the share of 'New' technology in industry g at time t; $w_{gt} = y^N_{g1t}/(y^O_{g1t} + y^N_{g1t})$. And $r^T_{gj}, T \in \{O, N\}$ are functions of parameters (α^T_{gj}, ρ). See Appendix A.1 for the derivation.

Equation (10) 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 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). Those papers categorise labor input by education, whereas here it's

[31] For counterfactual analysis, we do need assumptions about how $y(a, s, j)$ would change under the counterfactual. For example, I currently assume that it does not depend on the counterfactual policy. This would be the case if the distribution of μ_i conditional on (a,s) does not depend on the counterfactual policy.

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

3.1 Equilibrium characteristics and effect of a supply-side shift

I define the equilibrium as log task prices ($\log \mathbf{p_t} = \{\log p_{1t}, ...p_{Jt}\}$) 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 can be adopted if their unit costs are equal. Here $\omega_{gt} = Y^N_{gt}/(Y^N_{gt} + Y^O_{gt})$, the share of output produced by the new technology.[32]

In Appendix A.3, we derive the following condition which makes firms indifferent between the two technologies in industry g.

$$\sum_j [(\alpha^N_{gj})^{\frac{1}{1-\rho}} - (\frac{A^N_{gt}}{A^O_{gt}})^{\frac{\rho}{\rho-1}} (\alpha^O_{gj})^{\frac{1}{1-\rho}}] p_{jt}^{\frac{\rho}{\rho-1}} = 0 \tag{11}$$

This equation is linear in $p_{jt}^{\frac{\rho}{\rho-1}}$. If the TFP ratio is very far from 1, there might be no prices that can satisfy this condition. In that case, one technology will dominate in that industry. When the TFP ratios are not extreme, there are likely infinitely many points in the $(p_{jt} > 0, 1 \leq j \leq 9)$ space that would equalise the unit costs between the two technologies in all 7 industries.

Denote δ^T_{jg} as unit input, that is, the amount of task j required by tech T to produce one unit of output in industry g. Note it's a function of all task prices $\mathbf{p_t}$.

Given all task prices $\mathbf{p_t}$, the demand for task j is

$$\begin{aligned} &\sum_g [\delta^N_{jg}(\mathbf{p_t}) Q_g(\mathbf{p_t}) \omega_{gt} + \delta^O_{jg}(\mathbf{p_t}) Q_g(\mathbf{p_t})(1 - \omega_{gt})] \\ &= \sum_g (\delta^N_{jg}(\mathbf{p_t}) - \delta^O_{jg}(\mathbf{p_t})) Q_g(\mathbf{p_t}) \omega_{gt} + \sum_g \delta^O_{jg}(\mathbf{p_t}) Q_g(\mathbf{p_t}) \end{aligned} \tag{12}$$

Industry output Q_g is a function of $\mathbf{p_t}$ through industry goods prices. It does not depend on $\mathbf{w_t}$.

(12) shows that task demand is not uniquely pinned down by task prices. Instead, movements in $0 \leq \omega_{gt} \leq 1$ allows task demand to move within the cone of diversification. The cone of diversification has as many dimensions as the number

[32] ω_{gt} is 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.

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

Market clearing requires:

$$\sum_g (\delta^N_{jg}(\mathbf{p_t}) - \delta^O_{jg}(\mathbf{p_t}))Q_g(\mathbf{p_t})\omega_{gt} + \sum_g \delta^O_{jg}(\mathbf{p_t})Q_g(\mathbf{p_t}) - LS_j(\mathbf{p_t}) = 0 \qquad (13)$$

where task supply $LS_j(.)$ follows (9).

Given all task prices $\mathbf{p_t}$, these market-clearing constraints are a system of 9 linear equations: they are linear in the 7-element vector $\omega_{gt}, 1 \leq g \leq 7$.

When some supply-side shock shifts the supply curve (for example if the density $f(a, s)$ in (9) changes), it's possible that a change in ω_{gt} will clear the markets without any change in task prices $\mathbf{p_t}$. 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.[33] Recall that individuals' occupational choice probabilities are functions of their 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 much (Figure 3). The small amount of occupational downgrading observed within education groups could be interpreted as the education-specific distribution of skills having deteriorated slightly. In short, through technology shifts, an increase in the supply of skills can leave the task prices unchanged, and the occupation destinations conditional on skills unchanged.

3.2 Identification of technology shares

We don't observe technology share 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 (10). 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} :

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

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

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 have 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. Thus, the observed change in z_{gt} at the industry level reveals the shift towards the New technology within this industry.

$$z_{gt}=(1-\tilde{w}_{gt})z_O+\tilde{w}_{gt}z_N \tag{14}$$

where $\tilde{w}_{gt}$ is the scale of new technology adopters relative to the entire industry.

In practice, we will use several measures of z. We observe z over time and at the industry level. 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 (14). In section 4.3, we will assume a latent factor model to impute w_{gt}.

3.3 Identification of model parameters

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

On the supply side, the unknowns are: η_j, the utility for working in occupation 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 to estimate other supply-side parameters such as the returns to skills. The reason was explained around equation (9).

η_j is the preference for working in occupation j, and we normalise $\eta_1=0$. The higher η_j, the more people will select into occupation 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. By contrast, we have fixed η_j, so that no changes in employment will be attributed to unobservable preference shifts. Empirically, I search for η_j to match the observed employment shares in 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.

The joint skill distribution comes from the numeracy score and the literacy score in the British Cohort Studies (BCS), measured at age 34. They are summarised to 7 points of support in each dimension.[34] 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.[35] We obtain the distribution from the BCS for each gender-education, get gender-education weights from the Labour Force Survey for each year, and aggregate up. Thus, the shift in skills distribution over time comes from the changing composition of gender and education in the UK workforce.

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 them by p_{jt} gives us the expected task output conditional on skills $y(a, s, j)$.

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

I calibrate $\rho = -0.1$, which corresponds to Goos et al. (2014)'s estimate of the substitution elasticity between tasks at 0.9.

Recall equation (10):

$$\ln(\frac{p_{gjt}}{p_{g1t}}) - (\rho - 1) \ln \frac{y_{gjt}^O + y_{gjt}^N}{y_{g1t}^O + y_{g1t}^N} = (1 - w_{gt}) r_{gj}^O + w_{gt} r_{gj}^N \tag{15}$$

$$= r_{gj}^O + (r_{gj}^N - r_{gj}^O) w_{gt} \tag{16}$$

Given ρ, we can calculate the LHS of (16) directly for all g, j, t. The RHS is a linear function of w_{gt} with unknown parameters. So, regressing the term on w_{gt}

[34] 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.

[35] 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.

by industry and occupation will give us r^O_{gj} as the constant and $r^N_{gj} - r^O_{gj}$ as the slope. Given $r^T_{gj} = (\alpha^T_{gj}/\alpha^T_{gj})^{1/(1-\rho)}$, and that $\sum_j \alpha^T_{j,g} = 1$, we can back out all α^T_{gj} from r^T_{gj}.

A^T_{gt} can be identified using the equation below. This equation comes from the F.O.C in firm's profit maximisation. Its derivation is in Appendix A.4.

$$(p_{gt}A^T_{gt})^{\frac{\rho}{\rho-1}} = \sum_j [\frac{p_{jt}}{(\alpha^T_{gj})^{1/\rho}}]^{\frac{\rho}{\rho-1}} \tag{17}$$

This equation gives A^T_{gt} as a function of $(\alpha^T_{gj}, p_{jt}, p_{gt}), \forall j$. 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} such that $y^T_{gjt} > 0$. For industry g where tech T was not adopted at time t, (17) gives the upper bound of A^T_{gt}.[36] We identify 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}. Given the CES utility function, the relative trends of B_{gt} are:

$$\ln B_{gt} - \ln B_{1t} = \frac{1}{\sigma}(\ln Q_{gt} - \ln Q_{1t}) + \ln(p_{gt} - p_{1t}) \tag{18}$$

σ is unknown. Industry-level prices and outputs can be obtained from the ONS.[37] We estimate σ by assuming $\ln B_{gt} - \ln B_{1t}$ 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 output measures. This is because my model does not include capital explicitly, the industry output based on observed employment in my model will be lower than actual output in more 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 (18) gives the relative demand trends.

