

Blundell, Richard; Green, David A.; Jin, Wenchao

Working Paper

The UK wage premium puzzle: How did a large increase in university graduates leave the education premium unchanged?

IFS Working Papers, No. W16/01

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Blundell, Richard; Green, David A.; Jin, Wenchao (2016) : The UK wage premium puzzle: How did a large increase in university graduates leave the education premium unchanged?, IFS Working Papers, No. W16/01, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2016.1601>

This Version is available at:

<https://hdl.handle.net/10419/173939>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The UK Wage Premium Puzzle: How did a Large Increase in University Graduates Leave the Education Premium Unchanged?

IFS Working Paper W16/01

Richard Blundell
David A. Green
Wenchao Jin

The Institute for Fiscal Studies (IFS) is an independent research institute whose remit is to carry out rigorous economic research into public policy and to disseminate the findings of this research. IFS receives generous support from the Economic and Social Research Council, in particular via the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP). The content of our working papers is the work of their authors and does not necessarily represent the views of IFS research staff or affiliates.

The UK Wage Premium Puzzle: How Did a Large Increase in University Graduates Leave the Education Premium Unchanged?*

Richard Blundell[†]

David A. Green[‡]

Wenchao Jin[§]

May 2016

Abstract

Since the early-1990s the UK experienced an unprecedented increase in university graduates. The proportion of people with a university degree by age 30 more than doubled from 16% for those born in 1965-69 to 33% for those born ten years later. At the same time the age profile of the graduate premium remained largely unchanged across cohorts. This paper first establishes the facts using a detailed analysis of micro-data on wage and employment patterns over the last two decades, benchmarked against the US economy. We then show that the stability of the age profile in the premium across different birth cohorts is unlikely to be explained by either composition changes or selection on unobservables. We also argue that it is inconsistent with skill-biased technical change affecting all advanced economies in the same way. We further rule out explanations based on factor price equalisation. Our resolution of the puzzle is a model in which increases in level of education induce firms to transit toward a decentralised technology in which decision-making is spread more widely through the workforce. We provide empirical support for this view.

Keywords: Wage premium, education differential, skill biased technical change, general purpose technology.

*We thank seminar participants at UBC and IFS for helpful comments on earlier drafts. This research is part of the research programme of the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS. We are grateful to the ESRC, grant number RES-000-23-1524, for financial support. The usual disclaimer applies.

[†]University College London and Institute for Fiscal Studies

[‡]Vancouver School of Economics, UBC

[§]University College London and Institute for Fiscal Studies

Introduction

In the period extending from the mid-1990s to the present, the UK economy experienced a substantial transformation in educational attainment. Approximately 16% of the birth cohort born between 1965 and 1969 (who turned 25 in the early 1990s) held a university degree by age 30. For the cohort born just one decade later (who turned 25 in the early 2000s), the percentage with a university degree had more than doubled to 33%.

The speed of this increase is unmatched in most developed economies and substantially faster than in the US, the comparison country we focus on here. Given estimates of the impact of cohort-specific supplies of higher education workers on wage differentials for the US (e.g., [Card and Lemieux \[2001\]](#)), one might reasonably predict the UK education shifts to have had large effects on the UK wage structure. This was not the case. For the 1965-69 birth cohort, at age 30, the ratio of the median hourly wage for those with a first degree or higher (referred to as “BA” henceforth) to that of those with a secondary-education qualification (referred to as “high-school” or HS henceforth) was 1.45.¹ In comparison, the BA-to-HS wage ratio at age 30 for the much higher educated 1975-79 birth cohort was 1.48. And this pair of cohorts is not an anomaly. As shown in [Figure 1](#), the entire profiles of BA-to-HS wage differentials mapped against age are all essentially right on top of one another for the cohorts born after about 1960.

The similarity of age profiles for wage differentials across cohorts is striking, considering that the percentage with a Bachelors or higher degree at age 30 ranges from a low of 13% for the 1960-64 birth cohort to a high of 37% for the 1980-85 cohort. The puzzle we seek to address in this paper is, how did the UK economy absorb such a large increase in the number of educated workers without experiencing a relative wage response?²

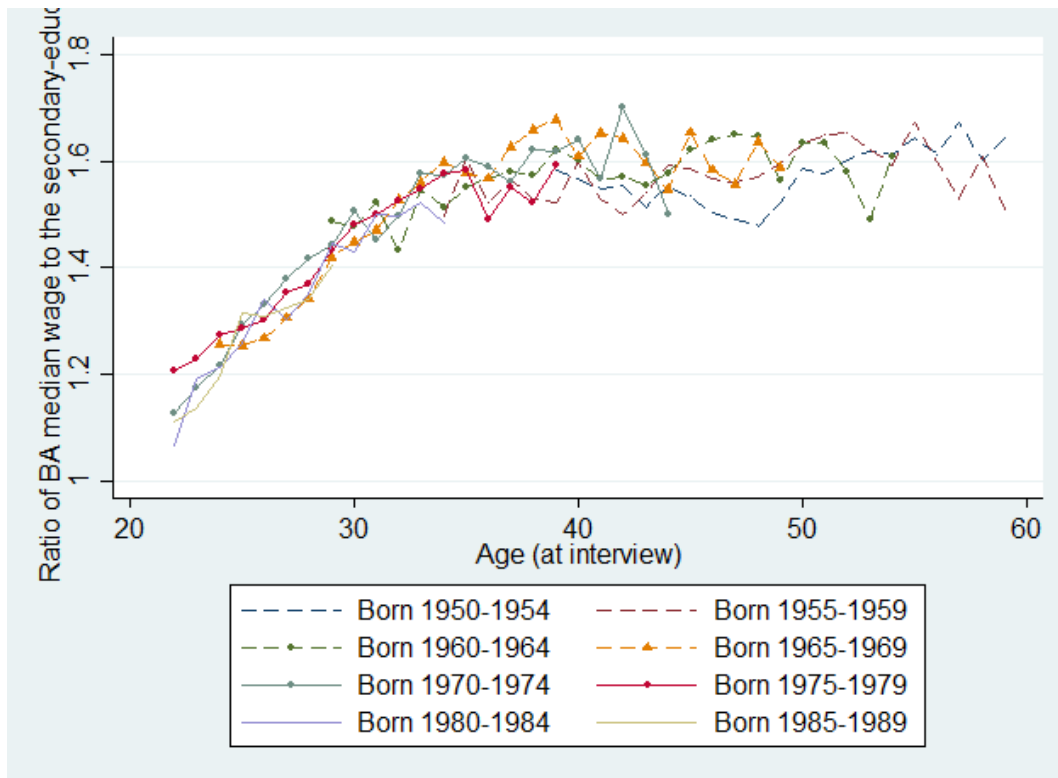
Our approach follows three steps. At the outset of the paper, we will establish a core set of facts using mainly UK Labour Force Survey (LFS) data. This will include the proportion of individuals with university degrees and more detail on the relative wage patterns just described. We present our results at the birth cohort level, finding that when we control for a common age profile, the education wage differentials are flat across the 1965-69 through 1975-79 birth cohorts. We find some decline in the post-1980 birth cohorts, those entering the labour market in the mid-2000s - a point to which we return when considering explanations. Note that the period investigated in the paper is after period when the BA-HS wage differential in the UK increased substantially [[Machin and McNally, 2007](#)].

In [section 1.3](#), we provide a benchmark for the UK trends using the same plots for

¹Throughout this paper, BA refers to people with first degrees or equivalent. High school (HS) refers to people who obtained GCSE or equivalent qualifications at age 16. People who obtain vocational or sub-degree-level qualifications after getting their O-levels or GCSEs but have no degree-level qualifications are categorized into HS.

²[Chen \[2013\]](#) shows that Taiwan underwent a similar sized increase in university attainment at approximately the same time as the UK and that it, too, experienced little change in its education wage differential. She argues that this can be explained with a model of skill biased technical change.

Figure 1: Ratio of BA median wage to that of high-school graduates by 5-year birth cohorts, UK



Note: authoer's calculation from the Labour Force Survey.

the US. In particular, we use the US data to estimate a more flexible version of the [Card and Lemieux \[2001\]](#) specification, regressing cohort specific changes in the BA/HS wage differential on a time trend and changes in the ratio of BA to HS employment. We employ those estimates to ask what one would predict for changes in the BA/HS wage differentials if the UK had faced the same elasticity of substitution between BA and HS workers and the same relative demand shift as the estimates imply for the US. We find that under this scenario, the UK BA/HS wage differential would have been predicted to fall by about 20 percentage points between the 1950-55 cohort and the 1985-89 cohort one, whereas in reality there is no significant difference between these two cohorts. Thus, a simple story based on an exogenous technical change affecting all advanced countries does not fit with the UK data patterns. In short, our puzzle is how the UK absorbed such large increases in educational attainment with the education wage differential being unchanged when a reasonable benchmark suggests that it should have fallen substantially.

One possible explanation is that compositional shifts across cohorts mean either that the increase in the supply of skills was not as large as it appeared or that compositional shifts disguised an actual decline in the education wage differential. If either of these were true then there is no real puzzle, that is, it is not really the case that large increases in educational attainment were accompanied by negligible changes in the education

wage differential. The second step in our analysis will be to try to uncover the extent to which the core patterns reflect compositional shifts across cohorts. Section 2 examines shifts associated with immigration (the UK shifted to being a substantial immigration receiving country in this period [Dustmann et al. \[2013\]](#)); the increase in the proportion of workers with above-Bachelors university degrees ([Lindley and Machin \[2013\]](#)); and potential differences between public and private sector workers. Neither the immigration shifts nor the increase in the proportion of workers with advanced degrees changes our core patterns. Looking at the private sector only, we continue to find an unchanging education wage differential for the birth cohorts with the largest increases in educational attainment (the 1970-74 and 1975-79 cohorts) but we find some added evidence of a larger decline in the education differential among the post-1980 cohorts that enter the labour market in the mid-2000s and after. A successful explanation will need to account for this combination of a lack of change in the differential in the cohorts with the biggest educational gains followed by a small decline in the differential for subsequent cohorts for which the increases in education are smaller.

Section 2.4 then discusses potential composition shifts in terms of unobserved abilities. With such a large shift of people from lower to higher education, there could have been considerable changes in the average and median abilities within each education group, and it is possible that the wage ratio patterns we described are ultimately due to offsetting selection effects. We assess this possibility by adopting different assumptions about who constituted the increased degree holders in a bounding exercise. Under an hierarchical ability model, we show that bounds on movements in the education differential continue to point to at most small changes in the education differential between the 1965-69 and 1975-79 cohorts. The bounds, however, do point to the possibility of larger declines in the education differential for the post-1980 cohorts.

In another exercise, we ask how much selection is needed to reconcile the UK relative wage movements with their US counterparts, finding that the required selection is, in some years, implausibly large. Moreover, the extent of the selection would have to vary considerably across time and age in a manner just sufficient to leave the BA-HS wage ratio unchanged. We cannot rule out such time varying composition shifts but we find it very unsatisfying as an explanation and search for economic models that generate the observed wage patterns endogenously.

Thus, in the third step we evaluate competing economic theories for why the UK would have experienced large increases in the supply of high education workers but no change in the education wage differential. One key possibility is that these patterns are reflections of factor price equalization from trade theory. To the extent this is true, we should observe shifts in industrial composition toward skill-intensive sectors. We look for this pattern in UK LFS data but do not find evidence in favour of it: the proportion of workers with a university degree increased substantially in all sectors. This result points to the possible influence of technological change in the form of a General Purpose Technology that affects production in all sectors of the economy.

The standard model of that General Purpose Technology arriving as an exogenous

technical change seems to us to provide an unsatisfying explanation. In particular, the combination of a substantial increase in the number of university degree holders and the lack of a change in their relative wages indicates the presence of a set of demand shifts that exactly offset the variable changes in educational supply across cohorts. In our view, explanations in which such demand shifts arrive exogenously are not very believable. Versions of exogenous skill-biased demand shift models based on tasks that have been used to discuss polarization can be shown to imply a similar specification for an educational wage differential regression to that derived from the canonical model and so do not improve on the explanation. Given this, our belief is that it is more reasonable to consider economic models in which the lack of change in the education differential arises endogenously.

The model we propose builds on the insights in [Caroli and Van Reenen \[2001\]](#) and [Bloom et al. \[2012\]](#) that the use of IT capital is accompanied by the use of decentralized organizational structures. [Bloom et al. \[2012\]](#) show that US multinationals use a more decentralized structure relative to UK firms even when both are observed operating in the UK. They argue that UK firms are laggards in adopting decentralized structures because of regulation based inflexibilities. We offer an alternative explanation: that at the time of the development of the new IT related structures, the low education level in the UK implied that it was less profitable for UK firms to adopt the new approach. Then, as the UK education level ramped up, the UK underwent a technological transformation.

To capture these ideas, we propose a model in which firms can choose between a *centralized* organizational structure, with a few managers directing large numbers of other workers, and a *decentralized* structure in which decision making and supervision is spread more widely through the workforce. Firms employ high educated and low educated workers, with the former being relatively more productive at decision-making and management related tasks. The structure of the model is the same as a standard 2-good, 2-factor trade model, implying that the factor price invariance result holds. In particular, as the proportion of high educated workers increases, firms shift toward using the decentralized technology but, as long as the economy is operating in the cone of diversification, the wages for the two groups do not change. Eventually, though, with continuing increases in the education level of the workforce, the economy will transit to using only the decentralized technology, after which the education wage differential will begin to decline. Notice that, in this framework, there is no clear distinction between graduate and non-graduate jobs. It is the set of tasks in a job that defines an appropriate match with a skill level. The bundle of tasks associated with a job will change endogenously with a change in the organisational structure induced by the change in the proportion of educated workers.

We argue that this model fits with the core patterns in the UK data. In particular, we show that underlying the lack of change in the educational differential are completely invariant HS and BA real wage levels - a remarkable pattern that is implied by the model. In addition, the model can account for the seemingly odd pattern of the wage differential declining for the cohorts after the ones with the biggest increases in edu-

cational attainment. We investigate further implications by examining the relationship between the education composition of the workforce and the extent to which workers feel they control how they do their own work using matched worker-workplace survey data from the UK Workplace and Employer (WERS) data. We show that the regions where the increases in the BA proportion were largest had the greatest uptake of decentralized organizational forms.

We are not the first researchers to note the substantial increase in degree-holding in the UK. For example, [Carpentier \[2004\]](#) documented the trend in student numbers from 1920 to 2002, which increased sharply around the early 90s. They also showed a reduction in university expenditure per student around the same time. Many other studies have also documented the substantial increase in the share of graduates in the 1990s or across cohorts [[OLeary and Sloane, 2005](#), [Walker and Zhu, 2008](#), [Green and Zhu, 2010](#), [Devereux and Fan, 2011](#)]. The cross-cohort movement has been used by some researchers as a source of variation for identifying the return to a university degree in the UK. [[Devereux and Fan, 2011](#)].

Previous papers have also noted the lack of a reduction in the college wage premium over time or across recent UK cohorts [[Machin and McNally, 2007](#), [McIntosh, 2006](#), [Walker and Zhu, 2008](#)]. However, those papers either appeal to the explanation that there have been offsetting relative demand shifts stemming from exogenous skill biased technical change or do not attempt to explain the lack of change in the relative wages at all. We add to the previous literature by providing a more thorough investigation of potential causes of the wage and employment patterns. We also present a potential explanation that does not rely on exogenous skill biased demand shifts that just happen to be the right size to match the change in educational attainment across each possible pair of cohorts.

1 The puzzle and the background

1.1 Changes to educational attainment

We begin with the two plots that define the puzzle we investigate. We will use the shorthand of calling the group with university degrees BA's, even though it includes other types of Bachelors degrees and more advanced degrees. For both the US and the UK, we summarize the data by plotting cohort effects corresponding to 5-year birth cohorts from a regression of either the ratio of median BA to median HS wages or the BA proportion on a fifth order polynomial in age and a complete set of cohort dummies.³ We focus on cohort effects because we believe this is the right level of variation for considering variation in education outcomes, allowing us to abstract from lifecycle patterns.⁴

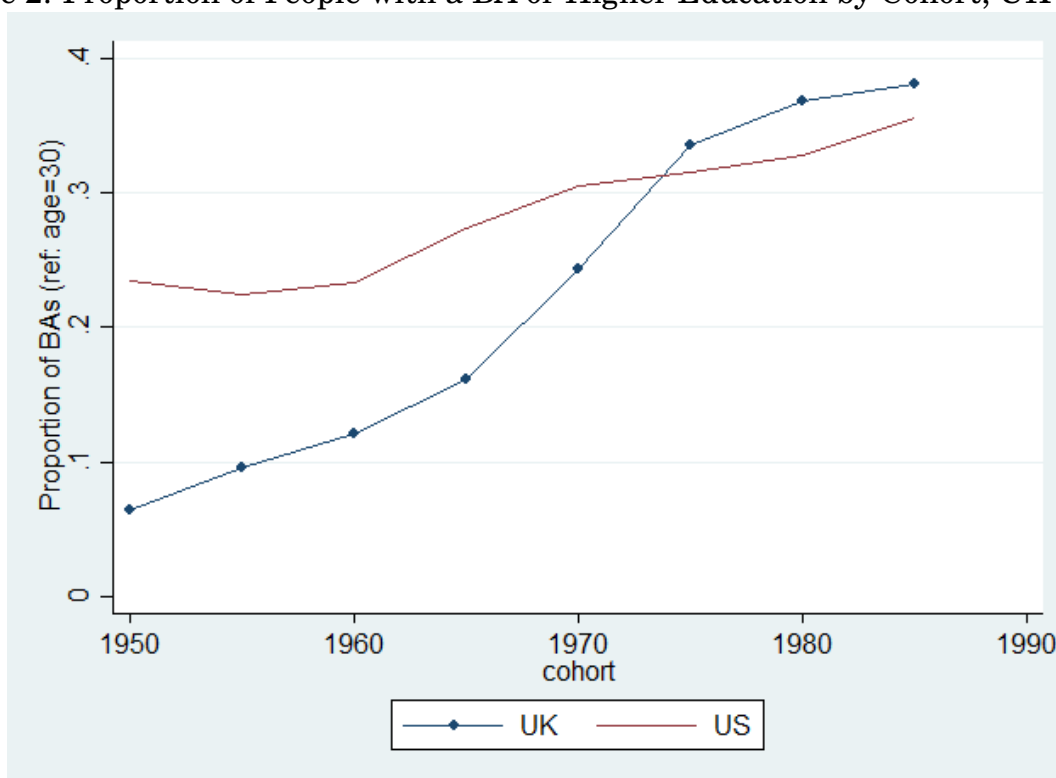
³For the BA plots, we normalize the age variables so that the intercept corresponds to age 30 and then add the estimated intercept to our cohort effects in the plots so that the heights correspond to effects at age 30.

⁴Previous studies have also selected cohorts as the appropriate level of variation (e.g., [Fortin \[2006\]](#)).

