The UK Education Expansion and Technological Change

Richard Blundell, David Green and Wenchao Jin*

January 23, 2018

Abstract

The proportion of UK people with university degrees by age 30 more than doubled between those born in 1965-69 and those born ten years later. However, the age profile of the graduate wage premium remained largely unchanged across cohorts. We show that these patterns cannot be explained by composition changes. Instead, we present a model in which firms choose between centralized and decentralized organizational forms and demonstrate that it can explain the main patterns. We also show the model has implications that differentiate it from standard exogenous technological change models and that UK data fit with those implications. We argue that our model of endogenous technological choice is plausible for the UK because of its status as a technological follower where workplaces only adopt frontier IT related technologies when educational levels in the labour force are high enough.

In the period extending from the early 1990s to the present, the UK economy experienced a dramatic 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%.

*Blundell: University College London and Institute for Fiscal Studies, Gower Street, London, WC1E 6BT, UK, r.blundell@ucl.ac.uk. Green: Vancouver School of Economics, UBC, and Institute for Fiscal Studies, 6000 Iona Dr, Vancouver, BC V6T 1L4, Canada, david.green@ubc.ca. Jin: University College London, Gower Street, London, WC1E 6BT, UK, w.jin.11@ucl.ac.uk. 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.

In this paper we examine the impact of that increase on the UK labour market, using our findings as a basis for contributing to the ongoing discussion about the interaction of educational attainment and technological change.

There is a strong consensus among economists that the Information Technology (IT) revolution has played a central role in determining wage and employment outcomes in many economies in the last four decades and that the effects of education should be viewed in conjunction with that revolution (see Acemoglu and Autor (2011) for a comprehensive review of the literature on these topics). Most famously, changes in the wage distribution have been described as a race between skill-biased demand shifts emanating from IT innovations and increases in skills generated by changes in education levels (Goldin and Katz, 2008). The core idea underlying this consensus is that the new technologies are complementary with skills. Intertwined with the broad literature on the effects of technology on the wage structure in general is a literature on skills, IT, and the organizational structure of the firm (e.g., Bresnahan et al. (1999); Caroli and Van Reenen (2001); Bloom et al. (2014)). This literature seeks to look inside the 'black box' of the firm to understand how skills complement IT, and its main message is that IT, by altering information flows and communications within firms, implies a shift in the optimal organization of the firm toward a form that is more decentralized and flexible. Decisions, information transfer, and co-ordination of tasks happen throughout the organization instead of through top-down direction as in the previous Taylorist form. The shift in organizational form is the channel through which more educated workers benefit from the broad technological change since human capital investment gives workers greater ability to deal with increased change and decision making and makes them relatively more productive in the new environment. On the other side of the ledger, a large literature on polarization argues that IT replaces routine tasks to the detriment of less educated workers (Acemoglu and Autor (2011)). Our approach, both theoretically and empirically, incorporates both the decentralization and routinization elements of the IT revolution and how they intersect with educational change.

Much of the empirical work and the clear majority of the theorizing on the interaction of education and technological change has been done on the US economy. However, there are good reasons to believe that the US is the technological leader in this period and, because of that, may exhibit special relationships between technological change and education that do not apply even to other advanced economies. Given this, we view the UK educational expansion as an opportunity to study the relationship between education and technological change in a technological follower, as we believe the UK has

been in terms of skill biased technologies and firm organizational forms. We will argue that taking this perspective has an impact on which model of technological change and education one adopts. In particular, we argue for a model in which firms choose among existing technologies rather than one in which new invention is the focus. Our claim is that these models provide a natural explanation for a remarkable fact for the UK: that its very substantial increase in education level was accompanied by no change in the university-high school wage differential. We present a model that captures this fact but also has further testable implications that we show are supported in the data.

Our key message is that technological change is not one size fits all. Many papers look for evidence of the importance of technological change in common movements in the relation between educational attainment and wage differentials across countries. The argument being that if new technologies are accessible in all developed countries then, conditional on mediation through relative skill supply shifts, it should act as a common force showing up in the same way in all developed countries. In contrast, differential movements in the combination of educational attainment and skill based wage differentials across countries is taken as evidence of the impact of other, non-technological factors (e.g., Caroli and Van Reenen (2001)'s examination of French and English data or Antoncyzk et al. (2010)'s 2010 assessment of education and wage movements in the Germany and the US). In contrast, we argue that the same changes in factor supplies interacting with the same technology can dictate quite different wage outcomes for two countries depending on whether they are leaders or followers in the adoption of that technology.