4 Sources of moments of data

[36] In that case, the TFP of the dominated technology is not identified, except that it must be below the upper bound.

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

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 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 until 2000, SOC2000 over 2001-2010, and SOC2010 from 2011 onwards. There are 300+ occupations within each SOC classification. When I bring the model to data, occupations are defined as the 9 major groups under SOC2000. The occupations are: 1, managerial, 2 professional, 3 associate professional and technical, 4 administrative and secretarial, 5 skilled trade, 6 personal services, 7 customer services, 8 process, plant and machine operatives, and 9 elementary.[38] 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 3-digit SOC1990 and individual's gender and education.[39] 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 wage bills, I add up all the actual weekly hours in the relevant cell (g, j, t), and multiply it by the mean hourly wage in that cell.[40]

For task price p_{jt}, I run a log wage regression every year on occupation dummies, gender-age interactions and detailed education dummies. I add the observed mean log wage in the reference occupation to the coefficient estimates on occupation dummies. The quantity of occupational labor y_{gjt} is simply the wage bill divided by the task price.

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

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

[39]There are 300+ occupations at the 3 digit level.

[40]This is because wages are not reported for all that report hours.

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

4.2 Skills distribution

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

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

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

4.3 Technology proxies

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

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

[42] Such pairs of (a,s) consitute 0.9% of the BCS sample.

Figure 4: Distribution of literacy and numeracy scores in BCS

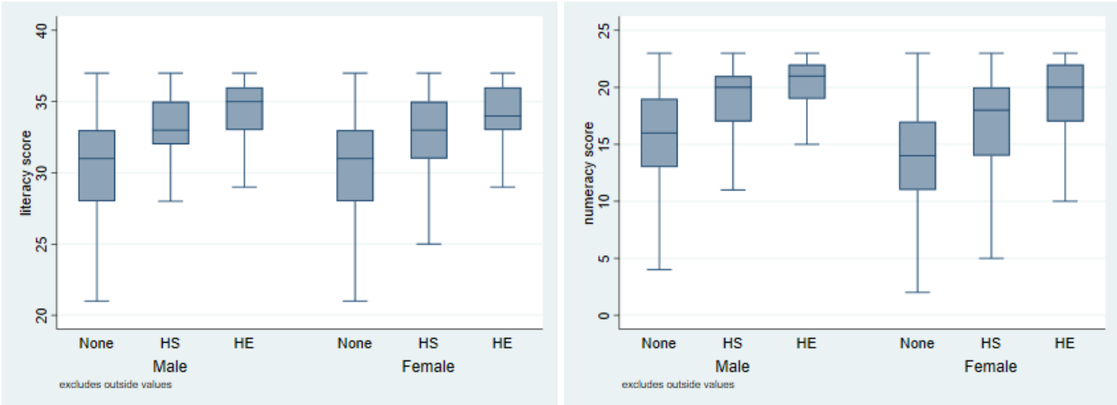

Note: from British Cohort Studies. The box edges correspond to the 25th percentile and the 75th percentile within the education and gender group. The line inside the box is the medium skill score. "HE" refers to higher education or above. "HS" refers to secondary school qualifications including A-levels,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

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

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

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

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

Table 1: Capital input composition and the graduate proportion

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

Standard errors in parentheses

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

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

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

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

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

$$z_{gt}^m = \zeta_g^m + \psi_g^m w_{gt} + \epsilon_{gt}^m \tag{19}$$

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

5 Corroborative evidence

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

Figure 5: Time trends in technology proxies in BSS

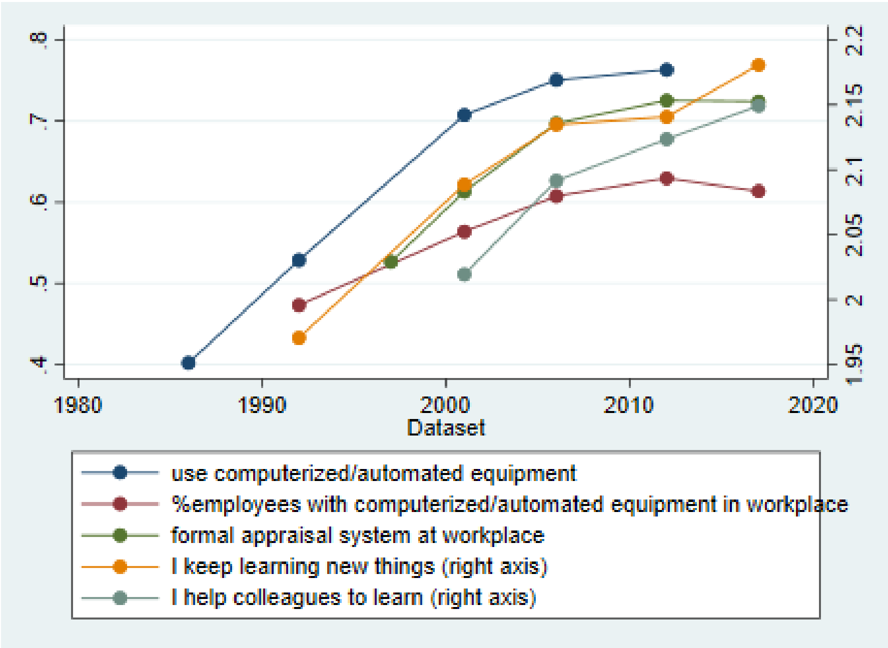

Note: the two learning measures take values between 0 to 3, 0 meaning 'strongly disagree' and 3 meaning 'strongly agree'. The other three are valued between 0 and 1.

Table 2: Proxies in BSS, correlation with graduate proportion

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

Standard errors in parentheses

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

Note: all the outcomes are aggregated to the industry-year-region level. Each regression is at the industry-year-region level, including year dummies, industry dummies, region dummies. 'own use PC' is binary on 'whether job involves use of computerised or automated equipment'. '%PC in workplace' is 'In your workplace, what proportion of employees work with computerised or automated equipment?'. 'appraisal' is binary for 'do you have a formal appraisal system at your workplace'. 'learn new thing' is the reported 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. 'help others learn' is similar, for the statement 'my job requires that i help my colleagues to learn new things'. 'BAprop' is the proportion of graduates at the industry-year-region level in the BSS.

In standard models, the demand curve is downward-sloping. Equation (20) below is a typical demand-side equation:

$$\ln(\frac{p_{gjt}}{p_{g1t}}) = (\rho - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + \theta_{gjt} \tag{20}$$

$(\rho - 1)$ is the negative reciprocal of the substitution elasticity between tasks in industry g. θ_{gjt} represents exogenous technological shift in the demand for task j in industry g. This equation (20) is similar to the canonical model of SBTC (Katz and Murphy, 1992), except that it is across occupations rather than education groups and that it's within industry g. In the exogenous SBTC literature, the last term θ_{gjt} would be an exogenous trend representing rising relative demand for skilled labour; and it is usually approximated by some polynomial of time. The coefficient $(\rho - 1)$ is estimated to be -0.7 in Katz and Murphy (1992) (implying an elasticity of 1.4).[43] Now we differentiate labour input by occupation, so the key elasticity is about the substitution across occupations (also called tasks here for brevity). If tasks are complements in production, we'd expect the coefficient $(\rho - 1)$ to be below -1. If tasks are substitutes, we'd expect $-1 < (\rho - 1) < 0$. In short, in models with exogenous technology, θ_{gjt} does not respond to supply-side shifts, and so a supply-induced increase in the task quantity ratio will reduce the task price ratio.