Constructing cohort effects in this way controls for differences in the age ranges over which the different cohorts are present in our data.

Figure 2 contains plots of the cohort effects for the BA proportion for both the UK and the US. The UK has strikingly lower proportions of people with a university education in the cohorts born before 1965; proportions that are less than half those in the corresponding US cohorts. But between the 1965-69 and the 1975-79 cohorts, the UK completely closes that gap, experiencing increases in the proportion with a BA of around 8 percentage points between each cohort. After the 1975-79 cohort, the UK proportion of BAs grows at a much slower rate that is more similar to the US growth rate. Remarkably, between the 1960-64 cohort and the 1980-84 cohort, the proportion of people with a BA in the UK tripled; between the 1965-69 and 1975-79 cohorts, alone, it approximately doubled.

Figure 2: Proportion of People with a BA or Higher Education by Cohort, UK and US



Note: sample restricted to age 22-59 and excludes full-time students. Each education-cohort cell has at least 100 observations.

The big increase in the UK proportion in the BA group between the 1965-69 and 1975-79 cohorts corresponds to a rapid increase in higher education enrolment from 1988 to 1994. This large increase has been documented in many studies (O'Leary and Sloane [2005], Carpentier [2006], Walker and Zhu [2008], Green and Zhu [2010], Devereux and Fan [2011]) and has been used as an arguably exogenous source of variation in studies of the causal impact of education [Devereux and Fan, 2011].

The expansion of higher education over the past few decades was a reflection of specific policy choices made by the UK government. Since the Robbins Report in 1963, policy related to the higher education sector has been moving toward implementation of the principle that university places should be available to all who are qualified by ability and attainment to pursue them and who wish to do so. The 1960s saw the foundation of more than 20 universities and dozens of polytechnics. Polytechnics were a form of higher education institution that taught both degree-level courses and below-degree-level courses, with their degrees certified by a chartered body called Council for National Academic Awards (CNAA). A CNAA degree from a polytechnic was technically equivalent to a university degree and we treat them as equivalent in our analysis. The Education Reform Act (ERA) of 1988 changed some block grants to tuition fees (paid by Local Education Authorities for each student). In response, polytechnics increased enrolment with lower funding per student. The other major education policy change in 1988 was the replacement of CSEs and O-Levels with GCSEs as the exams that students take at age 16. That reform led to an increase in education attainment at the secondary level and hence an increase in the proportion of the young with sufficient academic credentials for admission to universities. In 1992, polytechnics gained the right to issue degrees and became fully-fledged universities. The reclassification of polytechnics as universities led to a big jump in the number of university students in 1992; but the rapid increase in student numbers in higher education started in 1988 and continued until 1994.⁵ This corresponds to the 1970-74 and 1975-79 cohorts in our plots since those cohorts would have turned 18 between 1988 and 1997. In 1994, due to pressures on public expenditures and the desire to protect resources per student, the government introduced the maximum student number control. This limited the number of full-time undergraduates at individual universities every year. As a result, the growth in student numbers slowed down, though, as we can see from the post 1975-79 cohorts in Figure 2 it by no means stopped.

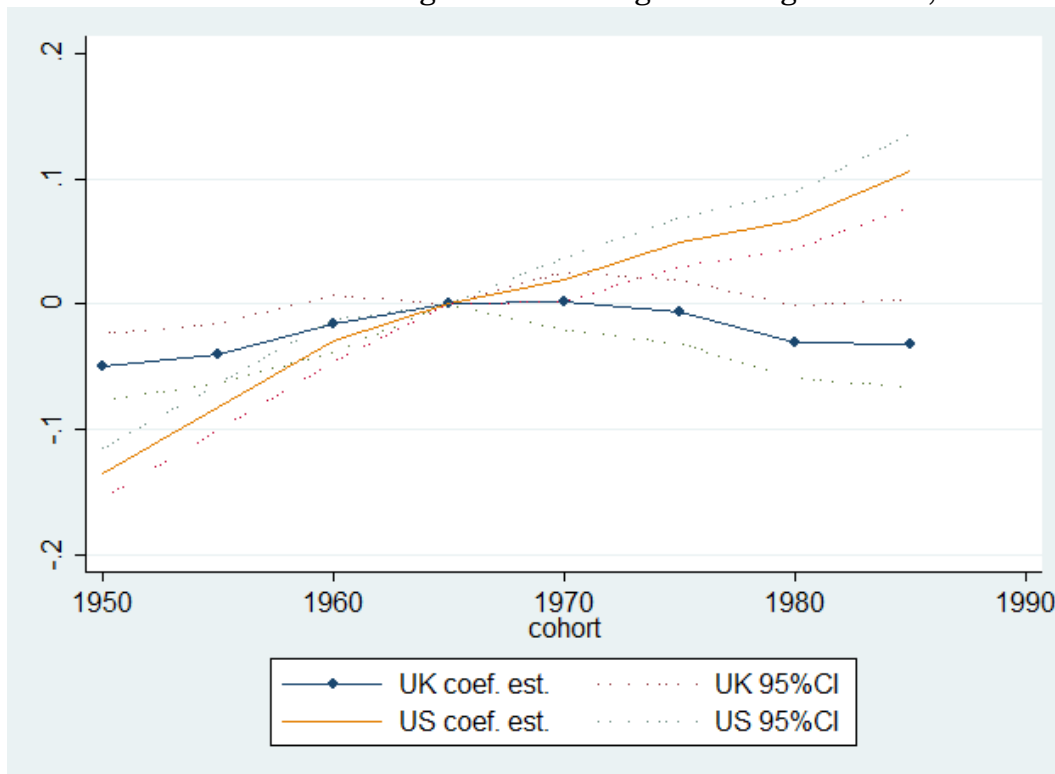
1.2 Changes in relative wages

The second main pattern relates to wages. Similar to Figure 2, we regress the median wage ratios on cohort dummies and an age polynomial, and plot the estimated cohort effects along with their 95% confidence intervals in Figure 3. We normalize the ratios to zero for the 1965-69 cohorts in both countries. In the US, the wage ratio increases rapidly for the 1954-59 and 1960-64 cohorts (cohorts that would have entered the labour market in the 1980s), grows more slowly until the 1970-74 cohort, and then resumes somewhat more rapid growth with the 1975-79 cohort (i.e., roughly with people entering the labour market after 2000). This fits with what has been previously documented about movements in the education differential in the US, albeit presented here at the cohort instead of the year level (see, e.g., Card and DiNardo [2002]). In comparison,

⁵This has been clearly shown in Figure 2 in Carpentier [2006]

the UK experienced some increase in the education differential across the pre-1965-69 cohorts (albeit smaller than for the US), fitting with results about the UK education differential before 1990 (Machin [2001]). But between the 1965-69 and 1975-79 cohorts, when the proportion of cohort members with a BA or higher degree doubled, the education differential is almost completely unchanged. There is some decline in the point estimates for the post-1980 birth cohorts but the decline is only about 5 percent relative to the 1965-69 cohorts and is not statistically significant at the 5 percent level.

Figure 3: Ratio of BA median wage to that of high-school graduates, cohort effects



The absence of significant changes to the relative wages between the 1965 and 1975 cohorts is consistent with previous studies which found the UK graduate wage premium to be stable in the 90s and early 2000s [Chevalier et al., 2004, McIntosh, 2006, Machin and Vignoles, 2006, Machin and McNally, 2007, Walker and Zhu, 2008]. Two earlier papers [OLeary and Sloane, 2005, Walker and Zhu, 2005], using data up to 2003, found the university premium to have fallen somewhat over the cohorts that experienced the higher education expansion. However, the authors later revised their cohort conclusions with more years of LFS data in Walker and Zhu [2008]. The latter paper reported that on average, there was little reduction to the university premium for men between the pre-expansion and post-expansion cohorts and a 10% rise for women. The change in results appears to stem from being able to follow cohorts beyond their mid-20s in the more recent data.

1.3 Establishing a Benchmark Using the Canonical SBTC Model for the US

The combination of the large increase in the supply of more educated workers and the lack of change in the relative wage in the UK in the post 1965 birth cohorts necessarily implies that the UK experienced a relative demand shift favouring the more educated in the last two decades.

In this sense, examining data at the national level for the UK in isolation, one could explain the wage and education patterns by appealing to a standard story of Skill Biased Technical Change (SBTC). Of course, without a direct measure of technological change such a story cannot be rejected⁶: the technological change becomes a residual that takes whatever value is needed to exactly offset the supply shifts. Often, though, technological change is assumed to have the same effect in all developed countries (Antonczyk et al. [2010]), and we can use that idea to form a benchmark against which to judge the UK experience. We do so using the US, which we see as a good benchmark both because of its similarities to the UK in terms of recent inequality movements and general structure of the economy, and because the SBTC model has been most fully investigated for the US. In particular, we will use the Card and Lemieux [2001] (CL henceforth) estimation framework. Their approach is useful for us because it is couched in cohort related terms.

CL start from an aggregate production function expressed as a function of college (C) and high school (H) workers:

$$y_t = (\theta_{Ht}H_t^\rho + \theta_{Ct}C_t^\rho)^{\frac{1}{\rho}} \quad (1)$$

where, θ_{Ht} and θ_{Ct} are factor specific productivity shifters and $\rho = 1 - \frac{1}{\sigma_E}$, with σ_E being the elasticity of substitution between the education groups. To allow for differences across cohorts, C_t and H_t are, in turn, written as CES aggregators of employment in the education group in specific age groups indexed by j . The cohort specific productivity shifters are given by α_j for the H group and β_j for the C group. They assume a common elasticity of substitution between the age groups, σ_A . In this context, skill biased technology changes are reflected in the relative movements of θ_{Ct} and θ_{Ht} .

Given this specification, we can write the relative wage equation as:

$$\log \frac{W_{j,t}^C}{W_{j,t}^H} = \log \frac{\beta_j}{\alpha_j} + \theta_t - \frac{1}{\sigma_E} \log \frac{C_t}{H_t} - \frac{1}{\sigma_A} (\log \frac{C_{j,t}}{H_{j,t}} - \log \frac{C_t}{H_t}) + e_{j,t} \quad (2)$$

where, $\theta_t = \frac{\theta_{Ct}}{\theta_{Ht}}$ and $e_{j,t}$ is some idiosyncratic measurement error or shock.⁷ CL assume that the skill biased technological change embodied in movements in θ_t follow a linear

⁶The SBTC literature has often used R&D and computer usage to proxy technological progress, as in Machin and Van Reenen [1998], Bresnahan et al. [1999].

⁷Note that θ_t can capture not only technology-driven changes to the productivities of different education groups (such as IT innovations), but also exogenous demand shifters like the liberalization of the financial sector and China entering the WTO. We thank Ian Walker for pointing this out.

trend (i.e., $\theta_t = \gamma_0 + \gamma_1 t$) - a standard assumption in the early part of the skill biased technical change literature. In the appendix, we report the results from estimating (2) with a linear trend for θ_t with UK and US data for the past two decades for various subsamples defined by the wage concept (hourly and weekly) and gender. For the US, the estimated substitution elasticity between different age groups, σ_A , is near the top of the range of estimates reported by [Card and Lemieux \[2001\]](#). But the estimate of σ_E , the substitution elasticity between the HS and BA labour inputs varies wildly by subsample and wage measure. For the UK, the estimates of σ_A have the wrong sign and the σ_E estimates are again highly variable. The model's failure to fit even the US data for our period (1993-2012) is not totally surprising. The CL model is essentially a richer version of [Katz and Murphy \[1992\]](#), and [Beaudry and Green \[2005\]](#) show that the Katz-Murphy model fits the US data (in the sense of having elasticity estimates of the right sign) less and less well the more recent the data that is used in the estimation. In addition, the US and UK estimates are so different that one can soundly reject the hypothesis that the data in the two countries are being generated from a common model of this form. Thus, the canonical model does not fit either country well and certainly cannot provide a common explanation for patterns in the two of them together.

We can resurrect the model for the US data by allowing for a nonlinear relative productivity trend. In particular, we re-estimated (2) with the US data allowing for the relative productivity trend θ_t to follow a 7th order polynomial.⁸ We then used the parameters from the US estimation to predict relative wage trends in the UK. That is, we ask what one would predict for the movement of UK wages based on UK relative supply movements but US estimates of the skill biased demand trend θ_t and US estimates of substitution elasticities and age effects in (2).

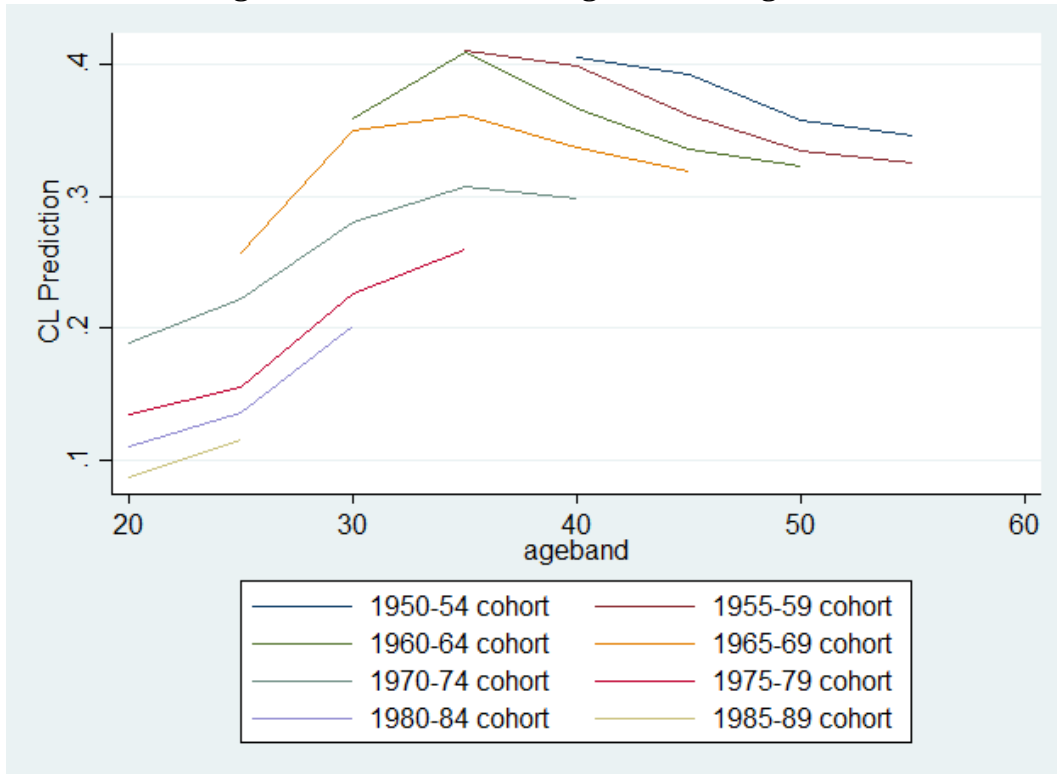
To compare the predicted wage ratios against the observed ones in [Figure 1](#), we convert the predictions at the level of age-band-year to the level of cohort-age-band.⁹ In contrast to [Figure 1](#), [Figure 4](#) shows the predicted relative wage would shift downwards across cohorts and especially between the 1965 and 1975 cohorts, which experienced the sharpest expansion of Higher Education.

To isolate cohort effects, we use the fitted wage ratios in a regression on a fifth-order age polynomial and a set of cohort dummies, as before. We plot the resulting cohort effects along with the UK cohort effects from [Figure 3](#) together in [Figure 5](#). Several points emerge from the figure. First, based on US derived demand patterns and UK supply movements, one would have predicted a roughly 20 percentage point decline in the BA-HS wage ratio in the UK between 1950 and 1985 cohorts; but the actual wage

⁸We found the higher-order terms are insignificant and that with a 7th order polynomial, we obtained estimates for the elasticities that were in the range of those reported in [Card and Lemieux \[2001\]](#) while for lower order polynomials we often obtained quite different (and sometimes economically nonsensical) estimates that varied considerably across groups defined by gender and different wage measures.

⁹Birth cohorts are five-years apart. As the correspondence between year and birth cohort conditional on any 5-year-age-band is not exact, we compute age-band-cohort-level predictions as a weighted average of age-band-year-level predictions.

Figure 4: Predicted UK log BA-HS wage ratio



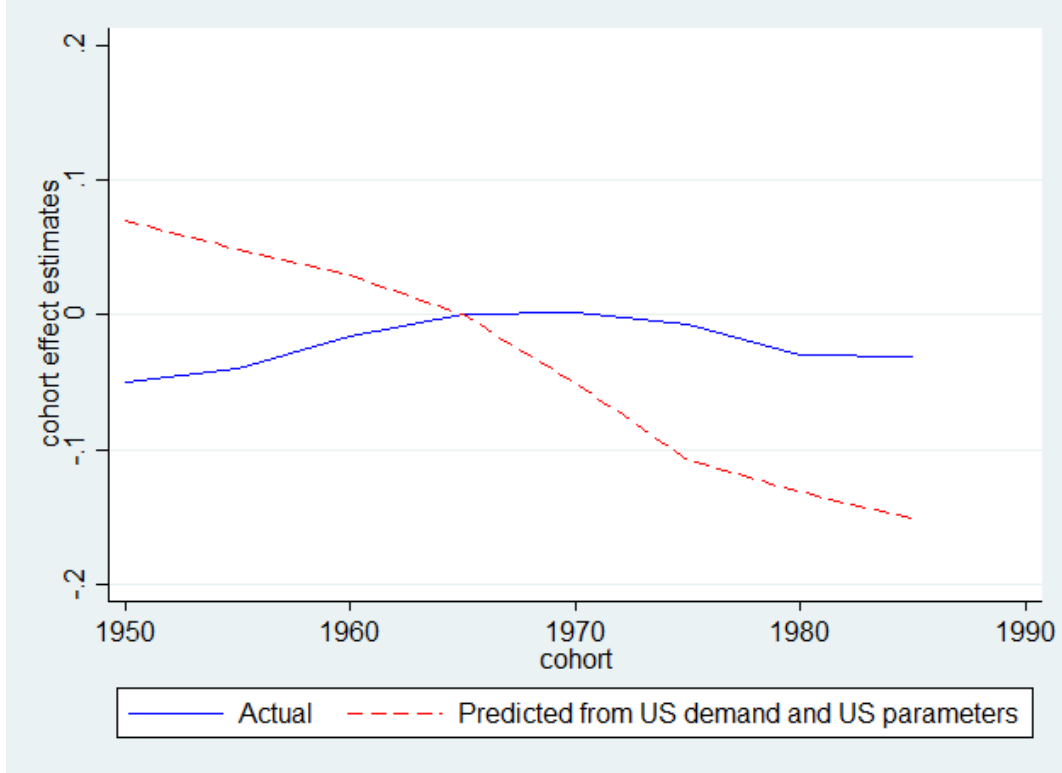
Note: this prediction is based on the estimates of US demand trend and US parameters.

ratio was no lower for the 1985 cohort than the 1950 one conditional on an estimated age profile. Second, the fall in the predicted wage ratio was especially large between 1965 and 1975 cohorts (a bit over 10 percentage points), when the supply increase was the fastest. But in reality, the wage ratio was quite constant between these two cohorts and only started to fall for the 1980 cohort. Third, the predictions imply a decline in the BA-HS wage ratio in the pre-1965 cohorts when both the actual line and the consensus in the literature point to an increase in the differential. Thus, based on relative demand shifts estimated from US data combined with the near tripling of the proportion of people in the UK with higher education, both the relative lack of change in the education wage differential and the timing of the changes that do occur are puzzling. In the remainder of the paper, we investigate potential explanations for that puzzle.