The paper proceeds in seven sections not including the introduction. In the first section, we provide a brief overview of different models of technological change and educational attainment as well as setting out our argument that the US has been a technological leader and the UK a follower in recent decades. In the second section, we establish the core patterns for the UK, relying largely on Labour Force Survey (LFS) data between 1992 and 2014. 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 the period when the BA-HS wage differential in the UK increased substantially (Machin and McNally, 2007). Our goal in the third section is to demonstrate that both the size of the educational change and the lack of movement in the wage differential are real: they cannot be explained as, for example, declines in the actual wage differential

that are masked by changes in composition. We consider compositional changes related to increases in the female participation rate, the shift toward more advanced university degrees over time, and the substantial increase in immigration. We also consider changes in unobserved abilities within educational categories that one might expect to arise as part of the substantial increase in the proportion of a given birth cohort who have a university degree. 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.

In the fourth section, we set out a model of technological change which focuses on the role of decentralization of decisions and information. The model is a variant of models in Rosen (1978) and Borghans and ter Weel (2005) and, as in the latter paper, incorporates technological change that induces polarization. Interestingly, in addition to the destruction of middle paying occupations, which is the commonly discussed form of polarization, the model implies that an additional type of polarization can arise with an increase in the relative supply of skilled workers: one in which more and more of the top end management tasks are done by the more educated while the less educated are relegated to the bottom end tasks (Acemoglu and Autor (2011)). The model is one of technological choice in which firms use skilled and unskilled labour and choose between an older, centralized mode of operation and a newer, decentralized mode. The model endogenously generates an unchanging skilled-unskilled wage ratio. This was the point of using this type of model and so that outcome provides no proof of the model's relevance. However, the model also generates testable added implications about the pattern of employment in manager positions for skilled and unskilled workers as the relative supply of skilled workers increases. 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 educational attainment.

We examine these empirical implications of our model in section five. In that section, we also investigate further implications by examining the relationship between the educational 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 establish that this is a causal relationship using an IV strategy using a combi-

nation of parental education and the proportion of the population made up of people born in the cohorts most affected by the educational increase in 1995 (i.e., before the entry of the most affected cohorts into the labour force). We view this as a credible strategy since the validity of the instrument just requires that differences in fertility rates across regions were not driven by changes in firm organizational forms twenty years later. Thus, the data fits with a model in which increased educational attainment induces more and more firms to choose a decentralized organizational form. One interesting implication of the model that is confirmed in this data is that increases in education levels in a region actually induce larger increases in individual decision making among less educated than more educated workers. This arises because under the old, centralized technology, more educated workers were disproportionately managers and were already making their own decisions. It is for the less educated that decentralization is a particularly big revolution.

In section six, we briefly consider competing explanations for the data patterns but argue that they are either rejected by our empirical results or require knife edge assumptions (such as that one requires exogenous skill biased demand shifts that just happen to vary to exactly offset different relative supply shifts in each cohort) which we find hard to justify. The seventh section of the paper contains conclusions.

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, showing that it increased sharply around the early 90s. He 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 OÆLeary and Sloane (2005); Walker and Zhu (2008); Green and Zhu (2010); 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, in part, by providing an 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. Instead, we present a model in which this pattern arises endogenously, which has ramifications for how we think about the interactions of technological change, factor supplies, and factor demand. We also differ from earlier studies in our explicit emphasis on the firm

organization part of the process - that is where our empirical work focuses. Combined, these give us new insights into how technological change affects economies. Overall, we view studying the UK as an opportunity to examine the impact of education policy on technological adoption and, through it, on wages in the situation that is likely relevant for most countries - being a technological follower.

1 Technological Leadership and Models of Technological Change

One can think of the interaction of increased human capital attainment with technological change (broadly defined so as to include changes in organizational form) in terms of three main models. The first is one in which the technological change is exogenous: a new technology is introduced for an unspecified reason and is so dominant in terms of cost savings over existing technologies that it is adopted on a wide scale. Wage differentials are then determined by the interaction of relative demand shifts arising from this technological change (hinging on the skill bias of the technological change) and shifts in supply. Early versions of this model that focused directly on the university/high school wage differential have generally been shown not to fit the data well (Beaudry and Green (2005); Card and DiNardo (2002); Acemoglu and Autor (2011)) but the more recent literature on polarizing changes in technology also have this broad form (e.g., Autor and Dorn (2013)). In all of these models, wage differentials reflect the classic race between technological change and education, with wages in higher skilled groups (defined by education or occupation) rising more if increases in the supplied labour in that group are smaller ((Goldin and Katz, 2008)).