By contrast, the demand curve could be flat in my framework. Recall equation (10):

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

This is similar to the exogenous technology formulation (20). In both equations, log price ratio equals $\rho - 1$ times log quantity ratio plus a term for technical change. In my model, the term for technical change is a weighted average between the Old and New technologies where the weight w_{gt} is *endogenous*.

How will wages respond to a supply-side shift in my model? 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 equalise demand and supply. More generally, the endogenous technological shift will tend to offset exogenous shocks on the supply side, so the resulting impact on wages would be smaller than in the case of exogenous technology. To see why, consider a positive supply shock that increases professional employment. There is a direct negative effect on professional relative wage through the first term $(\rho - 1) < 0$. If the new technology is more intensive in professional task $r_{gj}^N > r_{gj}^O$, the lower professional wage will make the New technology more cost-effective, and thereby causing a shift to the

[43] 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.

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

	Dependent var: log $wage_{gjt}/wage_{g1t}$			
	natural resources	manufacturing	construction	trade
log emp ratio	0.2773	0.0956	-0.4225	0.1665
	(0.4058)	(0.1323)	(0.2629)	(0.5868)
j-specific trend	yes	yes	yes	yes
Observations	200	200	200	200
	transport, information	finance, business serv	other services	
ln y_gjt/y_g1t	-0.8401	0.0048	0.3706*	
	(1.3341)	(0.2665)	(0.1446)	
j-specific trend	yes	yes	yes	
Observations	200	200	200	

Standard errors in parentheses

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

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

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

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

$$\ln(\frac{p_{gjt}}{p_{g1t}}) = (\rho - 1)\ln\frac{EMP_{gjt}}{EMP_{g1t}} + \sum_{k=0}^{5}\gamma^k_{gj}t^k + u_{gjt} \tag{21}$$

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

The results are reported in table 3. I find that the key estimate $(\rho-1)$ is small and not significantly different from zero in most industries. It is negative in only two out of seven industries, and it is significantly positive in one industry. The instruments are reasonably strong: the standard errors are small enough to rule out $(\rho - 1) < -1$ in most industries. Overall, the estimates suggest the demand curve is not as downward-sloping as would be expected from standard models. My framework with endogenous technical change offers an explanation as to why it may be flat.

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

6 Empirical results

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

Given the calibrated ρ, ζ, I estimate all the other structural parameters according to the methods discussed in section 3.3. Given all the parameters, I solve for the equilibrium $(\mathbf{p_t}, \mathbf{w_t})$ in each year. I search for the equilibrium that is closest to the observed and satisfies all the equilibrium constraints within tolerance[45].

6.1 Parameter estimates and model fit

First, let's compare the estimated task intensities between the two technologies. Figure 6 shows the task intensities $\alpha^O_{gj}, \alpha^N_{gj}$ in all 7 industries. Because their identification comes from changes in occupational employment shares and wage bill shares within industries, the comparison between α^O_{gj} and α^N_{gj} is fairly robust to the calibrated value of ρ.

Take the manufacturing industry for example. It 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

[44] On the other hand, in the context of the canonical SBTC model with two education groups, **?** shows that measurement errors and ad hoc functional form assumptions of the SBTC trend are to blame for the failure of the canonical model to fit data in other papers. After addressing those issues, they estimated a high elasticity of substitution between high and low skill workers in the US. By contrast, my model implies high elasticity through the endogenous choice of technology.

[45] Demand minus supply in any occupational employment share is at most 1e-4 in absolute value. The unit costs of two co-existing technologies can differ up to 1%. This is to capture fictions in adopting a new technology.

tasks. This is what we expect. And this is driven by the data: within manufacturing employment has shifted substantially away from manual routine to managerial. Meanwhile, in non-financial services, the new technology is less intensive in admin and elementary and more intensive in all 3 abstract tasks and personal service task.

Some patterns are common across industries. In all industries, the New technology is more intensive in professional task. In 6 out of 7 industries, the New technology is less intensive in admin task, and more intensive in managerial task. In the natural resources industry, the New technology compared to the Old technology mainly involves a shift from operatives to skilled trades. Other than that, for skilled trades, in the industries where it is sizeable, the new technology is either less or equally intensive in it than the Old technology. The same is true for machine operatives. Among the lower-skilled tasks (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 three routine tasks and towards managerial and professional tasks.

Next, we examine how the key endogenous variables in the model fit the actual trends.

Figure 7 shows the observed and predicted trends in occupational employment shares. For all of the 9 occupations, the model fit is quite good. Every occupation with an observed declining [increasing] trend has a predicted declining [increasing] trend. And the difference between observed and predicted employment shares is no greater than 1% of aggregate employment. Recall that the only time-varying exogenous factors in the model are TFP of both technologies, industry demand, and aggregate skills distribution. The parameters particularly important for employment shares such as the task intensities α_{gj}^T and the occupational amenities η_j are assumed to be constant. Therefore, the design of the model does not mechanically guarantee a good fit of employment trends.

Figure 8 shows similarly a good fit for log task prices. In Appendix B, we plot the fit for log industry prices (Figure 17) and for technology shares w_{gt} (Figure 16). Note that given the 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.[46] 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 unreasonable.