2 The Effects of Changes in Composition

The simplest answer to our puzzle is that compositional shifts are obscuring the true patterns. The potential importance of composition shifts becomes evident if we think of workers as bundles of efficiency units of tasks. More able workers supply a larger number of efficiency units per hour worked, and, in a standard neoclassical model, their observed wages will reflect this. As a result, observed average wages can increase either because of increases in the market price per efficiency unit or because the composition

Figure 5: Cohort effects in the predicted UK log BA-HS wage ratio



of workers shifts in the direction of a higher average number of efficiency units per worker. Since our result is that the observed BA-HS wage differential has not fallen relative to what we might expect, the scenario of greatest potential interest for us is one in which the price differential for BA versus HS tasks declines while the differential in their average efficiency increases.

First, we can rule out gender shifts among education groups as the main cause because the same puzzle exists within each gender separately. Figures 17 in the Appendix show that within each gender, there have been large increases in the proportion of BAs across cohorts but no substantial changes to relative wages. In the rest of the section, we will consider the following compositional shifts: immigration; exactly a BA versus advanced degrees; public versus private sector status; and composition in terms of unobservables.

2.1 Immigration

The proportion of UK workers born outside the UK has doubled over the past two decades. As immigrants are more likely to have university degrees, the large flows of immigrants contribute directly to the aggregate increase in the share of BAs in the workforce. But it is not clear whether we should count every immigrant with a university education as the equivalent of a university educated native born worker. As demonstrated in [Dustmann et al. \[2013\]](#), immigrants often work in jobs that do not match

their observed skills or qualifications, implying that a simple count of the number of immigrants with a university education may over-state the contribution of immigration to the effective supply of highly educated labour. Given the size of the increase in the immigrant proportion in the past 20 years, the positive bias in the measured supply of university labour may become substantial.

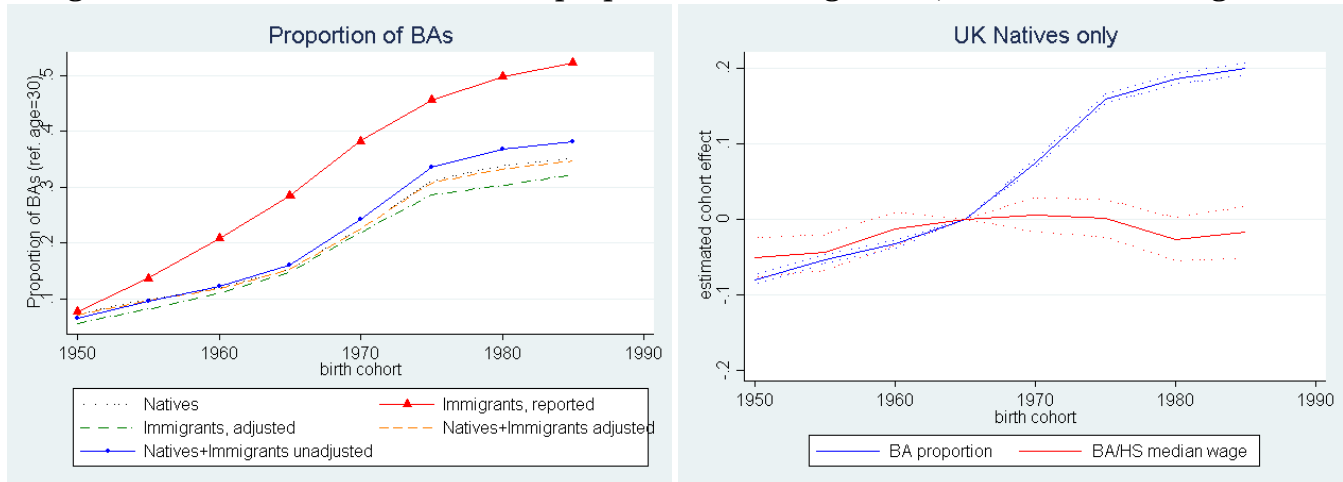
Following the spirit of [Dustmann et al. \[2013\]](#), we carry out a counterfactual exercise in which we re-assign immigrant education levels according to their wages relative to the natives' wage distribution. Specifically, within each calendar year and 5-year age band, we find the percentile of the native wage distribution that corresponds to the immigrant's wage, calculate the proportions of the native born in each education group at that percentile, age-band and year, and assign the immigrant a new education level drawn according to the native born probabilities. For non-working immigrants, we assign a new education level from the education distribution of non-working native born workers. We have done the adjustments for all immigrants, although one may argue that immigrants who arrive as children should be treated the same as natives and so our adjustment provides an upper bound on the extent of over-stated education. The top line in the left subfigure of [Figure 6](#) shows our standard birth cohort effects for the BA proportion for immigrants. For the post-1960 cohorts, it is typically at least 0.1 above the corresponding line for the native born. Echoing the results in [Dustmann et al. \[2013\]](#), the adjusted proportion of immigrants with a BA is much lower; indeed, it is everywhere lower than the native born line. However, the adjusted proportion with a BA obtained by combining the true proportion for the native born with the adjusted proportion for immigrants is not substantially different from the original, unadjusted proportion with a BA. The large increase in the proportion with a BA between the 1965-69 and 1975-79 cohorts, in particular, is relatively unaffected by the adjustment. This is because the proportion of the cohorts that are immigrants stays relatively small until the last few cohorts. The clear message is that the potential mis-reporting of effective immigrant education does not alter the main pattern.

Issues with immigrant skill measurement could also affect our wage measures. To check on whether this is an issue, in the right subfigure of [Figure 6](#) we plot both the BA proportion and the education wage differential (both normalized to zero in the 1965-69 cohort) for the native born alone. The main pattern of a strong increase in the BA proportion matched with little change in the education differential after the 1965-69 cohort is still present without immigrants. Based on these two exercises, we conclude that we cannot explain the combination of education growth and flat education differentials through composition effects related to immigration.

2.2 Postgraduate education expansion

In 1993, 2.3% of 25-34 year olds had qualifications higher than a Bachelor's degree; in 2014, that percentage had risen to 10.6% [[Lindley and Machin, 2013](#)]. The dark, solid line in [Figure 7](#) plots the proportion of people with a postgraduate degree conditional

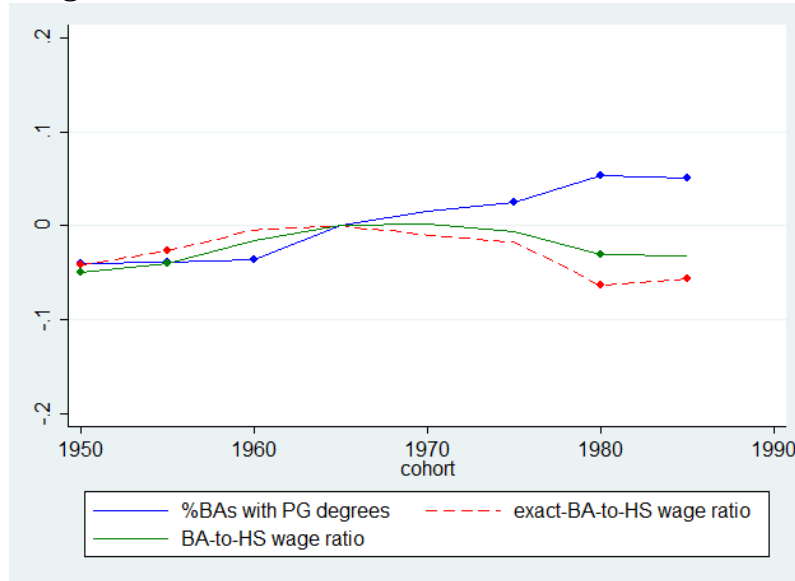
Figure 6: Cohort effects for the BA proportion and wage ratio, natives and immigrants



Note: the left sub-figure is based on the same specification and normalization as in Figure 2. In the right sub-figure, the cohort effects are normalized to zero at the 1965 cohort, as in Figure 3.

on having a university degree of some kind normalized to zero for the 1965-69 cohort. That proportion was also increasing across cohorts, particularly between the 1975-79 and 1980-84 cohorts - just after the substantial increase in the proportion of people with any university degree. The rise in the proportion of our “BA” sample with advanced degrees might be predicted to increase the median wage for university graduates as a whole; though this is not necessarily the case if the people who move into advanced degrees would have been above-median university earners even without the extra degree. To assess the impact of increasing advanced degrees on our data patterns, we examine the ratio of the median wage of exact BAs (excluding postgraduates) to that of HS workers. Strikingly, the wage ratio still shows very little movement between the 1965-69 and 1975-79 cohorts, which may reflect the fact that the biggest shifts toward advanced degrees happened in the subsequent cohorts. The post 1980 cohorts now show a statistically significant decline in the ratio relative to the 1965-69 cohort but at just over $-.05$, the decline is only slightly larger than what we observed for the combined university graduates plot in Figure 3. Thus, while the shift toward more advanced degrees may have played a role in mitigating the measured decline for the post-1980 cohorts, it cannot explain the result that the cohorts with the biggest increases in the proportion with a university degree saw no change in their wage differential relative to high school graduates.

Figure 7: Cohort Effects for the Proportion of University Graduates with Advanced Degrees and the BA to HS Wage Ratio



Note: marked points are significantly different from zero.

Meanwhile, the proportion of university graduates who have a "good" degree increased from 38% in 1985 to 47% in 1993. [Naylor et al., 2015] Here "good" refers to first-class and upper-second-class degrees as opposed to lower-second and third-class degrees under the British undergraduate degree classification. If there was no loosening in the criteria of degree classifications, which is a big if, then this could be interpreted as signs of a positive shift in the quality of graduates over time. But this shift would be small in magnitude given that the wage premium of a "good" degree is estimated to be below 10% [Naylor et al., 2015]. In short, observed categorizations of BAs by whether they have postgraduate degrees and by undergraduate degree class cannot explain the puzzle.

2.3 Public versus private sector

The public sector in the UK employs about a quarter of all workers. The public sector workforce is on average more educated, better paid, and proportionally more female. One possible explanation for the stability in the education differential is that the public sector expanded to absorb the extra graduates and set wages which do not reflect individuals' marginal productivities. At the very least, one might expect an increase in public sector employment associated with the expansion of post-secondary institutions. The public sector's employment share did increase somewhat until the recent austerity cuts. To gauge the impact of that expansion, we collect workers into three sectors (self-employed, paid private sector, and paid public sector) and perform a standard decomposition exercise, holding the proportion of workers in each sector constant over time. We present the results of that exercise in the Appendix but the simple summary is

that virtually none of the increase in the proportion of workers with a university degree between 1994 and 2014 (or within any sub-period) is accounted for by shifts across these three sectors. The increase in BA education in the workforce was ubiquitous within all three sectors.

A division between the public and private sectors may also be relevant to the extent we view public sector wages as shielded from the main forces operating in the economy. In that sense, an influx of graduates into the public sector might not be expected to reduce their wages relative to non-graduates. In Figure 8, we plot the cohort effects for the median wage ratio constructed as before but only for private sector employees. As in Figure 3, for private and public sector employees combined, there is little change in the education differential across the cohorts with the largest increases in the BA proportion. Between the 1965-69 and 1970-74 cohorts (when the BA proportion increased by over 8 percentage points) the wage ratio declines by about 1 percent, and between the 1970-74 and 1975-79 cohorts (when there was another 8 percentage point increase in the BA proportion), the wage ratio declined by a further 3 to 4 percent. But outside of those cohorts, the private sector patterns are quite different from the overall patterns. The pre-1965 cohorts show a decline instead of an increase in the education differential across cohorts while the post-1979 cohorts show a more substantial decrease in the wage differential than what is shown in Figure 3. The decline in the ratio across the older cohorts is troubling since it does not fit with previous work indicating that there was an increase in the return to education in the 1980s in the UK. What appears to be behind the pattern in Figure 8 is a negative time effect in the education differential affecting all cohorts in the same way in the last few years of our sample period. In Figure 9, we regress the wage differential at the year-age-band level on age dummies and year dummies and plot the year effects. It shows a relatively flat time trend until the last couple of years of the data. Since the last few years are more important for recent than earlier cohorts, this time effect disproportionately drags down more recent cohort effects. In fact, when we re-plot of Figure 8 using data only up to 2012 (see Figure 18 in Appendix), there is an increase in the education differential across the early cohorts and less of a decline in more recent cohorts (for example, the 1975-79 cohort effect is 2.5 percent below the 1965-69 value instead of 4.5 percent below including the data up to 2014). Given this, our conclusion is that the private sector continues to show a relative constancy of the education differential across the cohorts experiencing the largest increases in university attainment but also seems to have experienced some decline in the differential in the post 1980 cohorts. We will return to this pattern later in the paper.

2.4 Unobserved quality and relabelling

In this section, we use a variety of approaches to try to address the question of whether the composition of unobservable characteristics has shifted across education groups in a way that could explain the wage patterns.

Figure 8: Cohort effects for the BA-to-HS wage differential, UK private-sector only

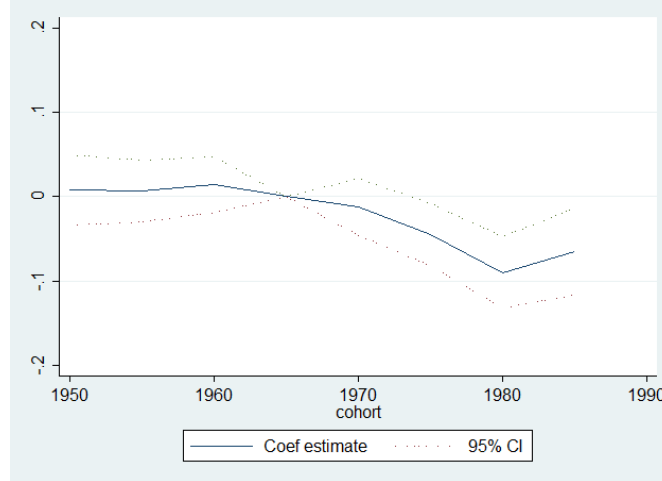


Figure 9: Time effect in BA-to-HS wage differentials, UK private-sector only



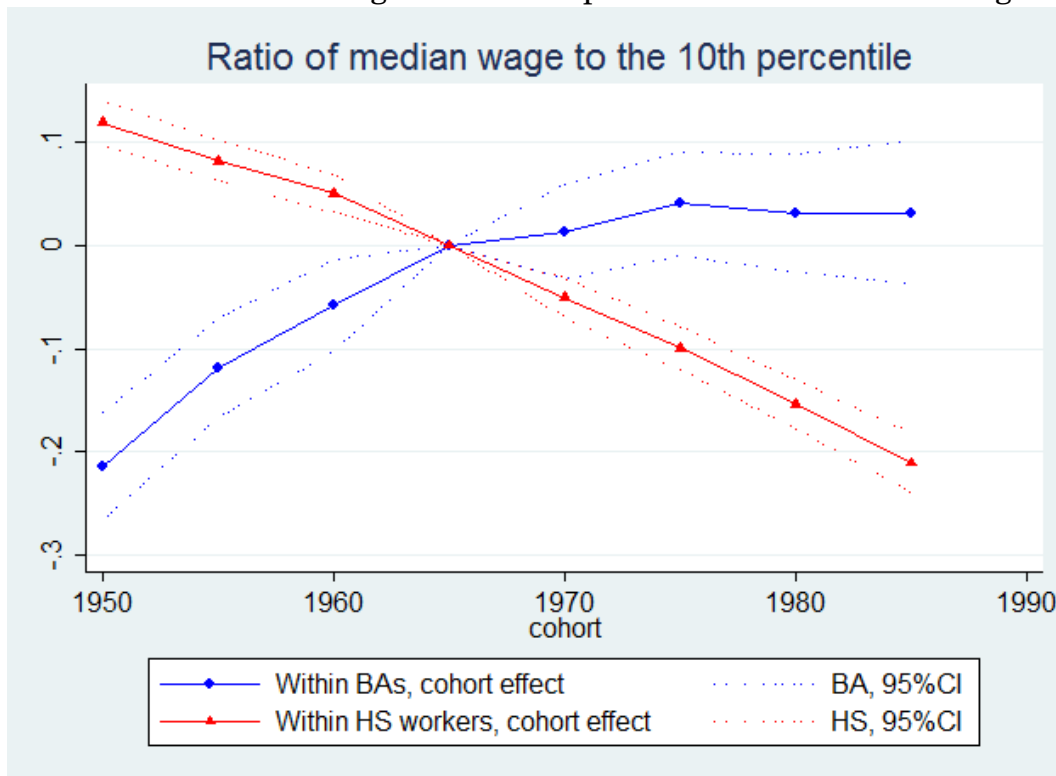
Note: as normalization, the average time effect of the period till 2010 is set to zero.

As higher education expands, it draws in pupils from an increasingly wider range of prior attainment and perhaps innate ability. The fall in per student resources that came along with the rapid expansion in the UK from 1988 to 1994 might also have had a negative impact on the quality of courses and hence of graduates. Thus, it seems possible that the average quality of BA workers has declined across cohorts. On the other hand, the quality of HS-educated workers is also likely to fall if the more able individuals in more recent cohorts now go to university and some of those who would have been HS dropouts previously now obtain secondary qualifications. If these scenarios are true, then it is theoretically ambiguous whether the ability-composition constant wage ratio of the education groups is greater or smaller than the observed one. This is important because it is the movement in the composition constant wage ratio (essentially, the ratio of prices for the high and low educated labour factors) that is relevant for investigations of the forces driving changes in the labour market.

2.4.1 Quantiles and Employment

The idea that BAs have a lower and wider range of quality after the higher education expansion has been advocated in [OLeary and Sloane \[2005\]](#) and [Walker and Zhu \[2008\]](#). Both papers use quantile regressions to estimate the university wage premium across different periods or cohorts, and they report a greater decline in the premium at lower quantiles than at higher quantiles. While it's tempting to interpret such results as evidence of declining quality of BAs at the lower end of the BA wage distribution, examining the wage distributions for BA and HS workers separately suggests a different conclusion. We regress the ratio of the median to the 10th percentile of the BA wage distribution on a fifth order age polynomial and cohort dummies, and plot the estimated cohort effects in Figure 10. Since the 1965 cohort, there has been no significant change in the dispersion of wages in the lower half of the BA wage distribution. In contrast, the same line for the HS group shows a strong decline in the 50-10 ratio across cohorts. This indicates that the relative decline in low end wages for BA's versus HS workers is driven by improvements for the latter group. One possible explanation for that improvement is the introduction of the National Minimum Wage (NMW) in 1999 and its rise relative to the median wage over the 2000s. But regardless of the explanation, these patterns make it difficult to conclude that the fall of the graduate premium at lower quantiles is due to a greater deterioration in the quality of BAs than HS workers at their respective lower ends.