The second model type is one in which the invention of new technologies is a function of movements in the relative factor endowments in an economy. Thus, an increase in the education level in an economy provides an incentive for inventors to create new technologies that are relatively intensive in the use of higher educated labour (Acemoglu (1998)). In this case, the relative increase in demand for skills is actually induced by the increase in their supply. Acemoglu (2007) shows that in cases where innovation is created by government funded research or by monopolistic or oligopolistic firms, if the elasticity of substitution between skilled and unskilled labour is high enough then the increase in the relative supply of skilled workers can actually induce an increase in the relative wage of the skilled workers. That is, if this is the relevant model for thinking

about the interaction of skill supply and innovation then attempts to combat inequality by increasing educational attainment could enhance rather than reduce wage inequality.

The third type of model is one in which the technologies already exist and firms choose among them. These endogenous choice models have the structure of a 2 sector by n factor trade model with the same implications. In particular, if n>2 and all factors are inelastically supplied then these models can yield the same implications as the induced invention models, i.e., that increases in the relative supply of skill can generate increases in the skilled wage differential (Beaudry and Green (2003); Beaudry et al. (2010)). On the other hand if all but two of the factors are perfectly elastically supplied (as one might expect if new organizational capital, for example, requires a one time investment but widely accessible information thereafter) then even large increases in the relative supply of educated labour will leave skill group wages unchanged if the economy remains within a region in which both the new and old technologies are in use (the cone of diversification) (Beaudry and Green (2003); Beaudry et al. (2010)).[1] We believe this class of models fits with the spirit of the literature on decentralization and organizational form which, starting with Milgrom and Roberts (1990)'s seminal contribution, often approaches organizational form as something firms optimally choose given existing options (e.g., Bresnahan et al. (1999); Caroli and Van Reenen (2001)).

Which of these models is relevant is potentially context contingent. There may be technologies that are so superior that the exogenous technical change model is clearly relevant (though we suspect those situations are extremely rare). Further, in economies that are technological leaders in time periods when new technological possibilities are opening up, the induced invention model may be more appropriate. However, for other, following economies (and even in the technological leaders in periods after the initial invention is complete) it seems to us more useful to think in terms of the endogenous technological choice models, with firms choosing from an already invented set of options.

Much of the theorizing about these different models has been done with the US economy and US stylized facts in mind. But the US context may be quite unique. In particular, we believe that there are good reasons to believe that the US has been a technological leader in the development of skill biased technologies and their associated organizational forms in recent decades. In terms of the induced investment model, the US has a work force with relatively high levels of advanced education. Indeed, in 1980,

[1]Note that this is different from the first, competitive model in Acemoglu (2007). In that model, firms choose among technologies but those technologies are included as another input in a standard, unitary production function. The technological choice models we are referring to involve choosing among completely different technologies with different substitution elasticities.

on the cusp of the computer revolution, 22% of the US populaton aged 25 to 64 had a tertiary education, which was by far the highest in the OECD (Barro-Lee(2017)).[2] Thus, incentives for innovators to generate human capital intensive technologies would have been highest in the US. Moreover, the US has had the highest ratio of investment in ICT (Information, Computers, and Technology) capital to total non-residential gross fixed capital throughout the 1985 to 2010 period (OECD(2017)). The idea that the US is the innovation leader in this area is also supported by evidence in Bloom et al. (2012) showing that US multinationals use a more decentralized structure relative to both domestic firms and multinationals from other countries even when all are observed operating in the same economy (the UK).

On the other side, there is also good reason to believe that the UK is a follower in the area of skill biased technologies and their associated organizational forms. Certainly, the UK was well behind the US in educational attainment at the beginning of the computer revolution. This can perhaps be most clearly seen in data organized by birth cohort. For the cohort born between 1955 and 1959 in the UK (and who would have turned 25 in the early 1980s, at the outset of the computer revolution), 10% held a university degree by age 30 compared to 22% for the same cohort in the US.[3] For the cohort born a decade later, the numbers were 16% for the UK and 27% for the US - the UK was still a laggard. Thus, viewed through the lens of the theory of induced invention, we would not expect the UK to have been a leader in skill-biased innovation. However, beginning in 1988, the UK undertook a series of reforms that greatly expanded access to a university education and by the cohort born between 1975 and 1979 (who turned 25 in the early 2000s), the UK had surpassed the US with 33% attaining a university degree in the UK compared to 31% in the US. That increase in the educational attainment of new labour market entrants in the UK could have provided the conditions for firms to move toward the new technologies developed in the US. Interestingly, the proportion of investment that was in ICT capital shot up in this decade in the UK, approximately doubling at the same time the proportion of new labour market entrants with a university education also doubled (OECD(2017)).[4] Further, the evidence in Bloom et al. (2012) about use of decentralized organizational forms also suggests that UK firms were following rather than leading the

[2]The next highest were Australia and New Zealand at about 15%, with the remainder of the OECD decidedly lower.

[3]These figures are computed from the UK LFS for the years 1992 to 2015 and the Outgoing Rotation Group sample from the US Current Population Survey for the same years.