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

Figure 6: Estimated task intensities in each industry

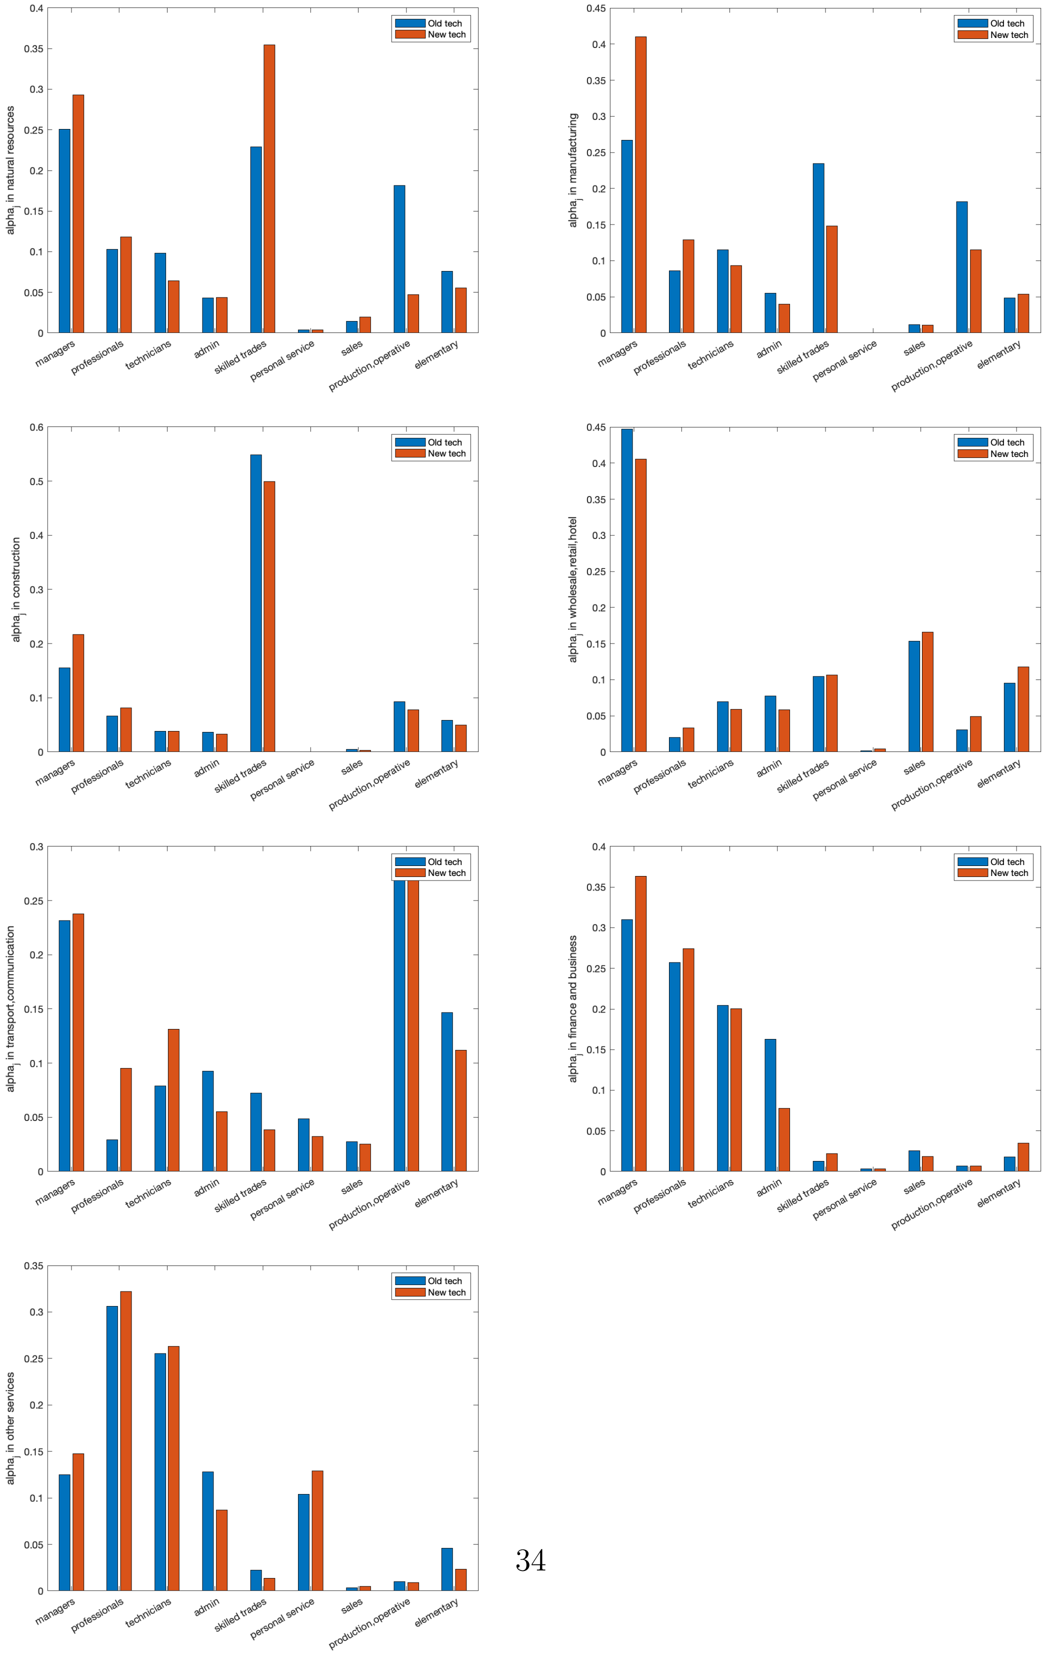

Figure 7: Fit of occupation employment share

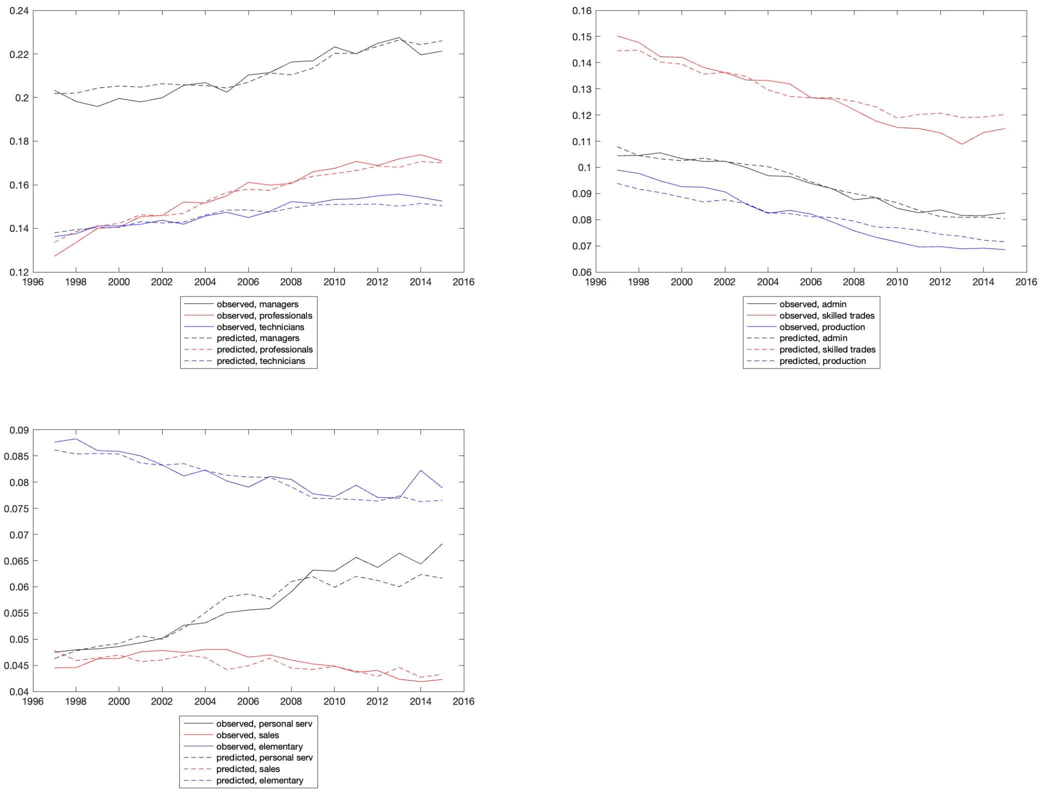

Note: The actual time trends of occupational employment shares are solid lines. The corresponding baseline predictions are dashed lines of the same colour.

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

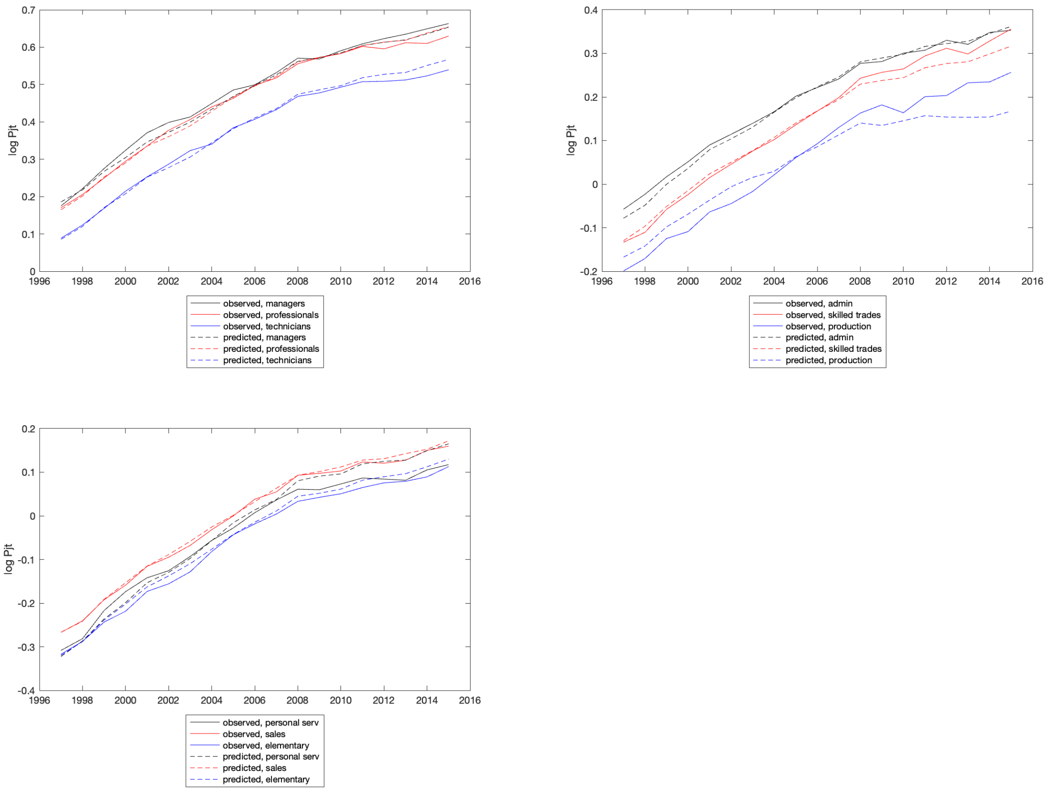

Note: The actual time trends of task prices are solid lines. The corresponding baseline predictions are dashed lines of the same colour.