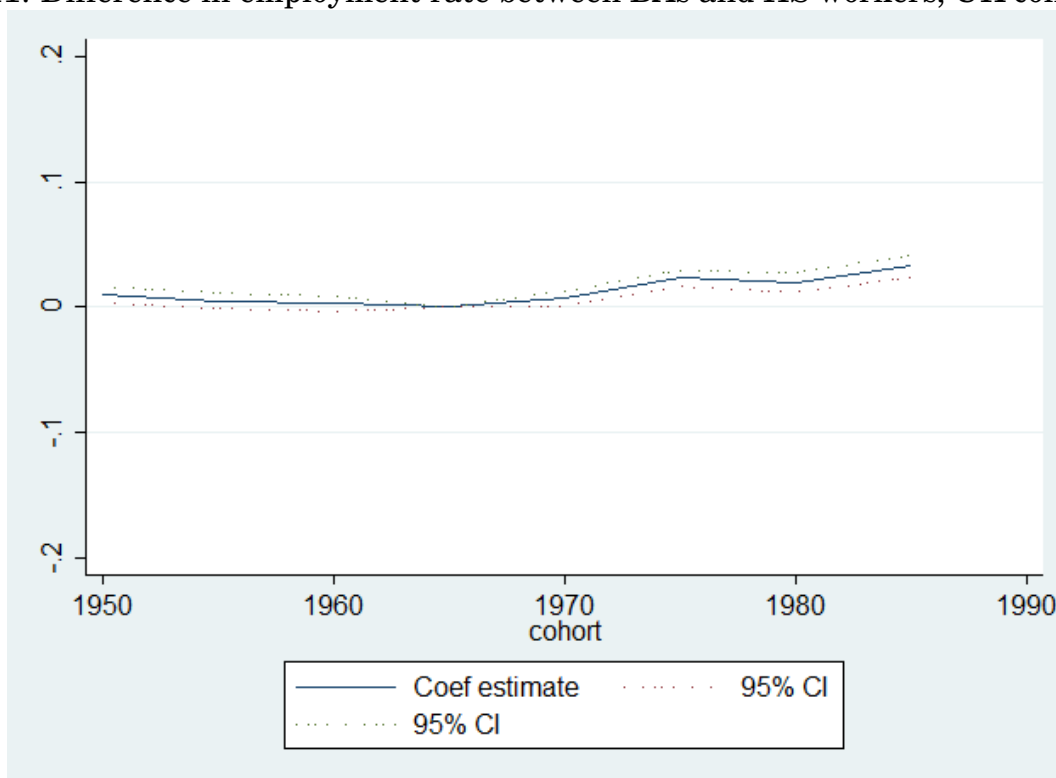
Figure 10: Ratio of median wage to the 10th percentile within education groups, UK



One place we might look for a compositional shift is at the extensive margin: if the

large increase in the relative supply of BAs combined with their constant relative wages induced a relative decline in the employment rate of BAs then this could imply changes in the relative “quality” of BA versus HS workers. In Figure 11, we plot the estimated cohort effects for the difference between the employment rate of BA’s and that of the HS population. The employment rate difference is actually 2 to 3 percentage points higher for the post-1975 cohorts compared to the 1965 base. Thus, the lack of a relative wage response to the educational supply shift was not offset by a relative decline in employment. The change in relative employment rates across cohorts is also small in the context of a 15 percentage point increase in the BA proportion between the 1965 and 1975 cohorts. Thus, we believe compositional shifts based on changes at the extensive margin are not a key driver of the main patterns.

Figure 11: Difference in employment rate between BAs and HS workers, UK cohort effects



2.4.2 Changes in Selection on Unobservables

The exercises we just described are somewhat indirect approaches for capturing changes in the composition of unobservable characteristics. The alternative is to try to address selection on unobservables directly. Two broad approaches are available for assessing the potential impact of selection on the wage patterns we investigate. The first is to estimate selection corrected wage regressions using, for example, a Heckman two-step approach. To be convincing, that approach requires an exclusion restriction in the form of a variable that affects education choice but does not directly determine wages. We do not have a candidate for such a variable in our data and so turn, instead, to a bounding

approach. Under these approaches (Manski [1994], Blundell et al. [2007], Lee [2009]), extreme assumptions on the ability of individuals who would have shifted from lower to higher education across cohorts allow us to put bounds on movements in the median wages within education groups and, hence, on the education wage differentials. We will implement a variant on a bounding approach developed in Gottschalk et al. [2014], who use it to examine movements in occupation task prices.

Implementation of a bounding approach rests on some (preferably minimal) assumptions about the model of wage determination. We will consider a simple but very standard model in which the wage for person i in education group j is given by:

$$\ln w_{ict} = \sum_{j=1}^3 D_{ijt} \beta_{cj} + \sum_{j=1}^3 D_{ij} f_{cj}(age_{it}) + \sum_{j=1}^3 D_{ij} \lambda_j \eta_i + \epsilon_{ict} \quad (3)$$

where c indexes the person's birth cohort, D_{ij} equals 1 if person i is in education group j , and zero otherwise, f_{cj} is a cohort-and-education-group-specific age profile of wages, normalized to 0 for age 30 and ϵ_{ict} is an idiosyncratic error that is independent across time and people and of all other right hand side components in the regression. The specification incorporates a person-specific ability factor, η_i , the effects of which differ across education groups according to loading factors, λ_j . Importantly, both the distribution of η_i and its factor loadings are stationary across cohorts. This model is extreme in its assumption of only one ability factor, but it is also very standard and allows us to see clearly the effects of selection.

We shall assume that the values of the λ 's and other parameters are such that for each cohort, the three education groups correspond to three contiguous, non-overlapping ranges of ability. In particular, the groups are defined by two cohort-specific thresholds A_{uhc}, A_{hdc} . University graduates are those with $\eta > A_{uhc}$; high-school grads have $A_{hdc} < \eta \leq A_{uhc}$; and high-school dropouts have $\eta \leq A_{hdc}$. In theory, such a hierarchical model of selection could be rationalized by a Roy model where individuals choose education levels by comparing their expected net present value of wages and of costs, and assuming $\lambda_u > \lambda_h > \lambda_d$ and that the costs of obtaining education are weakly decreasing in ability. In addition, the hierarchical model fits the idea that university admission is largely rationed by prior attainment.

Consider a situation in which the university proportion increases between cohorts c and $c+1$, because there is less rationing. This corresponds to a decline in the value of A_{uhc} . Importantly, some individuals who would not get a university degree if they were born with their respective ability in cohort c will get a degree if they belong to cohort $c+1$ but no one is induced to make the opposite switch. That is, there will be flows in only one direction. Let's call the set of individuals who would get a degree if they face the conditions in cohort $c+1$ but not if they were in cohort c , "joiners". Their ability distribution has a range with a top value of A_{uhc} and so it lies entirely below that of the rest of university graduates in cohort $c+1$. The latter group have abilities that are high enough for them to enter university even when the costs were higher (as they were for

cohort c). We will call them “stayers”.¹⁰ Obviously, the joiners’ ability distribution lies above that of those who remain in the HS group in cohort c+1.

Given these assumptions, we can obtain bounds on the composition-constant median wage for each education group. The detailed derivation is in the appendix and here we explain the intuition. The core idea is that if we could observe changes in the median wage for a composition constant group then that change would reflect only the change in the price per efficiency unit of workers with a given type of education ($\beta_{cj} + f_{cj}(age_{it})$ in (3)), not composition changes. We can’t uniquely identify a composition constant group but we can bound the wage movements for one: the stayers. The observed wage distribution of BAs in cohort c+1 is a combination of that of the joiners and that of the stayers. Under our assumptions, if the number of BA’s increases across cohorts then that must reflect an inflow of joiners but no outflow. That means we can use the observed median wage for BA’s in the first cohort as corresponding to the median wage of the stayers. In the second cohort, we can form two extremes based on what we assume about the joiners. In the first, we could assume that all the joiners have lower ability than the median stayer. We could then form one extreme estimate of the median wage for stayers by first trimming a number of observations equal to the number of joiners from the bottom of the observed wage distribution for the second cohort and then getting the median of the remaining observations. For example, if the size of the BA group increases from 20 to 30 percentage points of the population between cohort c and cohort c+1 at a given age, then we trim the bottom one third of the BA wage distribution of cohort c+1 and the median of the remaining distribution is the upper bound of the median of the stayers. Another extreme bound could be formed by similarly trimming the top third of the cohort c+1 distribution and getting the median for the remaining sample. However, under an hierarchical model of the kind we are discussing, the best the joiners could be is as good as the stayers (if they were better than the stayers, they would be in the sector already). If they are as good as the stayers then the observed median wage for BA’s in cohort c+1 would be the same as the median wage for the stayers. Thus, the observed median forms the other bound on the cohort c+1 median wage for the stayers. Differencing these extreme bounds for the stayers’ median wage in cohort c+1 from the observed median wage for cohort c then gives us bounds on the movements in the price for BA labour for a composition constant group.

We can implement a similar exercise for the HS group. In this case, though, if the BA group grows between cohort c and c+1 this must be directly matched with an emigration of individuals from the top of the HS ability distribution between those cohorts. In trimming terms, this means that one bound can be obtained by appending a number of workers equivalent to the increase in size of the BA group to the top of the cohort c+1 wage distribution for HS workers. At the same time, if the Drop-out group shrinks then,

¹⁰Calling them stayers and joiners is a slight abuse of terminology since we are considering different cohorts and so there are no individuals actually staying or joining. Instead, these groups correspond to different ranges in the stationary η distribution.

under the single factor Roy model, they must have moved to the bottom of the ability distribution in HS and we would trim a number of workers equivalent to the decrease in size of the Drop-out sector from the bottom of the cohort $c+1$ HS distribution. Doing both the BA and Drop-out related trimming and appending yields a new adjusted HS sample in cohort $c+1$ that corresponds to one bound on the wages for the HS group stayers. Taking the difference between the median wage in that sample and the actual median wage for HS workers in cohort c yields an upper bound on the change in the log wage profile at a given age for HS workers. Consider the benchmark case where the upper bound scenarios for the BA and HS workers correspond to one another (i.e., the movements out of the top of the HS distribution become the movements into the bottom of the BA distribution). We can then obtain one bound on the movement in the university - high school wage differential by taking the difference between the upper bound on the movement in the university median and the upper bound on the movement in the high school median. The other bound is the actual change in the median wage ratios shown in Figure 1.

We repeat the sample trimming exercise for each cohort using the 1965-69 cohort as the base of comparison (cohort c in our example). The left panel of Figure 12 shows age profiles of the ratio of adjusted (upper bound) BA-to-HS wage ratios for each cohort while the second shows cohort effects derived in the same manner as in the earlier figures. The cohort effects show an increase in the adjusted upper bound differential between the 1965-69 and 1970-74 cohorts. Given that the other bound is the actual change in the median wage ratio, the implication is that under this ability model, one cannot argue that selection on unobservables obscured what was actually a decline in the true wage differential. For the difference between the 1965-69 and 1975-79 differential, one bound shows a near zero decline and the other shows approximately a 5 percent decline. Thus, here there is some room to argue that selection is hiding a true decline in the ratio, but that decline is still very small compared to a doubling of the proportion of the population with a BA. For the post-1980 cohorts, the bounds include larger declines - over 10% relative to the 1965-69 cohort at the extreme. However, a glance at the profiles in the left panel suggests the need for some caution in interpreting the cohort coefficients. The age profiles for the various cohorts no longer look parallel once the extreme bound trimming is implemented, implying that the age at which we evaluate the cohort differences can alter our conclusions. But, overall, our conclusion from this exercise is that, under this model of ability, selection on unobservables cannot explain why we do not see a large decline in the education wage differential for the cohorts with the largest increase in their education level.

The previous exercise hinges on the assumed model of selection on ability. We, next, ask how big selection would have to be to account for the differences between the US and UK experiences that we described in the Section 1.3. In particular, we consider a scenario in which the university joiners' wages are always above the median of the HS stayers (as in the extreme version of the hierarchical model) but could be above or below the median of university stayers. The more the university joiners have wages above

the university stayers' median, the more will selection tend to induce an increase in the observed BA-HS wage differential. The question we pose is the following. Suppose that the UK actually faced the same skill biased demand shift as the US in this period, what proportion of the above-average HS workers who join the BA group would also have to be above the median of the BA group in order to rationalize the observed wage movements in the UK? To answer this question, we use our estimates from section 1.3 of the relative demand shifts for BA versus HS workers from the CL model with the seventh order polynomial fit using US data in our time period. Using those demand shifts along with the observed relative supply changes, we impute predicted wages for groups defined by year and 5-year wide age groups under the assumption that the UK faced the same relative demand shifts as the US.¹¹ Comparing the predicted wages with the observed wages in each age group by year, we back out the proportion of switchers that must have earned more than the latent BA median in order to rationalize the difference.¹²

Figure 13 contains this proportion plotted by age group and year. The first thing to note is its magnitude: 20–40% for 25-29 year olds in the early 2000s. Recall that we have made these calculations based on the assumption that all the university joiners would have higher wages than the median HS stayer. If we were to relax that assumption then we would need even larger proportions of the joiners being above the median BA. Depending on how much we relax this assumption, we would likely require over 50% of new university joiners to be of above median university ability.¹³ We view a claim that at least 20% and possibly over half of the new university joiners would have wages placing them above the median of those who were already going to university to be implausible - though, admittedly, we have no way of supporting that reaction. More importantly, the proportion varies widely by age group and year, even falling into negative territory for the 30s age groups in the most recent years.¹⁴ That variety suggests that it is difficult to explain the data patterns with a standard selection model. The unobserved compositional changes by education group would have to not only be substantial and in the right direction, but also vary substantially by year and age group in just the right way to rationalize the different US and UK wage and employment patterns. While we cannot definitively reject such a possibility, we do not know of any theory or intuition that implies those compositional changes. We prefer to look for an economic model that

¹¹The age groups we work with are 20-24, 25-29, ..., 55-59.

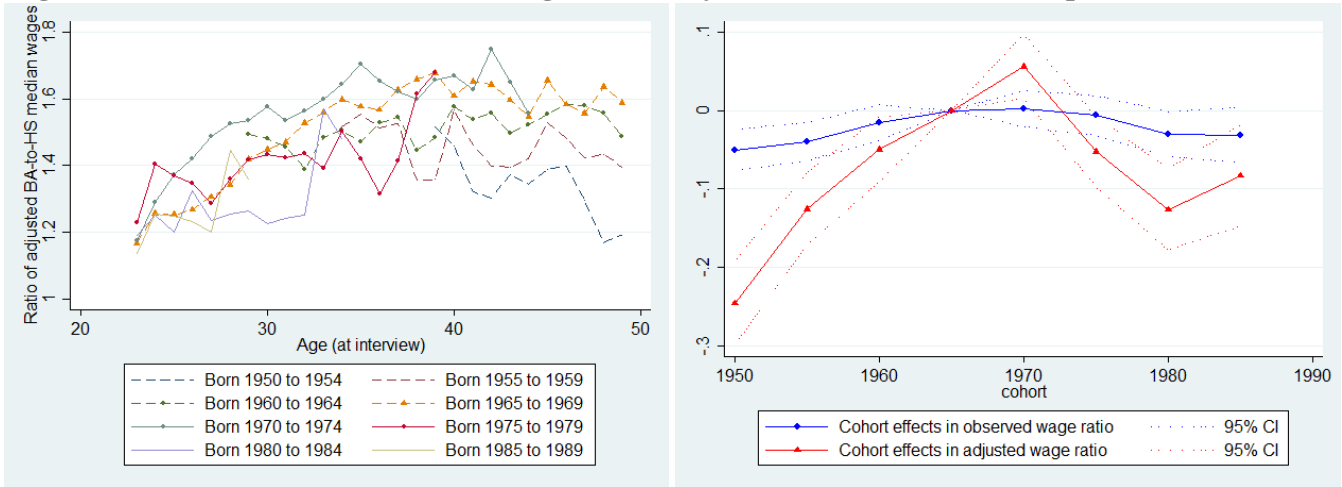
¹²We only do this calculation in age-year cells where the BA proportion is at least 5 percentage points higher than 1993, which is the reference year. When the BA proportion is too similar to the reference level, we have a problem of a small denominator and the imputed proportion can be outside the [0,1] range due to measurement error.

¹³In addition, our demand shift estimates are based on an assumption that the US wages do not reflect similar selection related to increases in the US proportion of workers with a BA. If there were such selection in the US then we would need to have even greater selection in the UK to account for differences between the two countries.

¹⁴The imputed proportion is negative when the predicted latent BA median is so high relative to the observed BA wage distribution that the density above the predicted median is less than half of the original proportion of BAs.

can generate the wage and employment patterns endogenously rather than rely on the existence of a whole vector of exogenous changes that just happen to fit the patterns.

Figure 12: UK Median BA-to-HS wage ratio, adjusted to the education split of 1965 cohort



Note: for each age and cohort, we adjust the wage distribution by using the proportions observed for 1965 cohort as reference points. For example, if the observed proportion of BAs is higher than that for the 1965 cohort at the same age, we would trim the bottom of the observed BA distribution.

Figure 13: implied proportion of switchers ending above the latent BA median



3 Plausible theories

At this point, we believe we have established two main points. First, that the UK underwent a very large expansion of the proportion of workers with a BA without an accompanying change in the BA-HS wage differential or substantial changes in relative employment rates. Second, that these patterns cannot be explained by composition shifts based on observable or unobservable characteristics. The main caveat to these claims is that they seem clearly true for the 1965-69 to 1975-79 birth cohorts (the ones experiencing the largest increases in education level). But there is some evidence from both restricting the analysis to the private sector and the bounds on the selection on unobservables to suggest that the wage differential may have declined for the post-1980 birth cohorts that entered the labour market in the mid-2000s. We now turn to evaluating quantitatively three competing theories as to how the UK economy absorbed such large increases in educated workers without much reduction to their relative wages.

3.1 Trade

In a standard Heckscher-Ohlin model with three production factors: BA labour, HS labour and capital (assuming capital is perfectly elastically supplied), the ratio of wages will be invariant to shifts in the domestic relative supply of the two types of labour. The economy would respond to changes in relative supplies by shifting production toward goods that are more intensive in the growing factor. The aggregate increase in BAs would be absorbed through an expansion of industries or sectors that tend to employ a highly educated workforce, and perhaps at the expense of industries that use more low-skilled workers.