6.2 Counterfactuals

The model contains three sources of exogenous time-varying factors: TFP of two technologies, industry demand, and the skills distribution. In this section, we will examine how each of them affected occupational prices and employment in the past.

In each counterfactual, only one exogenous factor changes over time while others stay the same as 1997. Because numerically there are multiple equilibria, for each year, I search for $(\log \mathbf{p_t}, \mathbf{w_t})$ that is closest to a benchmark subject to equilibrium constraints. The benchmark is the corresponding values in $t-1$ for $t > 1997$, and the observed values for the first year ($t = 1997$). I interpret the result as a lower bound on the effect of shifting that one factor.

Figure 9 considers the counterfactual where the skills distribution shifts over 1997-2015 while TFP and industry demand are constant. In this counterfactual scenario, the equilibrium task employment would shift significantly. Although I have used the same axis for employment and task price changes, the magnitude of changes should be interpreted differently. An increase of 0.01 in professional employment share is about a 10% increase from its initial employment share, whereas the 0.01 change in its log price is close to a 1% change. For skilled trades, the counterfactual employment share falls from 14.5% to 12.2%, that is -0.17 in log terms. Meanwhile, all the counterfactual wage changes are less than 3%.

Figure 9 also shows the actual changes as markers, so we can see that the supply shift alone could account for between a third and a half of the actual increase in the three abstract occupations over the 18 year period. It can also account for between one third and two thirds of the actual decline in manual routine occupations. The impact on admin occupational share is in the opposite direction to the observed change, and the impact on manual occupational shares is smaller.

Figure 10 examines the effect of industry demand shifts. We hold TFP and skills distribution constant, and let B_{gt} follow the actual trend. This counterfactual represents a shift in the demand curve. Because the industry demand shifts were strongly against manufacturing, and the manufacturing industry is very intensive in operatives and skilled trades, we see a large decline in both employment and task price for operatives and skilled trades in this counterfactual. Industry demand shifts alone can account for a third of the employment decline in machine operatives, and over half of the decline in skilled trades. It can also account for a third of the increase in professional employment and over half of the increase in technician employment.

Finally, figure 11 examines the effect of TFP shifts, holding industry demand and skill supply constant. This counterfactual leads to smaller changes in occupational employment, compared to the other two counterfactuals. It's worth noting that the counterfactual effects of the three exogenous factors do not add up to

Figure 9: Counterfactual: only skills distribution shifted

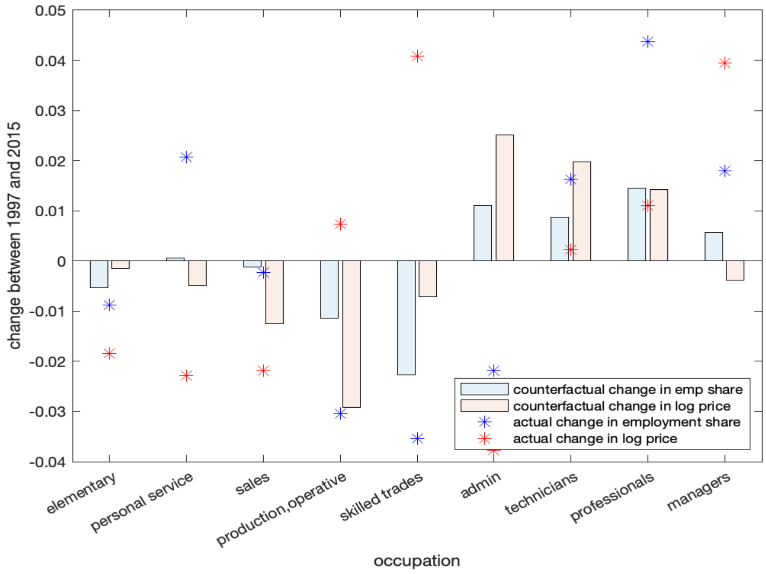

Note: using the lagged $\log P_{jt}, w_{gt}$ as the benchmark. For log task prices, we normalise the average change across 9 occupations to 0.

Figure 10: Counterfactual: only industry demand shifts

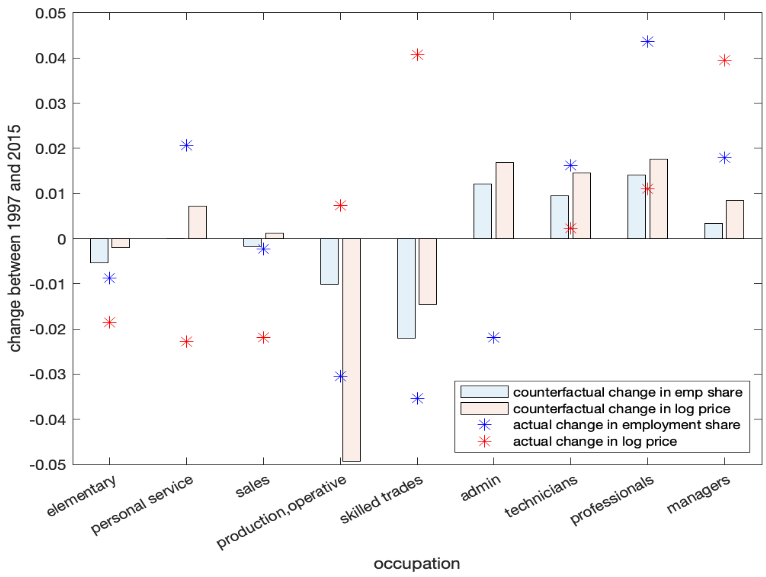

Note: using lagged $\log P_{jt}, w_{gt}$ as the benchmark. For log task prices, we normalise the average change across 9 occupations to 0.

Figure 11: Counterfactual: only TFP shifts over time

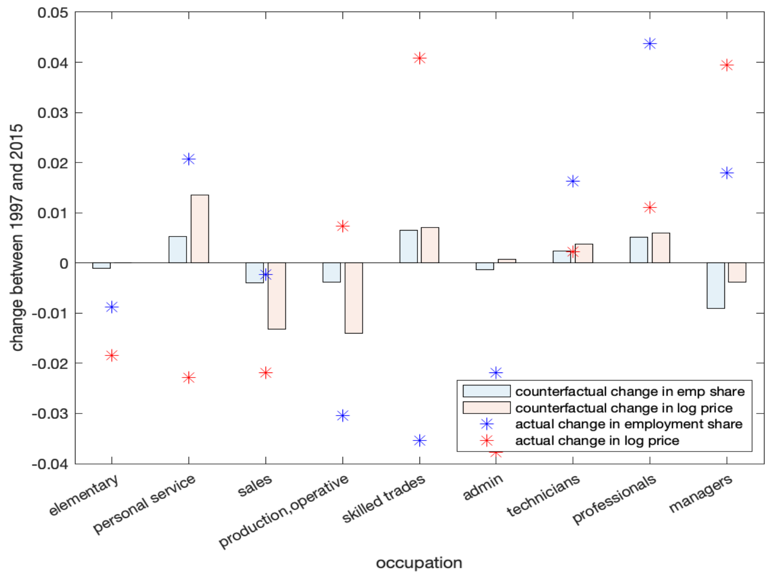

Note: using lagged log P_{jt}, w_{gt} as the benchmark. For log task prices, we normalise the average change across 9 occupations to 0.

the actual change, because there are interaction effects and potentially multiple equilibria. The results are best interpreted as the lower bounds on the effect of a single factor, holding other factors at 1997 levels.

7 Conclusion

This paper develops an equilibrium model of endogenous task-biased technological change that can simultaneous explains three notable phenomena in the UK labour market since the 90s. First, the UK has seen a large shift in employment from middle-paying occupations to high-paying ones. Second, changes in occupational wages are small and uncorrelated with employment changes. Third, there was relatively little occupational downgrading within education groups during a period of rapid increases in education. In addition to these three facts, I have provided regression analysis that supports my model and is at odds with the hypothesis of exogenous technical change.

This paper contributes to the polarisation literature by emphasising the endogenous nature of technology *adoption*. 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. To the extent that the skills distribution within graduates are stable, the model predicts little occupational downgrading within graduates.

The calibrated model can fit UK data well over 1997-2015. While the estimated direction of technical change varies across industries, the overall pattern is that the New technology is less intensive in all three routine tasks and more intensive in managerial and professional tasks, with less difference in other tasks. The shift in skills distribution alone can account for between a third and two thirds of the actual decline in routine manual occupations, and between a third and half of the increase in each of the three abstract occupations. The shift in industry demand can account for similar magnitudes of employment declines in routine manual occupations and increases in professionals and technicians.

While this paper focuses on the UK, it provides a promising framework to study issues around occupations and education in other advanced economies other than the US. Many of these countries share some of the key facts observed in the UK since the 90s. First, like the UK, employment growth has been strongest in high-paid occupations in most European countries. This is consistent with the New technology being more intensive in abstract tasks. Second, occupational wages did not polarise outside the US. And third, the US had the highest proportion of graduates in 1990 and a slower increase afterwards than many European countries. Among the European countries that saw large increases in higher education, the majority did not see a significant change on graduates' relative wage; and there has also been relatively little increase in the share of graduates in non-graduates occupations. These empirical differences between the US and the other advanced economies are intriguing, and worth further investigations.

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 becomes available and it's unambiguously better than the existing one so that all firms should adopt the new technology immediately in the absence of fixed costs or frictions. That scenario might be a good enough approximation of reality in some cases; but 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 the 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 the determination of occupational trends.

Finally, the proposed framework offers a data-driven approach to answer several policy questions about the labour market. By having analytical and social skills (instead of education) as determinants of worker productivity, it allows a lot of heterogeneity within education groups and opens up the possibility of modelling changes in the group-specific skills distribution over time. The approach also makes clear that for analysing any policies that shift labor supply, it is important to model potential changes in the distribution of skills that matter for productivity, rather than labels like education. Another interesting question left for future work 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. Future research can estimate the skill distributions of British workers, EU immigrants, and other immigrants, and simulate the policy of applying the same skill selection criteria to EU immigrants as the non-EU ones.

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

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

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

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 (2022). The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies 89*(1), 142–180. rdab034.

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

Böhm, M. J., H.-M. v. Gaudecker, and F. Schran (2019). Occupation, growth, skill prices, and wage inequality.

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.

Boustan, L. P., J. Choi, and D. Clingingsmith (2022). Automation after the assembly line: Computerized machine tools, employment and productivity in the united states.

Breemersch, K., J. P. Damijan, and J. Konings (2017). Labour market polarization in advanced countries. Technical Report 197, OECD.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002, 02). Information technology, workplace organization and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics 117*(1), 339–376.

Card, D. and T. Lemieux (2001). Can falling supply explain the rising return to college for younger men? a cohort-based analysis. *The Quarterly Journal of Economics 116*(2), 705–746.

Carneiro, P. M., K. Liu, and K. G. Salvanes (2018). The supply of skill and endogenous technical change: evidence from a college expansion reform. *NHH Dept. of Economics Discussion Paper* (16).

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.

Cavaglia, C. and B. Etheridge (2020). Job polarization and the declining quality of knowledge workers: Evidence from the uk and germany. *Labour Economics 66*, 101884.

Contractor, Z. and B. Taska (2022). the effect of software adoption on skill demand. Technical report.

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.

Green, F. and G. Henseke (2016). The changing graduate labour market: analysis using a new indicator of graduate jobs. *IZA Journal of Labor Policy 5*(1), 1–25.

Green, F. and G. Henseke (2021). Europe's evolving graduate labour markets: supply, demand, underemployment and pay. *Journal for Labour Market Research 55*(1), 1–13.

Hardy, W., R. Keister, and P. Lewandowski (2018). Educational upgrading, structural change and the task composition of jobs in europe. *Economics of Transition 26*(2), 201–231.

Hershbein, B. and L. B. Kahn (2018, July). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review 108*(7), 1737–72.

Humlum, A. (2019). Robot adoption and labor market dynamics. Technical report, Princeton University.

Johnson, M. and M. P. Keane (2013). A dynamic equilibrium model of the us wage structure, 1968?1996. *Journal of Labor Economics 31*(1), 1–49.

Kampelmann, S. and F. Rycx (2011). Task-biased changes of employment and remuneration: The case of occupations.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics 107*(1), 35–78.

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

Machin, S. and J. Van Reenen (1998). Technology and changes in skill structure: evidence from seven oecd countries. *Quarterly journal of economics*, 1215–1244.

Michaels, G., A. Natraj, and J. Van Reenen (2014, 03). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics 96*(1), 60–77.

Murphy, E. C. and D. Oesch (2018). Is employment polarisation inevitable? occupational change in ireland and switzerland, 1970–2010. *Work, employment and society 32*(6), 1099–1117.

Naticchioni, P., G. Ragusa, and R. Massari (2014). Unconditional and conditional wage polarization in europe.

Oesch, D. and G. Piccitto (2019). The polarization myth: Occupational upgrading in germany, spain, sweden, and the uk, 1992–2015. *Work and Occupations 46*(4), 441–469.

Patel, A. (2022). the role of firms in shaping job polarization. Technical report.

Salvatori, A. (2018). The anatomy of job polarisation in the uk. *Journal for labour market research 52*(1), 1–15.

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

A Appendix

A.1 Derivation of a demand-side equation

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

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

$$p_{jt} = p_{gt}\frac{\partial Y_{gt}^T}{\partial y_{gjt}^T} = p_{gt}\alpha_{gj}^T(y_{gjt}^T/Y_{gt}^T)^{\rho-1} \quad \forall j, g, t, T \in \{O, N\} \tag{22}$$

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

$$\frac{p_{jt}}{p_{1t}} = \frac{\alpha_{gj}^T}{\alpha_{g1}^T}(\frac{y_{gjt}^T}{y_{g1t}^T})^{\rho-1} \quad \forall j, g, t, T \in \{O, N\} \tag{23}$$

$$\frac{y_{gjt}^T}{y_{g1t}^T} = (\frac{p_{jt}\alpha_{g1}^T}{p_{1t}\alpha_{gj}^T})^{\frac{1}{\rho-1}} \quad \forall j, g, t, T \in \{O, N\} \tag{24}$$

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

$$\frac{EMP_{gjt}}{EMP_{g1t}} = \frac{y_{gjt}^O}{y_{g1t}^O + y_{g1t}^N} + \frac{y_{gjt}^N}{y_{g1t}^O + y_{g1t}^N} \tag{25}$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N}\frac{y_{gjt}^O}{y_{g1t}^O} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N}\frac{y_{gjt}^N}{y_{g1t}^N} \tag{26}$$

$$= \frac{y_{g1t}^O}{y_{g1t}^O + y_{g1t}^N}(\frac{p_{jt}\alpha_{g1}^O}{p_{1t}\alpha_{gj}^O})^{\frac{1}{\rho-1}} + \frac{y_{g1t}^N}{y_{g1t}^O + y_{g1t}^N}(\frac{p_{jt}\alpha_{g1}^N}{p_{1t}\alpha_{gj}^N})^{\frac{1}{\rho-1}} \tag{27}$$