However, the proportion of workers who have a BA actually grew very rapidly in all major industries in the UK (Table 3.1). Formally, we decompose the growth in the proportion of workers with a BA into between versus within industry components. During the period 1994-2014, the total within-industry component accounts for a 19.1 percentage point increase in the overall proportion of workers with a BA and the between-industry component accounts for a 2 percentage point increase. We can also do the decomposition analysis for a finer definition of cells: the interaction of industry(1 digit of SIC1992), sector(public or private) and occupation (three categories). We do so for separate subsamples by age and gender. As shown in table 2, the between effect is much smaller than the within effect for every subsample. Our conclusion is that the relative wage invariance is not due to trade-related responses; at least, not operating at the level of industrial aggregation in our data.¹⁵

¹⁵Dustmann and Glitz [2015], using German firm level data, show that approximately 80% of adjustment to an immigration-induced change in local labour supply occurs through adjustments in employment by skill type within firms as opposed to changes across firms. They show some sensitivity to aggregation to the industry level but not enough to alter the core result that most adjustment is not about shifts in output composition.

Table 1: Share of BAs among 16-59-year-old employees by industry SIC1992, 1994-2014

SIC92	shares		from 1994 to 2014	
	in 1994	in 2014	within effect	between effect
Agriculture, mining and fishing	0.075	0.258	0.002	0.000
Manufacture	0.082	0.245	0.027	0.006
Energy	0.177	0.289	0.001	-0.000
Construction	0.066	0.188	0.006	-0.000
Wholesale, retail	0.043	0.173	0.019	0.002
Hotel and restaurants	0.032	0.158	0.007	-0.002
Transport, storage, communication	0.066	0.234	0.010	0.000
financial intermediation	0.112	0.429	0.015	-0.000
real estate, business activities	0.225	0.485	0.028	0.006
public administration	0.151	0.391	0.017	-0.000
education	0.414	0.582	0.016	0.008
health and social work	0.111	0.366	0.033	0.000
other community, social	0.146	0.321	0.008	0.000
private and external	0.059	0.373	0.001	0.000
Total	0.124	0.336	0.191	0.020

Note: the within-industry effect of an industry is defined as $(x_1 - x_0)(w_1 + w_0)/2$ and the between-industry effect of an industry is defined as $(x_1 + x_0 - \bar{x}_1 - \bar{x}_0)(w_1 - w_0)/2$, where subscripts 0 and 1 denote 1994 and 2014 respectively, w denotes the industry’s share of total employment, x denotes the share of BAs in the industry’s employees and \bar{x} denotes the aggregate share of BAs.

3.2 Exogenous Technological Change: Polarization Version

In section 1.3, we showed that the same model of exogenous technical change could not fit both the US and UK data in our period. The version of the SBTC model we used is one that [Acemoglu and Autor \[2011\]](#) describe as the “canonical model” in which technical change is assumed to increase the relative productivity of more educated workers. But several authors have argued that version of the exogenous SBTC story does not fit recent data even for the US ([Card and DiNardo \[2002\]](#) and [Beaudry and Green \[2005\]](#)). Partially in response to this criticism, a more nuanced version of the SBTC theory has emerged emphasizing the differential impact of Information Technology (IT) on different sets of tasks. In particular, IT is argued to be complementary with cognitive tasks but to replace routine tasks for which computer equivalents are easy to programme. Manual or service tasks are neither complementary with or substituted by IT but are often argued to experience increased demand either because they are demanded by the now richer cognitive workers or because of consumption variety related arguments (e.g., [Autor and Dorn \[2013\]](#)). With cognitive, routine, and manual task occupations generally occupying

Table 2: Within and between components in the increase of BAs, 1994-2014

Gender	age range	share in 1994	share in 2014	within effect	between effect
both	16-59	.125	.336	.17	.0409
both	25-34	.159	.442	.229	.0547
both	35-49	.142	.36	.182	.0366
both	50-59	.0863	.259	.134	.0386
female	16-59	.106	.35	.198	.0457
female	25-34	.151	.492	.279	.0611
female	35-49	.111	.371	.213	.0452
female	50-59	.062	.251	.146	.0389
male	16-59	.141	.322	.142	.0385
male	25-34	.165	.4	.184	.0506
male	35-49	.168	.351	.151	.0307
male	50-59	.108	.267	.123	.0356

Note: the cell is defined by full interactions of sector, industry and occupation. The within effect of a cell is defined as $(x_1 - x_0)(w_1 + w_0)/2$ and the between effect of a cell is defined as $(x_1 + x_0 - \bar{x}_1 - \bar{x}_0)(w_1 - w_0)/2$, where subscripts 0 and 1 denote the start year and the end year respectively, w denotes the cell's share of total employment, x denotes the share of BAs in the cell and \bar{x} denotes the aggregate share of BAs. The within and between effects shown above are the sums across all cells.

the top, middle and bottom of the earnings distribution, the implication is that there will be a polarization of employment, with growth in employment at the top and bottom of the distribution relative to the middle.

Does the task based model affect our earlier conclusions about whether a SBTC story can fit both the US and the UK wage and education patterns? It is worth noting that [Acemoglu and Autor \[2011\]](#) show that with two skill levels and an assumption that the higher skill level has comparative advantage in more complex tasks, one can derive an empirical specification that is isomorphic to the canonical model specification that we estimated. In this sense, the conclusions of our earlier exercise continue to hold. Put another way, suppose that the observed stability of the UK education wage differential arises because of a polarization in demand with university workers disproportionately benefiting from increased demand for cognitive tasks and high school workers benefiting from a combination of moving out of routine tasks into cognitive tasks ([Cortes \[2016\]](#)) and increased wages for manual workers. This still doesn't solve the question of why this doesn't happen to the same degree as in the US. That is, in the polarization model as well, one would need UK specific demand shocks that just happen to match the supply shifts in just the right way to leave the education wage differential unchanged.

To look further into the role of polarization in the UK wage and employment struc-

ture, in Table 3 at age 30-34, we present average real wages (in the first column of the first panel) and proportions of employees (in the first column of the second panel) in each of 9 one digit occupations in the 1965-69 cohort. The occupations are ranked by their average real wage. In the second columns in each panel we present the actual change in either wages or proportions between the 1965-69 and 1975-79 cohorts. The second column for employment proportions shows an approximate U-shaped pattern, with growth in employment shares in the top three occupations, declines in the middle (largely routine) occupations and growth in personal services. The relationship is not perfect since the two occupations paid below personal services both show declines, but the pattern is broadly one of polarization across successive cohorts. However, when we hold the education composition constant between the cohorts, there is either negative or very small growth in employment in the top occupations and no change in processing and skilled trades in the middle. There is some added evidence of relative growth at the bottom of the distribution. The main conclusion, however, is that the right branch of the U-shape in employment growth in the UK across these cohorts is almost entirely due to the education shifts. That is, occupation shifts appear to us to be of secondary importance relative to education shifts in determining the changes in the wage structure across cohorts in the UK. Given that, we do not believe that polarization/task based versions of the exogenous skill biased technical change theory will provide a useful lens through which to understand the specific wage and employment patterns we are examining.

Table 3: Changes between 1965 and 1975 cohorts, by occupations, at age 30-34,

occupation	mean real wage		employment shares		
	w_{1965}	%change observed	$share_{1965}$	change observed	change reweighted
Professional occupations	18.70	0.067	0.128	0.059	-0.011
Managers and senior officials	17.66	0.074	0.159	0.010	-0.008
Associate professional and technical	15.96	0.087	0.159	0.035	0.008
Administrative and secretarial	11.94	0.021	0.133	-0.035	-0.026
Skilled trades	11.67	0.052	0.122	-0.027	0.001
Process, plant and machine operatives	10.26	0.041	0.081	-0.028	-0.007
Personal service	9.00	0.054	0.061	0.013	0.027
Sales and customer service	8.73	0.130	0.062	-0.009	0.003
Elementary occupations	8.06	0.043	0.095	-0.018	0.013

Note: Real wage is in 2012 prices, deflated by GDP deflator. The final column reweights the employment shares of occupations using the education split of the 1965 cohort of 30-34 year olds.

3.3 Endogenous Technological Change

The idea that technological innovations or adoptions depend on relative wages or skill supply has been around for some time [Acemoglu, 1998, Beaudry and Green, 2003]. Intuitively, an increase in the relative supply of skills can induce firms to adopt or invest in technologies that are skill-biased. For the US, Beaudry and Green [2005] show that such a model can capture national-level patterns that the canonical model misses, and Beaudry et al. [2006] show that IT adoption occurred first in cities with already high proportions of workers with BAs. For Norway, which experienced a substantial expansion of higher education in the 70s, Carneiro et al. [2014] found that the skilled wage premium increased in the areas most affected by the higher education expansion relative to other areas. They interpret these findings through a model of directed technical change in which firms responded to an increase in skill supply by adopting the skilled-biased technology. This is similar in spirit to our exercise except that they present a model with a threshold education level beyond which all firms shift technologies and discuss differences before and after the technological shift. In contrast, we emphasize a period of transition between two technologies. Results in the literature on the impact of immigration on wages also fits with a model of endogenous technological adjustment [Lewis, 2011, Dustmann and Glitz, 2015].

As the focus of our paper is to explain the lack of change in wage differentials when higher education expanded dramatically, we propose a model of endogenous technological change that allows for such a lack of change. As shown below, in our model, the relative demand shifts that would be required to offset the large UK relative supply shifts in order to leave the wage differential unchanged can arise endogenously as part of a transition to a skill-biased technology; a transition that is induced by the increase in the supply of BA's. In this sense, the model can explain the main facts we have presented as an outcome of the operation of the economy rather than as an exogenous event.

Our preferred representation of the skill-biased technology in this period is as a decentralized organizational form. Caroli and Van Reenen [2001] argue that firms in recent decades have been transitioning from a more traditional organizational form with rigid hierarchies to one with more decentralized decision making, delayering of managerial functions, and multi-tasking. They also show that the newer organizational form is a complement to education or skill. It seems plausible, in turn, that as the UK expanded its proportion of university educated workers, it may have undergone an accompanying shift in organizational form.

Formally, consider a model of production in which there are two existing technologies, centralized (C) and decentralized (D); and two tasks performed by workers: managerial (M) and basic labour (L). The two technologies produce outputs of the same good, Y_C and Y_D , and the good has price 1. Output from the centralized technology is produced according to the Cobb-Douglas production function:

$$Y_{Ct} = M_{Ct}^\alpha L_{Ct}^{1-\alpha} \quad (4)$$

where, α is a parameter. Similarly, output from the decentralized technology is produced according to:

$$Y_{Dt} = M_{Dt}^\beta L_{Dt}^{1-\beta} \quad (5)$$

We will assume that the decentralized technology is more managerial task-intensive, that is $\beta > \alpha$. Intuitively, firms can choose between a centralized technology, e.g., with one manager and many workers, or a decentralized technology with the production process broken up into many teams with more managers and fewer workers per manager.

Each task is performed by a combination of skilled and unskilled labour, with the labour aggregated through CES functions:

$$M_j = [aS_{Mj}^\sigma + (1-a)U_{Mj}^\sigma]^{1/\sigma}, j = C, D \quad (6)$$

and

$$L_j = [bS_{Lj}^\rho + (1-b)U_{Lj}^\rho]^{1/\rho}, j = C, D \quad (7)$$

where, $\frac{1}{1-\sigma}$ is the elasticity of substitution between skilled and unskilled labour in managerial tasks; $\frac{1}{1-\rho}$ is the elasticity in labouring tasks; a and b are parameters; S_{Mj} is the amount of skilled labour in the managerial task in sector j; U_{Lj} is the amount of unskilled labour in the basic labouring task in sector j, etc.. We assume that skilled labour is relatively more productive in the managerial tasks (i.e., $a > b$) and that skilled and unskilled labour are more substitutable in the labouring task (i.e., that $\rho > \sigma$).

We assume that the numbers of unskilled and skilled workers in the economy are given exogenously in any period and that each worker supplies a fixed endowment of labour inelastically (where the endowment is specific to their age and education group). Thus, the amounts of effective labour of each skill level in each sector are given by:

$$S_{kj} = \sum_a \theta_a^S q_{kja}^S, U_{kj} = \sum_a \theta_a^U q_{kja}^U, j = C, D, k = M, L$$

where q_{kja}^S is the number of skilled workers of age group a employed in task k in sector j and q_{kja}^U is defined analogously for unskilled workers. The θ_a^e parameters correspond to the number of efficiency units provided by a worker in education group e and age group a.

Market clearing in the labour market corresponds to the total number of workers with education e and age a in the economy being equal to the sum of the numbers employed in the various occupations and technologies:

$$Q_a^e = q_{MCa}^e + q_{MDa}^e + q_{LCa}^e + q_{LDa}^e$$

where q_{kja}^e is the number of workers with education e employed in task k with technology j. Thus, workers of each age-education type of worker are divided among managerial and labouring tasks in the two technologies. Because workers can choose freely among technology/task combinations, there will be one wage per efficiency unit for each type of worker: w_a^e . Because of the assumed substitutability between age groups, there are only two actual efficiency unit wages: w^S, w^U for one unit of skilled and unskilled labour

respectively. The observed wages for workers are the product of that effective wage times their endowment of effective labour $w_a^e = \theta_a^e w^e$. Finally, the total supply of effective units of skilled labour S is $\sum_a \theta_a^S Q_a^S$, and the total supply of effective units of unskilled labour is $U = \sum_a \theta_a^U Q_a^U$.

Total output in the economy corresponds to $Y_t = Y_{Dt} + Y_{Ct}$, and efficiency corresponds to allocating S and U among managerial and labouring tasks in the two technologies to maximize Y_t . The assumption that the two technologies produce the same final good is intended to capture the idea that these are General Purpose Technologies which can be used within any firm in any industry in the economy, as opposed to technologies related to specific sectors.

3.3.1 Theoretical Implications

We present the full derivation of results for the model in Appendix but here, for brevity and because this model has quite a standard form, we just describe its main theoretical and empirical implications. The set-up in the model is directly analogous to a 2x2 trade model with the two constant returns to scale technologies taking the place of the two goods in that model. Three standard results from the trade model are useful for us. First, the two technologies will be in use at the same time if the relative factor amount $\frac{S}{U}$ lies within the cone of diversification, i.e., if $\frac{S}{U}$ isn't so large that only the D technology is used or so small that only the C technology is used.¹⁶

Second, it is straightforward to show a factor price invariance result: when the economy is in the cone of diversification, both the skilled wage and the unskilled wage, w^S and w^U , are invariant to the factor quantity ratio, $\frac{S}{U}$.¹⁷ This means if the education composition of one age group shifts and if the resulting $\frac{S}{U}$ remains in the cone, then the skilled-to-unskilled wage ratio within each age group will remain unchanged. More than that, both the skilled and the unskilled wages, individually, will not change. Eventually, if $\frac{S}{U}$ rises enough, the economy will pass out of the cone of diversification and only the decentralized technology will be used. At that point, further increases in the relative skill supply will reduce the skilled wage premium. Under the specific form of the model here, once the economy is outside of the cone of diversification, the skilled-to-unskilled wage ratio will fall by the same proportion within all age groups, not just in the age group that experienced the education expansion.

The third trade model result is that, while the economy is still in the transition period, an increase in $\frac{S}{U}$ will imply an increase in $\frac{Y_{Dt}}{Y_{Ct}}$, i.e., that relatively more output will be produced with the skill and managerially intensive Decentralized technology. At the same time, an upgrading of the education level in the economy (i.e., an

¹⁶The Appendix derives the range of $\frac{S}{U}$ within which both technologies are used.

¹⁷With perfect competition, for each technology we have $C_j(w_S, w_U) = 1, j = C, D$ where $C_j(w_S, w_U)$ is the cost of producing one unit of the output using technology j . For both technologies to be in use, both must simultaneously satisfy this condition. Thus, we have two equations in two unknowns (w_S and w_U) which yields solutions for w_S and w_U that do not depend on $\frac{S}{U}$.

increase in S and reduction in U), will not change the ratios of inputs within each sector ($S_{Mj}, S_{Lj}, U_{Mj}, U_{Lj}$, for $j = C, D$) and so the per capita productivity within each sector will not change. As the Decentralized sector has higher labour productivity, the aggregate labour productivity will increase as the Decentralized sector expands while the Centralized sector shrinks.

3.3.2 Empirical Implications

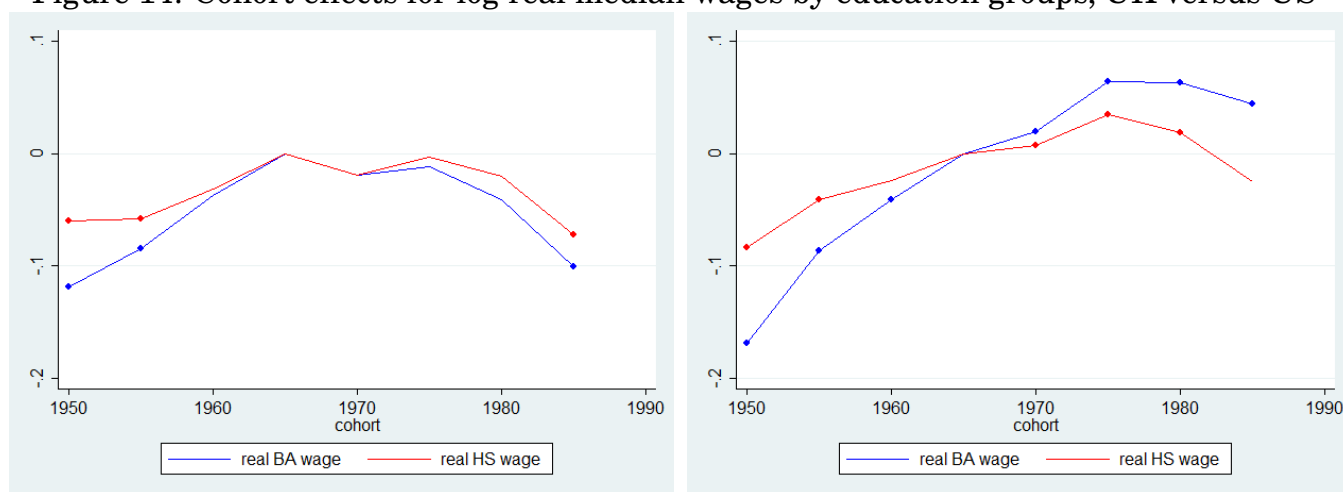
The second of the trade model-related implications (that the skilled wage premium is invariant to movements in the ratio of skilled to unskilled workers) is what makes this model appealing in our context. As we have documented, for the 1970-74 and 1975-79 UK birth cohorts, substantial increases in the proportion of workers with a BA or higher degree were accompanied by an unchanging skill premium. More importantly, the left side of Figure 3.3.2 shows that the stronger implication from the model, that the high and low education wages, individually, do not change while the economy is in the technology transition period also holds. In this figure, we plot the cohort coefficients from our standard specification run separately for high school educated and BA educated log real wages¹⁸, normalized to zero for the 1965-69 cohorts. In the figure, a symbol on the line indicates that the estimated difference between the effect for that cohort and the base (1965-69) cohort is significantly different from zero at the 5% level. Examining the figure, we see that between the 1965-69 and 1975-79 cohorts, the wages for both education groups are virtually unchanged and wage changes are not statistically significantly different from zero for these cohorts or for the 1980-84 cohort. In contrast, in a model with an exogenous skill biased demand shift, the real wages should increase over time. As we have seen, such a model would also require a set of technology shocks that just happen to be the right size in each cohort to offset the relative supply shifts at that point. In a model of endogenous technical change like ours, shifts in the supply of skills generate the changes in demand for skills in an exactly offsetting manner. This happens because of firms choosing between the more and less skill intensive technologies.

We present cohort effects by education group for the US on the right of Figure 3.3.2. Here the picture is quite different. For the US, the real wage for the BA educated rises faster than that for the HS educated across all cohorts up to the 1975-79 cohort. For the remaining cohorts, real wages for both education groups decline, with the high school wage falling faster. We believe that one can reconcile the wage and employment patterns for the US and the UK if we see the US as the technological leader, investing in generating the new technology in the 1980s and 1990s, while the UK is the technological follower, adopting the new technology only once the education level of its workforce makes it profitable for it to do so. The US was the technological leader relative to the UK because it attained higher education levels earlier. This fits with [Beaudry et al. \[2010\]](#) who show that the adoption of new computer technology was faster in US cities with higher initial education levels and that the return to education rose faster in those

¹⁸We deflate the nominal wage series using the CPI with 2012 = 100.

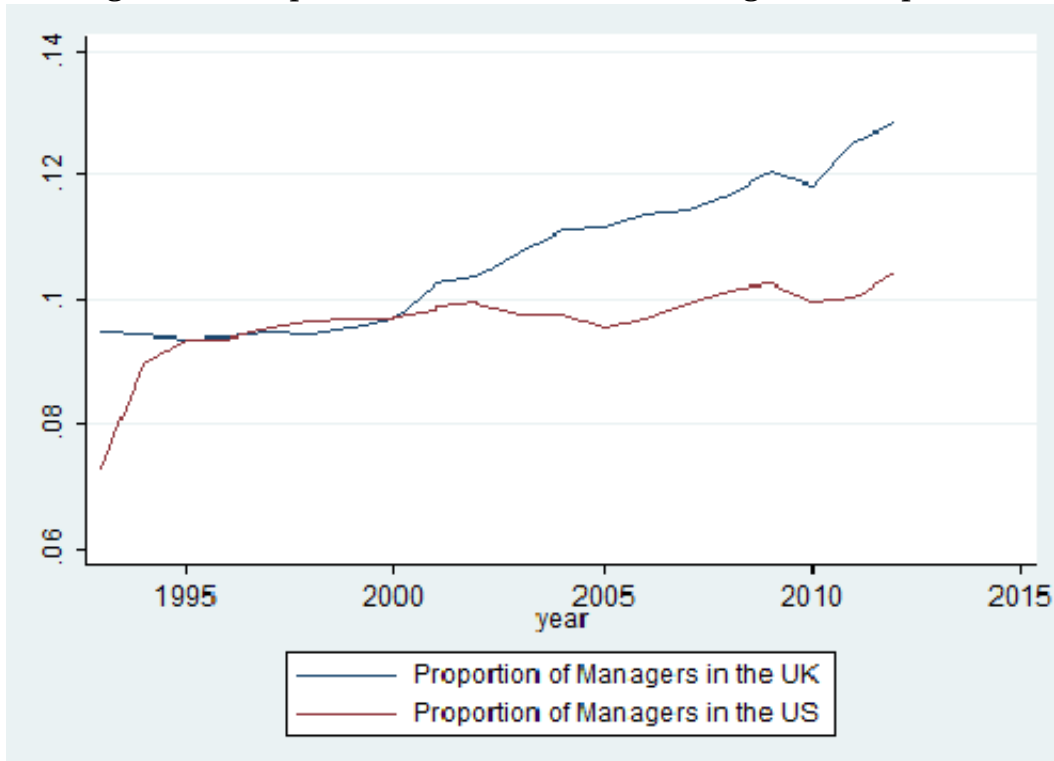
cities with the advent of the new technology than in other, lower education cities. The turn in the US wage series after the 1975-79 cohort in Figure 3.3.2 fits with the argument in Beaudry et al. [2013] that after the period of investing in initiating the new technology, the US economy entered a period with lower demand for skills relative to the investment period. They argue, further, that this decrease in demand for skills generated a movement of educated workers down the occupational structure, causing a supply effect in the occupations usually filled by low educated workers that could have driven the wages in those occupations down even further. They point to 2000 as the turning point between the investment and post-investment period and provide some evidence that investment in computer software and hardware turned at this point also. In terms of our cohorts, the 1975-79 cohort would have been entering the labour market in approximately the late 1990s so that the turn in our cohort plot fits with their turning point. Finally, Bloom et al. [2012] provide direct evidence of the US being the leader in innovation and implementation of a skill-intensive, decentralized organizational form. In particular, they show that US multinationals operating in the UK were more likely to use decentralized forms than UK firms operating in the same market.

Figure 14: Cohort effects for log real median wages by education groups, UK versus US



Note: The left subfigure is for the UK and the right one for the US. For each country, we regress age-cohort-education level wages on cohort dummies and 5-year-age-band dummies separately for the BA and the HS workers. The 1965 cohort is the reference cohort. The marked data points are significantly different from zero at 5% level.

Figure 15: Proportion of BA workers in managerial occupations



Our model can also rationalize the wage movements at the end of our set of cohorts. As we argued earlier, if the rapid increases in the proportion with a BA meant that the UK ultimately moved through the cone of diversification to complete adoption of the Decentralized organizational form then we should expect, eventually, to see further increases in the proportion BA associated with declines in the education wage differential. This seems to have started to emerge in the private sector in the last few years of our sample. The timing is not implausible since the UK has already surpassed the US in terms of the BA proportion among young people, and will soon catch up in terms of the aggregate BA proportion. In short, the model predicts the seemingly odd pattern in which the education differential is flat when the increase in education was largest, starting to turn down for later cohorts that experienced smaller educational upgrading.

While this match between empirical patterns and theory is intriguing, it is far from definitive. To examine the suitability of the model further, we look at the third implications: increases in the relative supply of skills should induce a shift toward a more decentralized organizational form. Direct evidence on this can be found by examining changes in the proportion of workers who list “manager” as their occupation. Figure 15 shows the change in the proportion of workers with a BA who are in management occupations relative to our first data year (1993) for both the US and the UK. The US shows some increase at the start of the period (possibly related to their adoption of the new organizational form), while that in the UK increases by more, and does so later, fitting with our notion that the UK lagged the US in its adoption of the new form.

To get more specific evidence, we examine employees’ self-reported level of influ-

ence at their workplaces in a micro dataset as a marker of a decentralized managerial structure. Our hypothesis is that in a more decentralized and de-layered organizational structure, workers will be given more autonomy and will report greater influence over their work. We are interested in whether an increase in the relative supply of education skills induces a shift toward a more decentralized organizational form as measured by this marker. We examine this question using the UK Workplace Employment Relations Survey (WERS). The WERS is a survey of workplaces that includes questionnaires both for the manager as well as for a subsample of employees.¹⁹ We focus on employees' responses to three questions:

“How much influence do you have about the following?”

1. “The range of tasks you do in your job”,
2. “the pace at which you work”
3. “how you do your work”.

The responses for each question range from 1 “A lot” to 4 “None”. These questions are included in the cross-sectional WERS surveys for 1998, 2004, and 2011. Rather than use these questions separately we implement a principal components analysis to compute an index of the ability of workers to influence their own work. We define the index as 4 minus the first principal component, so that the index is higher where more employees report having a lot of influence. The index accounts for approximately 80% of the total covariance among the three questions.

To investigate the role of skill supply in choice of organizational form, we examine the relationship between the local supply of workers with BAs and the average value of our influence index at the area level. Area here refers to Travel To Work Areas (TTWA), which were developed to capture local labour markets using data on commuting flows in 1991.²⁰ There were around 300 such areas in the UK in the 1998 through 2011 period. We collapse all three waves of WERS to this TTWA level to form a short panel. Meanwhile, we derive from the LFS the proportion of workers in the TTWA who have a BA or above for the two calendar years up to and including the WERS survey year.²¹

Table 3.3.2 reports the results from regressions of the influence index on the local BA proportion across a range of specifications. In all the specifications, we pool together the data from the three waves and we weight by the size of the area as measured by employment in the LFS data. In the first column, we report the results from an OLS regression with the proportion of BA's in the area and year dummy variables as the only regressors. The estimated year effects indicate a secular trend toward organizational

¹⁹The WERS surveys 25 employees per workplace. When there are fewer than 25 employees at the workplace, they are all given the questionnaire.

²⁰For further information on TTWA, see <http://www.ons.gov.uk/ons/guide-method/geography/beginner-s-guide/other/travel-to-work-areas/index.html>

²¹For example, for the WERS outcome measured in 2011, the BA proportion is measured from LFS 2010-2011.

forms with greater worker control. This may reflect a response to the general increase in the education level of the workforce but more direct evidence on whether such a relationship exists is found in the estimated effect of the proportion of workers with a BA. We estimate that a 10 percentage point increase in the proportion of BAs in an area is associated with a 0.06 increase in the influence index. This result fits with the idea that firms in areas with a higher proportion of educated workers use more decentralized organizational forms. The estimated magnitude of the effect is not trivial. The employment-weighted average local BA proportion increased by 13 percentage points over the sample period. With a coefficient of 0.6, the 13 percentage point increase translates into a 0.08 increase of the index, which is about a third of the standard deviation in 1998.

In the next set of columns, we check the robustness of this result across a series of specifications. In the second column, we introduce controls for the proportion of workers by industry, by workplace size, and by the size of the organization.²² Notably, the size and significance of the BA proportion coefficient remains very similar to what was observed in column 1. This implies that the association between the level of education of the population and the organizational form happens within industries (as one would expect with a General Purpose Technology) rather than through shifts in the industrial structure. In column 3, we repeat the specification including the same controls but drop the observations from London out of concern that these could be driving our results. Again, the size and significance of the estimate is unchanged.

Whether the estimated association between the BA proportion and the average influence index value in these regressions represents a causal effect of the level of education is unclear. More educated workers may migrate to areas where firms have more decentralized organizational structures, implying a reverse causality. Alternatively, there could be a third unobserved factor prevalent in some areas that both increases the attractiveness of using a decentralized form and is attractive to more educated workers. We find it difficult to determine what form such a factor would take given that we are already controlling for industrial structure and firm size. In addition, the fact that our results hold up when we drop London (which is a strong candidate as a place where more educated workers migrate to with the aim of working for the most up-to-date firms) is weak evidence against the first endogeneity channel. Nonetheless, we are concerned that there is remaining endogeneity, which we address in two main ways. The first potential response is to take advantage of the panel nature of the data to introduce area specific fixed effects. These would control for persistent features of a location that might affect both its education level and the technology choice of firms residing in the area. However, when we estimate a version of our regression in first differences, the sign on the BA proportion variable continues to be negative but both it and the other included

²²More specifically, industry is measured by the first digit of Standard Industrial Classification 1992; we have 5 categories of workplace size: ;25,25-49,50-249,250-999,1000+. Whereas workplace size refers to the number of employees at the specific site, the organization may have multiple sites and hence many more employees. We have 5 categories of organization size: ;50,50-249,250-999,1000-9999,10000+.

controls have very large standard errors. There does not appear to be enough within-area over-time variation to identify the effects we are trying to investigate. In place of a fixed effect estimator, we implement a pooled OLS specification in which we include the proportion of workers in an area who had a BA in 1995/6. One can interpret this variable as a parameterization of location fixed effects that uses only the part of the fixed effect that is correlated with the historic education composition. This allows us to control for persistent factors that have been directly related to education. We present the results from the specification including start of sample education composition in column 4. The start of sample BA proportion, itself, enters with a small effect that is much smaller than its standard error. The effect of the current BA proportion continues to have the same size and significance, suggesting that concerns about endogeneity related to historic patterns of education are not warranted.

Our other approach to potential endogeneity problems is to implement an instrumental variables (IV) estimator. In particular, we make use of variation across areas that relates to the expansion of education. Thus, one possible instrument is the change in the proportion of workers with a BA between the cohort of workers born in 1965-69 and those born in 1975-79 - the cohorts defining the largest increase in higher education. This instrument would be valid to the extent those increases arose from exogenous differences in the expansion of the higher education system. However, since we see workers at older ages, the instrument might still reflect differential mobility of high education workers across cohorts - potentially in response to the location of firms with more decentralized organizational structures. In response to this, we use as instruments the proportion of the population born in the years 1970-74 and the proportion born between 1975 and 1979, measured in 1995-96.²³ The underlying idea is that the proportion of the population with a university degree expanded substantially for the 1970s cohorts. As a result, areas with a high concentration of people of university age at the time of the expansion in the higher education system would be predicted to have a more educated population later to the extent that people have some tendency to stay where they grew up. The requirement that this instrument is valid is that parents in the previous generation did not have a tendency to have more children in areas which would later turn out to have more decentralized organizational structures. We believe that it is likely that this requirement was met. Moreover, the first stage for this instrument is strong: the F-stat for the joint significance of the two IVs is 17.9. We also report a specification in which we use these demographic variables along with the change in the BA proportion between the 1965-69 and 1970-74 cohorts and between the 1970-74 and 1975-79 cohorts as instruments. Together, these four IVs are highly significant in the first stage, having an F-stat of 63.1.

Column 5 of the table contains the results from using the demographic variables alone as instruments. The BA proportion effect remains positive and highly statistically significant. In size, it is approximately double that obtained using OLS, implying that

²³The denominator for the proportions is the population born between 1940 and 1979.

the 13 percentage point increase in the percentage of workers with a BA resulted in a two-thirds of a standard deviation increase in the personal influence index. The result is very much the same in column 6 where we use both demographics and the change in the proportion with a BA as instruments.²⁴ The fact that the effect estimated by IV is larger than the OLS estimate suggests that we may be estimating a local average treatment effect where there is heterogeneity in area responsiveness in terms of organizational structure to education. This could fit with the idea that the decentralized structure has the quality of a new General Purpose Technology. Firms in areas with historically high levels of education may have had little impetus to shift to the new organizational structure since the general increase in education would have had little impact on them. In contrast, firms in the areas where levels of education increased the most may have faced more incentives to make the organizational shift. Thus, the results reflect the interaction of a relatively newly available organizational structure with a recent increase in education. But regardless of the explanation for the relative sizes of effects from different estimators, the overall message is clear. An increase in the proportion of workers with a BA causes an increase in the extent to which workers feel they control their own work. This fits with results in [Caroli and Van Reenen \[2001\]](#) where they use UK and French data to show that a relative shortage of educated workers in a local labour market, as reflected in a higher education wage differential, implies that the firms in that market are less likely to implement organizational change. We view these results and ours as corroborating evidence for our model in which the large increase in the education level of new cohorts born after the late-1960s generated a shift in organizational structure toward a more decentralized structure in which workers had more control over their own tasks. As we have seen, in such a model, the technological shift can be accomplished without a change in the wage differential between more and less educated workers.

4 Conclusion

This paper highlights two empirical patterns: first, the UK underwent a dramatic increase in the proportion of people with a BA across successive cohorts beginning with those born in the late-1960s; second, the age profile of the BA-to-HS wage differential was essentially unchanged across that same set of cohorts. In addition, relative employment rates also change little across cohorts, suggesting that the invariance in relative wages was not simply due to institutional rigidities that forced the reaction to the supply increase to emerge in terms of quantities.

We argue that the absence of any obvious reduction in the wage differential in response to the large shift in relative supply is puzzling when we benchmark against US wage and employment patterns. Based on US estimates of skill biased demand shifts

²⁴These results are robust to excluding London and to how we weight the data (including not weighting it at all).

Table 4: Regressions of employee influence index

	(1)	(2)	(3)	(4)	(5)	(6)
Current % of BAs	0.573***	0.535**	0.600**	0.642***	1.306**	1.205***
	[0.123]	[0.272]	[0.238]	[0.134]	[0.585]	[0.251]
wave04	0.186***	0.188***	0.185***	0.177***	0.153***	0.152***
	[0.0237]	[0.0261]	[0.0255]	[0.0287]	[0.0353]	[0.0310]
wave11	0.293***	0.298***	0.291***	0.261***	0.185**	0.196***
	[0.0277]	[0.0401]	[0.0335]	[0.0322]	[0.0727]	[0.0415]
% of BAs in 1995-6		0.0626				
		[0.406]				
Current % of HS?			0.0395			
			[0.292]			
Constant	0.398***	0.396***	0.369*	1.073**	0.648***	0.654***
	[0.0241]	[0.0292]	[0.215]	[0.434]	[0.224]	[0.185]
further controls*	no	no	no	yes	yes	yes
instruments	na	na	na	na	cohort structure	4 IVs
Observations	670	670	670	670	670	580
R-squared	0.295	0.295	0.295	0.388	0.365	0.381

Note: all regressions are at the TTWA level, weighted by employment in the area. *Further controls include the current proportions of workplaces in the area by industry, by bands of workplace size, and by bands of organization size.

Table 5: Distribution of the employee influence index

	Mean	Std dev.
1998	0.48	0.24
2004	0.69	0.24
2011	0.85	.26

Note: the level of observation here is travel to work areas. They are weighted by employment size.

under the framework in [Card and Lemieux \[2001\]](#), one would predict declines in the education wage differential of approximately 20 percent across the full set of cohorts. Instead, it declined by less than 5 percent. The decline that did occur for the UK happened for cohorts after the ones that experienced the largest increases in education. Thus, both the size and the timing of the decline in the education differential are puzzling when viewed through the lens of the US experience in the same period.

One possible explanation for this puzzle is that it did not happen at all - that the wage differential really did fall but that coincident changes in composition concealed that decline. We argue that this is not the case. The combination of a large relative supply shift and the lack of movement in the relative wage differential is observed even after taking account of composition changes in terms of gender, public versus private sector, immigrant status, and whether a person had an advanced degree above a BA. Of course, it is possible that the relevant composition shift was in terms of unobserved individual heterogeneity, i.e., that the average ability of university graduates declined as the proportion of people with a BA grew. This mechanism is plausible but it is important to note that to explain the puzzle, we would require that unobserved worker ability declined significantly more among high school educated workers than among those with a BA. Using a bounding exercise based on a standard version of a Roy model of education choice, we show that the wage and employment movements do not fit with such an outcome. Even in more general selection frameworks, to explain the puzzle based on unobserved composition shifts one would need the extent of relative ability shifts to vary across time and age groups in a way that exactly offsets a widely varying set of relative supply shifts. While we cannot prove this did not happen, we view it as very unlikely. Thus, we conclude that the puzzle is real: large relative supply shifts were really met with little coincident change in the education wage differential while later, after the main increase in educational attainment was over, there was a small decline in the education wage differential.