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

$$r_{gj}^O = (\alpha_{gj}^O/\alpha_{g1}^O)^{1/(1-\rho)} \tag{28}$$

$$r_{gj}^N = (\alpha_{gj}^N/\alpha_{g1}^N)^{1/(1-\rho)} \tag{29}$$

Equation (27) simplifies to

$$\frac{EMP_{gjt}}{EMP_{g1t}} = (\frac{p_{jt}}{p_{1t}})^{\frac{1}{\rho-1}}[(1 - w_{gt})r_{gj}^O + w_{gt}r_{gj}^N] \tag{30}$$

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

$$\ln(\frac{p_{jt}}{p_{1t}}) = (\rho - 1) \ln \frac{EMP_{gjt}}{EMP_{g1t}} + (1 - \rho) \ln[(1 - w_{gt}) r^O_{gj} + w_{gt} r^N_{gj}] \tag{31}$$

A.2 Derivation of task supply equation

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

$$\begin{aligned} y(a, s, j) &= E[y(i, j)|a_i = a, s_i = s] && (32) \\ &= k_j e^{\beta_{aj} a + \beta_{sj} s} E[e^{\mu_i}|a_i = a, s_i = s] && (33) \end{aligned}$$

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

Going back to (7) and using (33) to substitute for $k_j e^{\beta_{aj} a + \beta_{sj} s}$, we get

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

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

Given task prices, the supply of task j is

$$\begin{aligned} LS_j(\mathbf{p}) &= \sum_i \pi_j(a_i, s_i, \mathbf{p}) y(i, j) && (34) \\ &= \int \int \pi_j(a, s, \mathbf{p}) y(a, s, j) f(a, s) da ds && (35) \end{aligned}$$

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

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

This section derives equation (11).

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

$$unitcost^T_{gt} = [\sum_j (\alpha^T_{gj})^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}]^{1-1/\rho} / A^T_{gt} \tag{36}$$

The ratio of unit costs between the two technologies is:

$$\frac{unitcost_{gt}^N}{unitcost_{gt}^O} = \frac{A_{gt}^O}{A_{gt}^N}[\frac{\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}}]^{1-1/\rho} \tag{37}$$

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

$$\begin{aligned}
&[\frac{\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}}]^{1-1/\rho} = \frac{A_{gt}^N}{A_{gt}^O} \\
\Rightarrow &\frac{\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}}{\sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}}} = (\frac{A_{gt}^N}{A_{gt}^O})^{\frac{\rho}{\rho-1}} \\
\Rightarrow &\sum_j (\alpha_{gj}^N)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}} - (\frac{A_{gt}^N}{A_{gt}^O})^{\frac{\rho}{\rho-1}} \sum_j (\alpha_{gj}^O)^{\frac{1}{1-\rho}} p_{jt}^{\frac{\rho}{\rho-1}} = 0 \\
\Rightarrow &\sum_j [(\alpha_{gj}^N)^{\frac{1}{1-\rho}} - (\frac{A_{gt}^N}{A_{gt}^O})^{\frac{\rho}{\rho-1}} (\alpha_{gj}^O)^{\frac{1}{1-\rho}}] p_{jt}^{\frac{\rho}{\rho-1}} = 0
\end{aligned}$$

A.4 Derivation of an equation to identify TFP terms

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

$$\begin{aligned}
&p_{gt} A_{gt}^T [\sum_j \alpha_{gj}^T (y_{gjt}^T)^\rho]^{\frac{1-\rho}{\rho}} \alpha_{gj}^T (y_{gjt}^T)^{\rho-1} = p_{jt} \\
\Rightarrow &p_{gt} A_{gt}^T [\sum_j \alpha_{gj}^T (y_{gjt}^T)^\rho]^{\frac{1-\rho}{\rho}} [\alpha_{gj}^T (y_{gjt}^T)^\rho]^{\frac{\rho-1}{\rho}} (\alpha_{gj}^T)^{1/\rho} = p_{jt} \\
\Rightarrow &[\frac{1}{[\sum_j \alpha_{gj}^T (y_{gjt}^T)^\rho]}]^{\frac{\rho-1}{\rho}} [\alpha_{gj}^T (y_{gjt}^T)^\rho]^{\frac{\rho-1}{\rho}} = \frac{p_{jt}}{p_{gt} A_{gt}^T (\alpha_{gj}^T)^{1/\rho}} \\
\Rightarrow &\frac{\alpha_{gj}^T (y_{gjt}^T)^\rho}{[\sum_j \alpha_{gj}^T (y_{gjt}^T)^\rho]} = [\frac{p_{jt}}{p_{gt} A_{gt}^T (\alpha_{gj}^T)^{1/\rho}}]^{\frac{\rho}{\rho-1}} \\
\Rightarrow &1 = \sum_j [\frac{p_{jt}}{p_{gt} A_{gt}^T (\alpha_{gj}^T)^{1/\rho}}]^{\frac{\rho}{\rho-1}} \\
\Rightarrow &(p_{gt} A_{gt}^T)^{\frac{\rho}{\rho-1}} = \sum_j [\frac{p_{jt}}{(\alpha_{gj}^T)^{1/\rho}}]^{\frac{\rho}{\rho-1}}
\end{aligned} \tag{38}$$

Figure 12: Employment and wage growth by ISCO major group

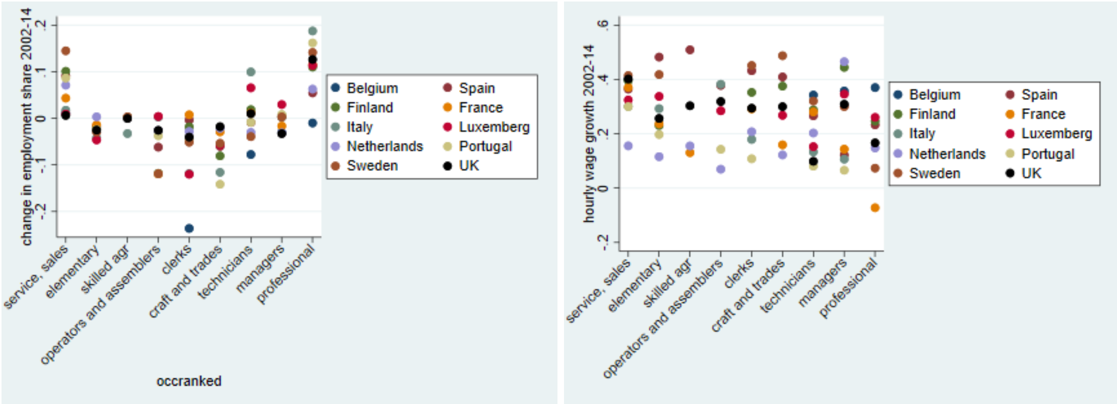

Source: SES 2002 and 2014. To compute the change in hourly wages, we exclude cells where the occupation's employment share has more than tripled or halved because those cases may involve large compositional changes.

B Additional figures

The flip side of the flat proportion of abstract occupations within graduates is that almost all of the aggregate increase in abstract occupations' share can attributed to the increase in education. Using the LFS (1997-2015), I decompose the change in occupational employment shares into within-gender-education-group component and between-group component. Figure 13 suggests that all of the increase in abstract employment is between-group, and almost all of the decline in skilled trades and operative employment is between-group.