There do exist standard models in which the relative supply shifts essentially engineer exactly offsetting demand shifts that leave relative wages unchanged. The best known such model is a Heckscher-Ohlin trade model with equal numbers of factors and goods. In the context of that model, shifts in relative factor supplies imply shifts in industrial structure toward goods with production that is intensive in the increasingly abundant factor. In addition, relative factor prices do not change. However, that model cannot explain our data patterns because we show that the relative supply shifts happen within not between industries. That is, the UK did not respond to the large increase in the education level of its workers by shifting toward education intensive goods. Instead, the change in the relative employment of high versus low educated workers happened within all industries.

The latter result implies something akin to a shift in a General Purpose Technology. More explicitly, we argue for a model in which firms in any sector can choose to produce using a centralized or a decentralized organizational structure as discussed in papers such as [Caroli and Van Reenen \[2001\]](#) and [Bloom et al. \[2012\]](#). In the decentralized

structure, workers need to be able to take individual initiative and control their own work - characteristics that we view as fitting more with higher educated workers. The model has a similar construction to a trade model in that the economy responds to a shift in the relative supply of more educated workers by shifting toward greater use of the decentralized organizational structure. And, as in the trade model, there is no adjustment in terms of relative wages or wage levels. We show that the UK did experience organisational change, including using more managers in production, in this period. Moreover, we show that areas in the UK which had more substantial increases in education levels are also areas where workers report having more control over their own work - something we see as a marker of a decentralized workplace. Importantly, this pattern occurs within industries not because of shifts in the industrial structure and is robust across a range of specifications, including the use of cohort-based instruments.

While we do not claim that our empirical results for the organisational change explanation are definitive, we believe that they do provide a coherent explanation of the remarkable stability of the education wage differential from the early 1990s until the mid-2000s in the UK, that occurred despite unprecedented increases in the share of entry workers with degree level education over the same period. This points to the UK responding to the substantial increase in university education through an adjustment in the organizational structure of work. We caution that it is dangerous to extrapolate. The UK has already surpassed the US in the BA proportion for the cohorts born after 1975, and soon it will probably exceed the US in the BA proportion for the entire workforce. It is plausible that the organisational technology is fully utilised so that a further educationalexpansion, in the absence of the arrival of a new technology, would result in declines in the education wage differential. There is already some sign of this decline in the private sector. The wage differential, though, remains substantial.

References

- Daron Acemoglu. Why do new technologies complement skills? directed technical change and wage inequality. *Quarterly journal of economics*, pages 1055–1089, 1998.
- Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics*, 4:1043–1171, 2011.
- Dirk Antonczyk, Thomas DeLeire, and Bernd Fitzenberger. Polarization and rising wage inequality: Comparing the us and germany. Technical report, ZEW, 2010.
- David H. Autor and David Dorn. The growth of low-skill service jobs and the polarization of the us labor market. *The American Economic Review*, 103(5):1553–1597, 2013.
- Paul Beaudry and David A Green. Wages and employment in the united states and germany: What explains the differences? *The American Economic Review*, 93(3): 573–602, 2003.

- Paul Beaudry and David A Green. Changes in us wages, 1976-2000: Ongoing skill bias or major technological change? *Journal of Labor Economics*, 23(3):609–648, 2005.
- Paul Beaudry, Mark Doms, and Ethan Lewis. Endogenous skill bias in technology adoption: City-level evidence from the it revolution. Technical report, National Bureau of Economic Research, 2006.
- Paul Beaudry, Mark Doms, and Ethan Lewis. Should the personal computer be considered a technological revolution? evidence from us metropolitan areas. *Journal of Political Economy*, 118(5):988–1036, 2010.
- Paul Beaudry, David A Green, and Benjamin M Sand. The great reversal in the demand for skill and cognitive tasks. Technical report, National Bureau of Economic Research, 2013.
- Nicholas Bloom, Raffaella Sadun, and John Van Reenan. Americans do it better: Us multinationals and the productivity miracle. *American Economic Review*, 102(1):167–201, 2012.
- Richard Blundell, Amanda Gosling, Hidehiko Ichimura, and Costas Meghir. Changes in the distribution of male and female wages accounting for employment composition using bounds. *Econometrica*, 75(2):323–363, 2007.
- Timothy F Bresnahan, Erik Brynjolfsson, and Lorin M Hitt. Information technology, workplace organization and the demand for skilled labor: Firm-level evidence. Technical report, National Bureau of Economic Research, 1999.
- David Card and John E DiNardo. Skill biased technological change and rising wage inequality: some problems and puzzles. Technical report, National Bureau of Economic Research, 2002.
- David Card and Thomas Lemieux. 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, 2001.
- Pedro Carneiro, Kai Liu, and Kjell G Salvanes. How does labor demand react to changes in the supply of skills? evidence from schooling reforms. 2014.
- Eve Caroli and John Van Reenen. Skill-biased organizational change? evidence from a panel of british and french establishments. *Quarterly journal of economics*, pages 1449–1492, 2001.
- Vincent Carpentier. Historical statistics on the funding and development of the uk university system, 1920-2002. 2004.

- Vincent Carpentier. Funding in higher education and economic growth in france and the united kingdom, 1921-2003. *Higher Education Management and Policy*, 18(3):1–26, 2006.
- Uniko Chen. The impact of the expansion in post-secondary education on the university wage premium in taiwan. Technical report, University of Michigan, 2013.
- Arnaud Chevalier, Colm Harmon, Ian Walker, and Yu Zhu. Does education raise productivity, or just reflect it?*. *The Economic Journal*, 114(499):F499–F517, 2004.
- Matias Cortes. Where have the middle-wage workers gone? a study of polarization using panel data. *Journal of Labor Economics*, 34:63–105, 2016.
- Paul J Devereux and Wen Fan. Earnings returns to the british education expansion. *Economics of Education Review*, 30(6):1153–1166, 2011.
- Christian Dustmann and Albrecht Glitz. How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics*, 33(3 Part 1):711–750, 2015.
- Christian Dustmann, Tommaso Frattini, and Ian P Preston. The effect of immigration along the distribution of wages. *The Review of Economic Studies*, 80(1):145–173, 2013.
- Nicole Fortin. Higher-educaton policies and the college wage premium: cross-state evidence from the 1990s. *American Economic Review*, 96(4):959–987, 2006.
- Peter Gottschalk, David A. Green, and Benjamin M. Sand. Taking selection to task: Bounds on occupational task prices for the us: 1984-2013. Technical report, University of British Columbia, 2014.
- Francis Green and Yu Zhu. Overqualification, job dissatisfaction, and increasing dispersion in the returns to graduate education. *Oxford Economic Papers*, 62(4):740–763, 2010.
- Lawrence F Katz and Kevin M Murphy. Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.
- David S Lee. Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies*, 76(3):1071–1102, 2009.
- Thomas Lemieux. Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *The American Economic Review*, pages 461–498, 2006.
- Ethan Lewis. Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics*, 126(2):1029–1069, 2011.
- Joanne Lindley and Stephen Machin. The postgraduate premium: Revisiting trends in social mobility and educational inequalities in britain and america. 2013.

- Stephen Machin. The changing nature of labour demand in the new economy and skill-biased technology change. *Oxford Bulletin of Economics and Statistics*, 63(s 1):753–776, 2001.
- Stephen Machin and Sandra McNally. Tertiary education systems and labour markets. *Tertiary Review*, pages 1–57, 2007.
- Stephen Machin and John Van Reenen. Technology and changes in skill structure: evidence from seven oecd countries. *Quarterly journal of economics*, pages 1215–1244, 1998.
- Stephen Machin and Anna Vignoles. *Education Policy in the UK. CEE DP 57*. ERIC, 2006.
- Charles F Manski. The selection problem. In *Advances in Econometrics, Sixth World Congress*, volume 1, pages 143–70, 1994.
- Steven McIntosh. Further analysis of the returns to academic and vocational qualifications*. *Oxford bulletin of economics and statistics*, 68(2):225–251, 2006.
- Robin Naylor, Jeremy Smith, and Shqiponja Telhaj. Graduate returns, degree class premia and higher education expansion in the uk. *Oxford Economic Papers*, page gpv070, 2015.
- Nigel C OLeary and Peter J Sloane. The return to a university education in great britain. *National Institute Economic Review*, 193(1):75–89, 2005.
- Ian Walker and Yu Zhu. The college wage premium, overeducation, and the expansion of higher education in the uk. *IZA Discussion Paper*, No. 1627, 2005.
- Ian Walker and Yu Zhu. The college wage premium and the expansion of higher education in the uk. *The Scandinavian Journal of Economics*, 110(4):695–709, 2008.

5 Appendix (for online publication)

5.1 Data

Most of the analysis in this paper is based on the demographic, education, employment, wage, and occupation variables in the UK Labour Force Survey (LFS). The LFS is a representative quarterly survey of approximately 100,000 adults that is the basis for UK labour force statistics. It is similar in nature to the US Current Population Survey (CPS) which we use as a comparison. We make use of UK LFS data running from the first quarter of 1993 to the last quarter of 2014.

Consistent definitions of education levels over time are obviously important in our investigations. The LFS asks respondents their highest level of educational qualification, with the potential categories changing over time. We take advantage of detail in the potential responses to construct six more aggregate categories that are consistent over time. For our main discussion, we then further aggregate those categories into three broader groups: a university degree level or above; secondary or some tertiary education below a university degree level; and below secondary qualifications. We draw the bottom line of secondary education as Grade C in the General Certificate of Secondary Education (GCSE), which are exams that students take at age 16, their final year of compulsory education. We consider this to be equivalent to High School graduation (HS) in the US because a substantial proportion of people have just GCSEs and the proportion of people strictly below the threshold in the UK is close to the proportion of HS drop-outs in the US.²⁵ We restrict our samples to people between ages 16 to 59 because the education qualification question was not asked of women over age 60 before 2007 unless they were working at the time of the survey. We carry out much of our investigation in terms of cohorts defined by the calendar year of birth.

Wages are surveyed in the first and fifth quarters an individual is in the survey. We use the hourly wage derived from the weekly wage in the main job and actual weekly hours. We recode hourly wages above £200 as missing. Our sample contains 30,000-60,000 wage observations per year. As we are interested in the real cost of labour to firms, we deflate wages by the GDP deflator²⁶. We are worried that student wages may include distortions related, for example, to co-op programmes and so drop all individuals who are part-time or full-time students in the survey week.

For comparative purposes, we look at the U.S. CPS, a large representative sample that is used in generating labour force statistics. We again use individuals aged 16 to 59 who are not full or part-time students in the survey week. The data is from the Outgoing Rotation Group samples of the CPS. Following Lemieux [2006], we do not use observations with allocated wages when calculating wage statistics. Wages and employment

²⁵For example, 10.6% of 25-34 year olds in the US are HS drop-outs in 2012. Coincidentally, the proportion of this age group in the UK who do not have qualifications equivalent to or higher than GCSE grade C is also 10.6% ; and 19.8% have qualifications equivalent to GCSE grade C and no higher qualifications.

²⁶Source: OECD

status refer to the week prior to the survey week, and we only use wage and occupation data on individuals who are currently employed in the reference week. We aggregate the U.S. workers into three education groups: high school drop-outs; high school graduates (which includes workers with some or completed post-secondary education below a Bachelor's degree); and university degree holders (Bachelors and higher).

5.2 The Card and Lemieux exercise

Following CL, we divide all workers into high school or university equivalents and count the number of hours worked by each.²⁷ The estimation has two stages. At the first stage, we regress the log wage ratio of university to high school workers within each age category on age dummies, year dummies, and $\log(C_{jt}/H_{jt})$, where C_{jt} is the total hours for university educated workers in age group, j , in year t and H_{jt} is the hours for high school educated workers. This provides us with an estimate of σ_A , the substitution elasticity between different age groups of the same education level. Using this, we can construct estimates of the total hours of university and high school labour, C_t and H_t , that take substitution across age groups into account. At the 2nd stage, the log ratio of the wages of the two different education groups is regressed on age dummies, year (linearly), $\log(C_t/H_t)$ and $(\log(C_{jt}/H_{jt}) - \log(C_t/H_t))$ to get an estimate of the elasticity of substitution between the education groups, σ_E , and a second estimate of σ_A . The coefficient on the linear year variable is interpreted as an estimate of the growth rate of skill biased demand shifts.

Assuming a linear time trend in demand shifts, we run the estimation for the US and UK separately and for various subsamples defined by gender and wage measure (weekly versus hourly). Table 6 contains the estimates of the elasticities, σ_A and σ_E . For the US, the estimated substitution elasticity between different age groups, σ_A , is near the top of the range of estimates reported by Card and Lemieux [2001]. But the estimate of σ_E , the substitution elasticity between the HS and BA labour inputs varies wildly by subsample and wage measure. For the UK, the estimates of σ_A have the wrong sign and the σ_E estimates are again highly variable. The model's failure to fit even the US data for our period (1993-2012) is not totally surprising. The CL model is essentially a richer version of Katz and Murphy [1992], and Beaudry and Green [2005] show that the Katz-Murphy model fits the US data (in the sense of having elasticity estimates of the right sign) less and less well the more recent the data that is used in the estimation. In addition, the US and UK estimates are so different that one can soundly reject the hypothesis that the data in the two countries are being generated from a common model of this form.

²⁷See Card and Lemieux [2001] for details. Essentially, workers who are exactly high school graduates contribute their reported number of hours to the high school category count and workers whose highest degree is a BA contribute their reported hours to the university category. Hours for high-school drop-outs are counted as a fraction of a high school graduate hour, with the fraction being determined by their relative wages. The hours for workers with a graduate degree are counted as more than one hour of a BA worker, with the relative wages again determining the ratio. The hours of workers with post-secondary degrees less than a BA are divided between the two categories according to relative wages. The wages used are for workers with exactly a high school graduation as their highest level of education and for workers with exactly a BA.

Table 6: Estimated labour substitution elasticities following Card and Lemieux [2001]

country	gender	wage measure	$\hat{\sigma}_A$	$\hat{\sigma}_E$	$\widehat{1/\sigma_E}$	s.e. of $\widehat{1/\sigma_E}$
assuming θ_t is linear in t						
US	male	weekly	9.29	-3.02	0.33	0.13
US	male	hourly	8.02	-5.22	0.19	0.12
US	both	weekly	12.91	-265.39	0.00	0.13
US	both	hourly	9.31	8.48	-0.12	0.12
US	female	weekly	9.43	7.70	-0.13	0.11
US	female	hourly	9.11	3.46	-0.29	0.12
UK	both	weekly	-46.05	11.51	-0.09	0.20
UK	both	hourly	-197.90	8.81	-0.11	0.18
UK	male	weekly	-30.46	5.50	-0.18	0.15
UK	male	hourly	-42.52	6.79	-0.15	0.14
UK	female	weekly	-19.88	10.51	-0.10	0.19
UK	female	hourly	-106.63	38.70	-0.03	0.19

Note: exclude workers working fewer than 30 hours a week. Age range 20-59. Regression is not weighted. Following C&L, the estimation has two stages. At the first stage, we regress the log wage ratio on age dummies, year dummies, and $\log(C_{jt}/H_{jt})$ to get an estimate of σ_A , the substitution elasticity between different age groups of the same education level. We back out the parameters α_j, β_j and construct C_t, H_t . At the 2nd stage, the log wage ratio is regressed on age dummies, year (linearly), $\log(C_t/H_t)$ and $(\log(C_{jt}/H_{jt}) - \log(C_t/H_t))$ to get an estimate of σ_E and a second estimate of σ_A . The reported $\hat{\sigma}_A$ is the 2nd-stage estimate.

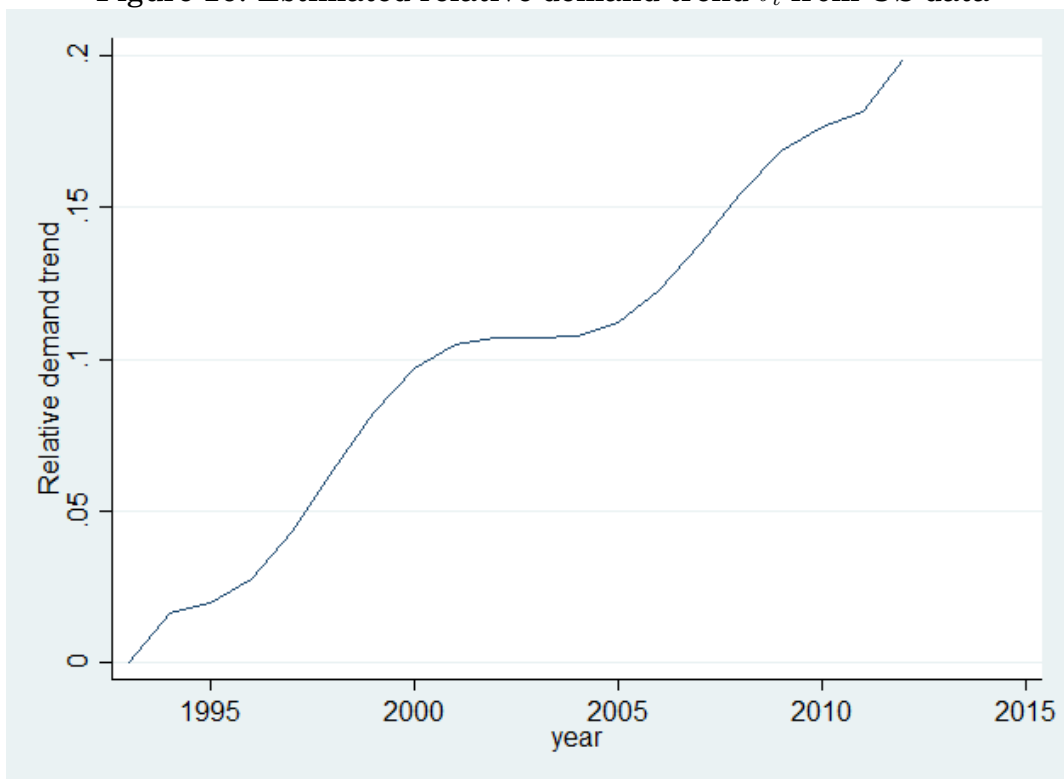
Next, we allow the demand shift to follow a 7th-order polynomial in year. The estimation on US data of hourly wage gives us $\hat{\sigma}^A = 9.5, \widehat{\text{sigma}}^E = 3.1$ and demand shifters that increase gradually over time as plotted in Figure 16. The same regression for the UK generates rather imprecise estimates as reported in Table 7. We proceed by using US estimates of parameters and demand trend and the UK labour supplies to generate predictions of graduate relative wages. Then, we extract cohort effects from these predictions and discuss them in the main body of the paper.

Table 7: Estimated labour substitution elasticities with non-linear demand trend

assuming θ_t is 7th order polynomial in t						
country	$\hat{\sigma}_A$	$\widehat{(-1/\sigma_A)}$	s.e. of $\widehat{(-1/\sigma_A)}$	$\hat{\sigma}_E$	$\widehat{(-1/\sigma_E)}$	s.e. of $\widehat{(-1/\sigma_E)}$
US	9.49	-0.11	0.02	3.18	-0.31	0.21
UK	-199.69	0.01	0.02	4.89	-0.20	0.35

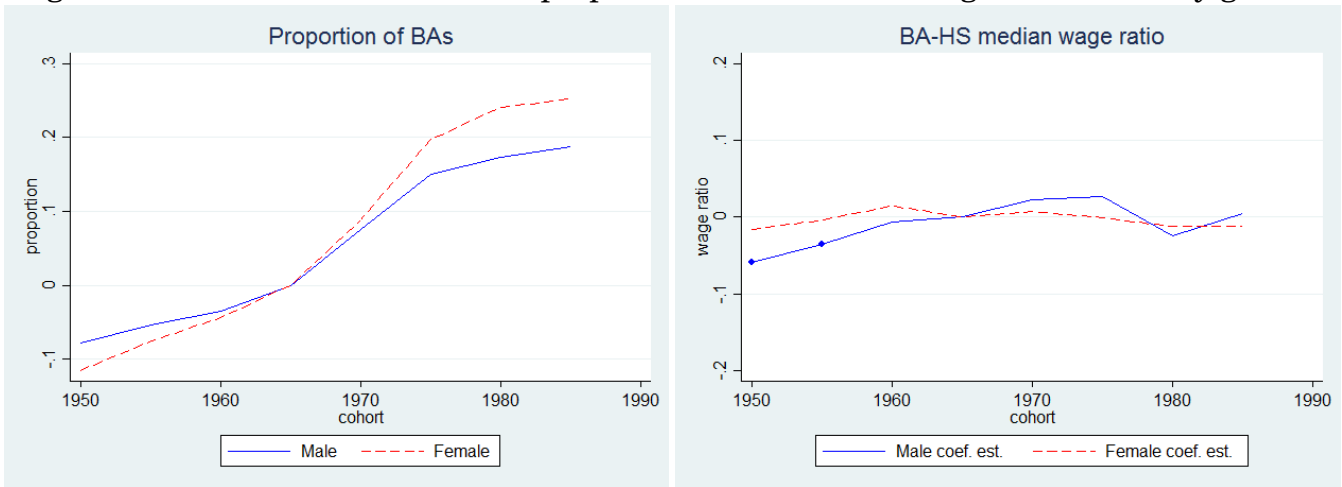
Note: the real wage measure is hourly wage. Both genders are included.

Figure 16: Estimated relative demand trend θ_t from US data



5.3 Compositional explanations

Figure 17: UK cohort effects on BA proportion and BA-to-HS wage differential by gender



Note: same specification as for figure 2.

Figure 18: Cohort effects for the BA-to-HS wage differential, UK private-sector only and till 2012



Table 8: Share of BAs among 16-59-year-old employees by sector

sector	shares		from 1994 to 2014	
	in 1994	in 2014	within effect	between effect
self employed	0.121	0.311	0.025	-0.000
private sector employees	0.091	0.284	0.121	-0.000
public sector employees	0.213	0.487	0.066	-0.001
Total	0.125	0.335	0.212	-0.001

from 1994 to 2009				
sector	share in 1994	share in 2009	within effect	between effect
self employed	0.121	0.254	0.017	0.000
private sector employees	0.091	0.219	0.079	0.000
public sector employees	0.213	0.400	0.048	0.002
Total	0.125	0.271	0.144	0.002

Note: the within-sector effect of a sector is defined as $(x_1 - x_0)(w_1 + w_0)/2$ and the between effect of a sector is defined as $(x_1 + x_0 - \bar{x}_1 - \bar{x}_0)(w_1 - w_0)/2$, where subscripts 0 and 1 denote the start year and the end year respectively, w denotes the sector's share of total employment, x denotes the share of BAs in the sector and \bar{x} denotes the aggregate share of BAs.

Because people (or, more properly, ability values) can be induced to switch into or out of higher education but not both at the same time, we can decompose the distribution function for BA wages in cohort $c+1$ into a component related to the distribution function for the “stayers” and a component for the “joiners”:

$$\Pr(\ln W_{uc+1} < w | \eta > A_{uhc+1}) = p_{uc+1} \Pr(\ln W_{uc+1} < w | \eta > A_{uhc}) + (1 - p_{uc+1}) \Pr(\ln W_{uc+1} < w | A_{uhc} \geq \eta > A_{uhc+1}), \forall w \quad (8)$$

where, p_{uc+1} is the proportion of the university educated in cohort $c+1$ who are stayers. Equation (8) holds for any wage level w , but we are interested in a particular level: the median wage in cohort $c+1$ for the university sector stayers, denoted as \tilde{w}_{uc+1} .

We can write \tilde{w}_{uc+1} as,

$$\tilde{w}_{uc+1} = \beta_{c+1u} + f_{c+1u}(age_{it+1}) + \lambda_u med(\eta_i + \epsilon_{ic+1t+1} | \eta_i > A_{uhc}) \quad (9)$$

Assuming stationarity of the η and ϵ distributions across cohorts, differencing this relative to the median conditional university wage in cohort c at the same age, age^* would yield,

$$\tilde{w}_{uc+1} - med(\ln W_{uct} | \eta_i > A_{uhc}) = \beta_{c+1u} + f_{c+1u}(age^*) - \beta_{cu} - f_{cu}(age^*) \quad (10)$$

That is, by comparing wage movements for people with the same set of η 's (the ones corresponding to choosing to get a university degree under either set of costs), we could obtain an estimate of the change in the actual wage profile across cohorts.

We cannot observe \tilde{w}_{uc+1} because we are comparing across cohorts and so cannot see who has ability levels that would result in their choosing the university degree in the different rationing situations. But we can obtain bounds for it. Returning to equation (8), we can obtain an estimate of p_{uc+1} based on changes in the size of the u group between cohort c and c+1 combined with the argument that people (or, rather, ability levels) either enter or leave the group but not both. We know that the second term on the right hand side of (8) ($\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | \eta > A_{uhc})$) equals 0.5 by the definition of \tilde{w}_{uc+1} , and the left hand side corresponds to a quantile of the conditional distribution of wages for the u group in the c+1 cohort, and so is calculable from the data. That only leaves the last term ($\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | A_{uhc} \geq \eta > A_{uhc+1})$) unknown and unknowable. However, since it is a probability, we can bound it on one side as $\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | A_{uhc} \geq \eta > A_{uhc+1}) = 1$, which corresponds to the marginal people who obtain a degree in cohort c+1 but would not have done so in cohort c having wages that place them below the median wage for the group who would get a degree in either cohort. Based on this, we can get an upper bound on \tilde{w}_{uc+1} by solving,

$$\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | \eta > A_{uhc+1}) = \frac{1}{2}p_{uc+1} + (1 - p_{uc+1}), \quad (11)$$

This is equivalent to trimming the bottom $(1 - p_{uc+1})$ proportion of observations from the c+1 university wage distribution and obtaining the median of the remaining sample.

Since the abilities of university “joiners” between cohort c and c+1 are assumed to be entirely below the abilities of the “stayers”, a joiner’s wage can be higher than a stayer’s only when the joiner has a particularly positive shock ϵ_{it} or the stayer has a particularly negative shock. As the idiosyncratic shock is assumed to be independent of ability, it follows that the joiners’ wage distribution is first order stochastically dominated by that of the stayers. Mathematically,

$$\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | A_{uhc} \geq \eta > A_{uhc+1}) \geq \Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | \eta > A_{uhc}) \quad (12)$$

Using the right side of this expression as the lower bound on $\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | A_{uhc} \geq \eta > A_{uhc+1})$ in (8) implies that the right hand side of (8) just equals 0.5. That is, the other bound is the c+1 median itself.

5.4 Model of endogenous technological change

As outlined in the paper, there are two production technologies: Centralized and Decentralized, and two tasks: Managerial and Labouring. Each task is produced by a combination of skilled and unskilled workers.

$$Y_{Ct} = M_{Ct}^\alpha L_{Ct}^{1-\alpha} \quad (13)$$

$$Y_{Dt} = M_{Dt}^\beta L_{Dt}^{1-\beta} \quad (14)$$

$$M_j = [aS_{Mj}^\sigma + (1-a)U_{Mj}^\sigma]^{1/\sigma}, \quad j = C, D \quad (15)$$

$$L_j = [bS_{Lj}^\rho + (1-b)U_{Lj}^\rho]^{1/\rho}, \quad j = C, D \quad (16)$$

Within skill groups, workers of different ages supply different quantities of effective labour and so there are only two effective wages: w^S for one unit of skilled labour and w^U for one unit of unskilled labour.

First, we will provide an alternative demonstration (in addition to the heuristic one given in footnote 19) that when both sectors $j \in \{C, D\}$ are in business, the equilibrium wage ratio w^S/w^U is determined by the parameters and is independent of the aggregate skill supply S/U .

For sector C, normalize the price of output to 1, then the FOCs are:

$$w^S = \alpha \left(\frac{M_C}{L_C}\right)^{\alpha-1} [aS_{MC}^\sigma + (1-a)U_{MC}^\sigma]^{\frac{1}{\sigma}-1} aS_{MC}^{\sigma-1} \quad (17)$$

$$= (1-\alpha) \left(\frac{M_C}{L_C}\right)^\alpha [bS_{LC}^\rho + (1-b)U_{LC}^\rho]^{\frac{1}{\rho}-1} bS_{LC}^{\rho-1} \quad (18)$$

$$w^U = \alpha \left(\frac{M_C}{L_C}\right)^{\alpha-1} [aS_{MC}^\sigma + (1-a)U_{MC}^\sigma]^{\frac{1}{\sigma}-1} (1-a)U_{MC}^{\sigma-1} \quad (19)$$

$$= (1-\alpha) \left(\frac{M_C}{L_C}\right)^\alpha [bS_{LC}^\rho + (1-b)U_{LC}^\rho]^{\frac{1}{\rho}-1} (1-b)U_{LC}^{\rho-1} \quad (20)$$

The ratios of the F.O.Cs \Rightarrow

$$\frac{w^S}{w^U} = \frac{a}{1-a} \left(\frac{S_{MC}}{U_{MC}}\right)^{\sigma-1} \quad (21)$$

$$= \frac{b}{1-b} \left(\frac{S_{LC}}{U_{LC}}\right)^{\rho-1} \quad (22)$$

Denote

$$x = \left(\frac{(1-a)w^S}{aw^U}\right)^{\frac{1}{\sigma-1}}, \quad (23)$$

$$y = \left(\frac{(1-b)w^S}{bw^U}\right)^{\frac{1}{\rho-1}} \quad (24)$$

We have $\frac{S_{MC}}{U_{MC}} = x$, $\frac{S_{LC}}{U_{LC}} = y$.

Substitute back to the first two FOCs,

$$w^S = \alpha \left(\frac{M_C}{L_C}\right)^{\alpha-1} [a + (1-a)x^{-\sigma}]^{\frac{1}{\sigma}-1} a \quad (25)$$

$$w^S = (1-\alpha) \left(\frac{M_C}{L_C}\right)^\alpha [b + (1-b)y^{-\rho}]^{\frac{1}{\rho}-1} b \quad (26)$$

We can cancel out M_C/L_C to get an expression for w^S in terms of $(\alpha, a, b, \sigma, \rho, \frac{w^S}{w^U})$:

$$\begin{aligned}
w^S &= \alpha^\alpha (1-\alpha)^{1-\alpha} [a + (1-a)x^{-\sigma}]^{\frac{(1-\sigma)\alpha}{\sigma}} a^\alpha [b + (1-b)y^{-\rho}]^{\frac{(1-\rho)(1-\alpha)}{\rho}} b^{1-\alpha} \\
&= \alpha^\alpha (1-\alpha)^{1-\alpha} [a + (1-a)x^{-\sigma}]^{\frac{(1-\sigma)\alpha}{\sigma}} (a^{\frac{\sigma}{1-\sigma}})^{\frac{(1-\sigma)\alpha}{\sigma}} [b + (1-b)y^{-\rho}]^{\frac{(1-\rho)(1-\alpha)}{\rho}} (b^{\frac{\rho}{1-\rho}})^{\frac{(1-\rho)(1-\alpha)}{\rho}} \\
&= \alpha^\alpha (1-\alpha)^{1-\alpha} [a^{\frac{1}{1-\sigma}} + (1-a)^{\frac{1}{1-\sigma}} (\frac{w^S}{w^U})^{\frac{\sigma}{1-\sigma}}]^{\frac{(1-\sigma)\alpha}{\sigma}} [b^{\frac{1}{1-\rho}} + (1-b)^{\frac{1}{1-\rho}} (\frac{w^S}{w^U})^{\frac{\rho}{1-\rho}}]^{\frac{(1-\rho)(1-\alpha)}{\rho}}
\end{aligned} \tag{27}$$

Given we had normalized the output price to 1, (27) is actually the ratio of skilled wage to the output price in the Centralized sector, written as a function of $(\alpha, a, b, \sigma, \rho, \frac{w^S}{w^U})$.

Similarly, if the Decentralized sector is in operation, we can write the ratio of skilled wage to the output price in the Decentralized sector as a function of $(\beta, a, b, \sigma, \rho, \frac{w^S}{w^U})$.

$$\frac{w^S}{P_D} = \beta^\beta (1-\beta)^{1-\beta} [a^{\frac{1}{1-\sigma}} + (1-a)^{\frac{1}{1-\sigma}} (\frac{w^S}{w^U})^{\frac{\sigma}{1-\sigma}}]^{\frac{(1-\sigma)\beta}{\sigma}} [b^{\frac{1}{1-\rho}} + (1-b)^{\frac{1}{1-\rho}} (\frac{w^S}{w^U})^{\frac{\rho}{1-\rho}}]^{\frac{(1-\rho)(1-\beta)}{\rho}} \tag{28}$$

Both technologies are in business when their output prices are the same, that is, $P_D = 1$. Then equating (27) and (28)

$$\alpha^\alpha (1-\alpha)^{1-\alpha} = \beta^\beta (1-\beta)^{1-\beta} [a^{\frac{1}{1-\sigma}} + (1-a)^{\frac{1}{1-\sigma}} (\frac{w^S}{w^U})^{\frac{\sigma}{1-\sigma}}]^{\frac{(1-\sigma)(\beta-\alpha)}{\sigma}} [b^{\frac{1}{1-\rho}} + (1-b)^{\frac{1}{1-\rho}} (\frac{w^S}{w^U})^{\frac{\rho}{1-\rho}}]^{\frac{(1-\rho)(\alpha-\beta)}{\rho}} \tag{29}$$

This implicitly defines the equilibrium skill wage ratio $\frac{w^S}{w^U}$ as a function of parameters $(\alpha, \beta, a, b, \sigma, \rho)$. Note that this equation does not contain the aggregate labour supplies S or U .

We can consider the right-hand-side of (27) and (28) as two curves of $\frac{w^S}{w^U}$. Suppose the parameters are such that the two curves cross just once and that after the intersection the curve represented by (28) lies above the other. This means when the wage ratio exceeds the equilibrium ratio, the Decentralized technology will yield a lower output price and the centralized technology will not be in use. The reverse is true when the wage ratio is below the equilibrium.

So far we have shown there is an equilibrium wage ratio $\frac{w^S}{w^U}$ under which both the Centralized and Decentralized technologies are in use. Now we work out the range of aggregate relative skill supply S/U that is compatible with that equilibrium wage ratio.

Divide (26) by (25) \Rightarrow

$$\alpha [a + (1-a)x^{-\sigma}]^{\frac{1}{\sigma}-1} a = (1-\alpha) \frac{M_C}{L_C} [b + (1-b)y^{-\rho}]^{\frac{1}{\rho}-1} b \tag{30}$$

By definition, we have

$$\begin{aligned}
M_C &= S_{MC} [a + (1-a)x^{-\sigma}]^{\frac{1}{\sigma}} \\
L_C &= S_{LC} [b + (1-b)y^{-\rho}]^{\frac{1}{\rho}}
\end{aligned}$$

Substitute into (30) \Rightarrow

$$\alpha [a + (1-a)x^{-\sigma}]^{\frac{1}{\sigma}-1} a = (1-\alpha) \frac{[a + (1-a)x^{-\sigma}]^{\frac{1}{\sigma}} S_{MC}}{[b + (1-b)y^{-\rho}]^{\frac{1}{\rho}} S_{LC}} [b + (1-b)y^{-\rho}]^{\frac{1}{\rho}-1} b \tag{31}$$

After some manipulation, we get

$$\frac{S_{MC}}{S_{LC}} = \frac{\alpha}{1-\alpha} \frac{a[b+(1-b)y^{-\rho}]}{b[a+(1-a)x^{-\sigma}]} \quad (32)$$

Thus, in the Centralized sector, the four quantities of labour demand ($S_{MC}, S_{LC}, U_{MC}, U_{LC}$) are proportional to $(\alpha a[b+(1-b)y^{-\rho}]xy, (1-\alpha)b[a+(1-a)x^{-\sigma}]xy, \alpha a[b+(1-b)y^{-\rho}]y, (1-\alpha)b[a+(1-a)x^{-\sigma}]x)$. The ratio of skilled to unskilled labour in sector C is

$$\frac{S_{MC} + S_{LC}}{U_{MC} + U_{LC}} = \frac{\alpha a[b+(1-b)y^{-\rho}]xy + (1-\alpha)b[a+(1-a)x^{-\sigma}]xy}{\alpha a[b+(1-b)y^{-\rho}]y + (1-\alpha)b[a+(1-a)x^{-\sigma}]x} \quad (33)$$

Symmetrically, the ratio of skilled to unskilled labour in the Decentralized sector is

$$\frac{S_{MD} + S_{LD}}{U_{MD} + U_{LD}} = \frac{\beta a[b+(1-b)y^{-\rho}]xy + (1-\beta)b[a+(1-a)x^{-\sigma}]xy}{\beta a[b+(1-b)y^{-\rho}]y + (1-\beta)b[a+(1-a)x^{-\sigma}]x} \quad (34)$$

Denote the right-hand-side of (33) and (34) as R_C, R_D . And denote $U_j = U_{Mj} + U_{Lj}, S_j = S_{Mj} + S_{Lj}, j \in \{C, D\}$.

As long as $R_C < \frac{S}{U} < R_D$, the relative size of the Centralized and Decentralized sectors will adjust to:

$$\begin{aligned} U_C &= \frac{R_D - \frac{S}{U}}{R_D - R_C} U, & S_C &= \frac{(R_D - \frac{S}{U})}{R_D - R_C} U R_C \\ U_D &= \frac{\frac{S}{U} - R_C}{R_D - R_C} U, & S_D &= \frac{(\frac{S}{U} - R_C)}{R_D - R_C} U R_D \end{aligned}$$

Thus, both technologies are in use when the aggregate skill ratio is within a range. Within this range, as the aggregate skill supply $\frac{S}{U}$ increases, the employment share of the decentralized sector will increase, and the skilled wage ratio $\frac{w^S}{w^U}$ stays constant as it is determined by the parameters alone.