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

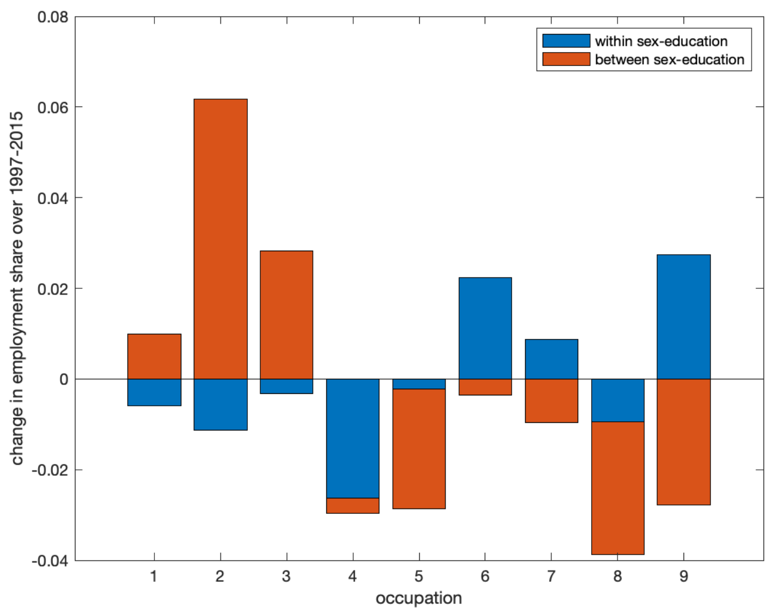

Source: UK Labour Force Survey

Figure 14: Adjusting occupational wage for classification changes

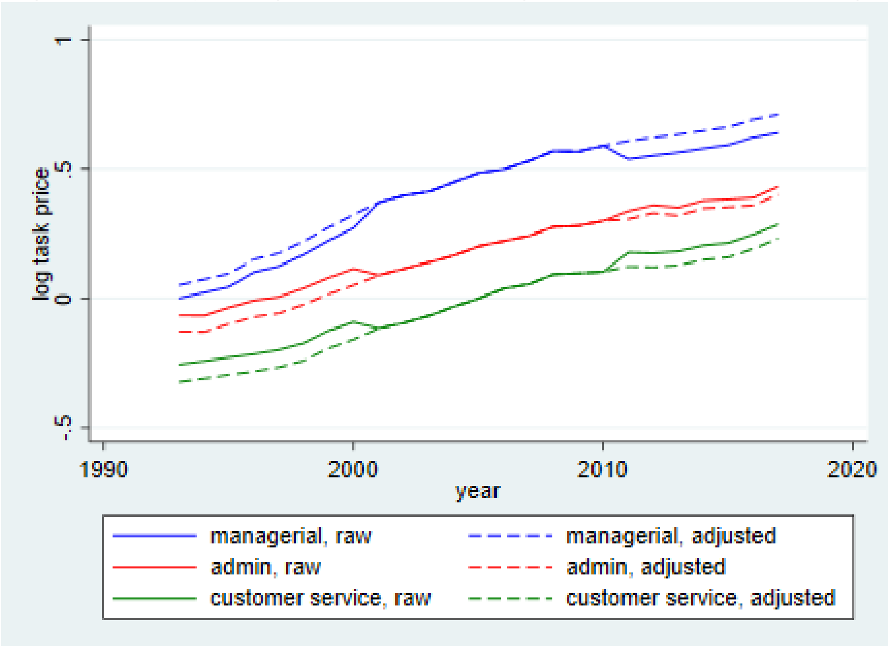

Source: UK Labour Force Survey 1993-2017.

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

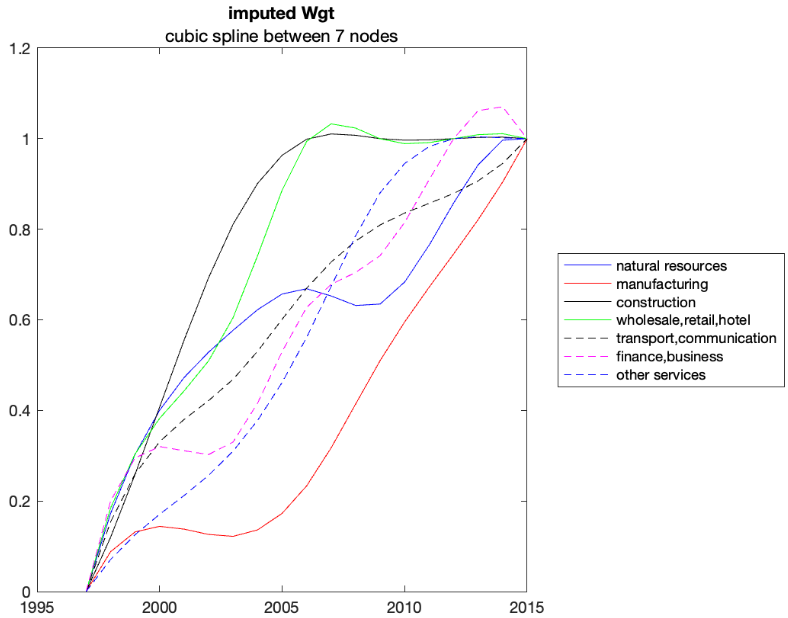

Note: We have 4 measures of capital composition from 1997 to 2015 annually and 5 measures from the BSS available at 4-5 points between 1992 and 2017. Because the different measures have different scales, I standardise each measure within industry so that when I minimise the sum of squared ϵ^m_{gt}, they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

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

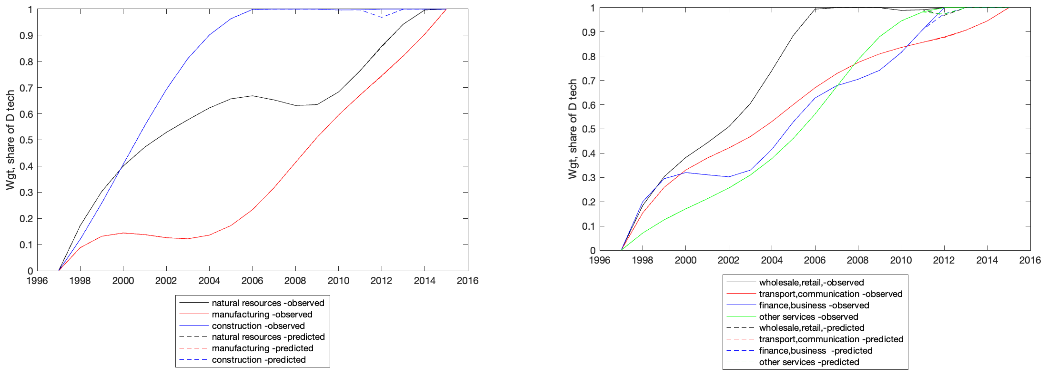

Note: The actual time trends of technology shares are solid lines. The corresponding baseline predictions are dashed lines of the same colours.

Figure 17: Fit of log industry prices P_{gt}

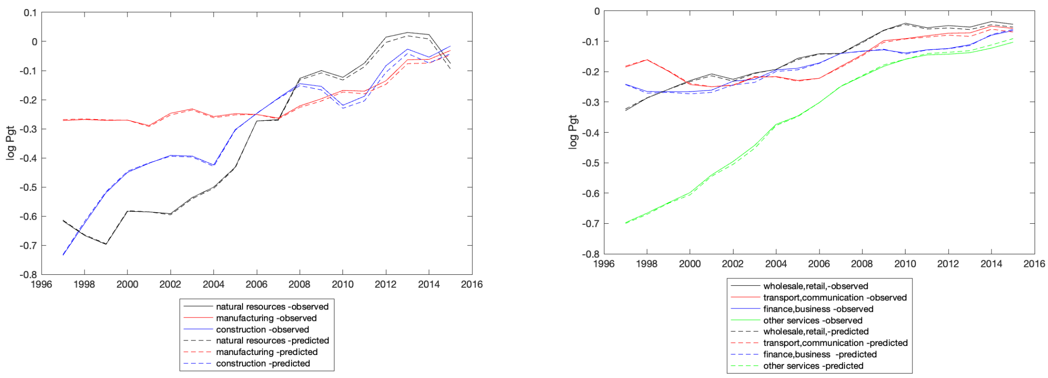

Note: The actual time trends of industry prices are solid lines. The corresponding baseline predictions are dashed lines of the same colour.