[4]The proportion of total non-residential fixed capital investment in ICT increased by 88% in the UK between 1990 and 2000. Only Finland and South Korea had faster growth in this proportion in this decade. In comparison, the proportion grew by 37% in the US.

charge. They argue that UK firms were 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 lower 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. We think that these patterns fit most naturally with models of technological choice.

There is another reason to think of the UK experience in terms of the technological choice model - indeed, it is the reason we considered such a model in the first place - and that is the behaviour of UK wages during the period of educational expansion. We turn to investigating that behaviour and the movements of educational attainment in more detail in the next section.

2 Data and Core Patterns

2.1 Data

Our main empirical work 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.[5] 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 deflator[6]. 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 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).

2.2 UK Wage and Educational Attainment Movements

2.2.1 Changes to educational attainment

We begin with a figure showing the level of university attainment over time for the UK, with the US as a benchmark. 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 the BA proportion on a fifth order polynomial in age and a complete set of cohort

[5]For 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.

[6]Source: OECD

dummies.[7] 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.[8] Constructing cohort effects in this way controls for differences in the age ranges over which the different cohorts are present in our data.

Figure 1 contains plots of the cohort effects for the BA proportion for both the UK and the US. The UK had strikingly lower proportions of people with a university education in the cohorts born before 1965; proportions that were less than half those in the corresponding US cohorts. But between the 1965-69 and the 1975-79 cohorts, the UK completely closed 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 grew 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.

[7] For 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.

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

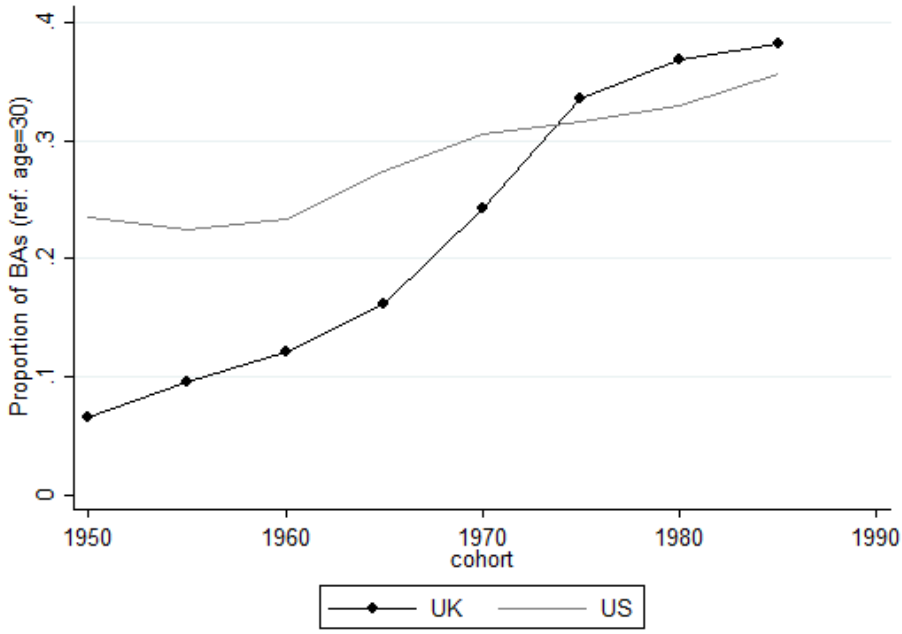

Figure 1: Proportion of People with a BA or Higher Education by Cohort, UK and US
The sample is restricted to age 22-59 and excludes full-time students. Each cohort-age-education cell has at least 100 observations. The cohort effects are estimated according to the description in text. Authors' calculation of the Labour Force Survey and Current Population Survey.

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 (OÆLeary 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 educational 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.[9] 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 1 it by no means stopped.

2.2.2 Changes in relative wages

The second main pattern relates to wages. In Figure 2 we plot the ratio of BA to high school median hourly wages at each age for the set of 5-year birth cohorts that are present in our UK LFS data. Perhaps the most striking feature of this figure is the similarity of age profiles for wage differentials across cohorts. With only one exception, every pair of cohorts with profiles in age ranges that overlap considerably cross more than once, and in most cases, multiple times.[10] For the 1965-69 birth cohort, at age 30, the ratio of the BA-to-HS wage ratio was 1.45. 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.

[9] This has been clearly shown in Figure 2 in Carpentier (2006)

[10] The only exception is the 1950-54 cohort and the 1965-69 cohort. The profiles for these cohorts only cross once, with the profile for the later cohort being otherwise always above that for the earlier one. This potentially fits with findings in the literature of an increase in the returns to education in the UK during the 1980s (Machin (2001)).

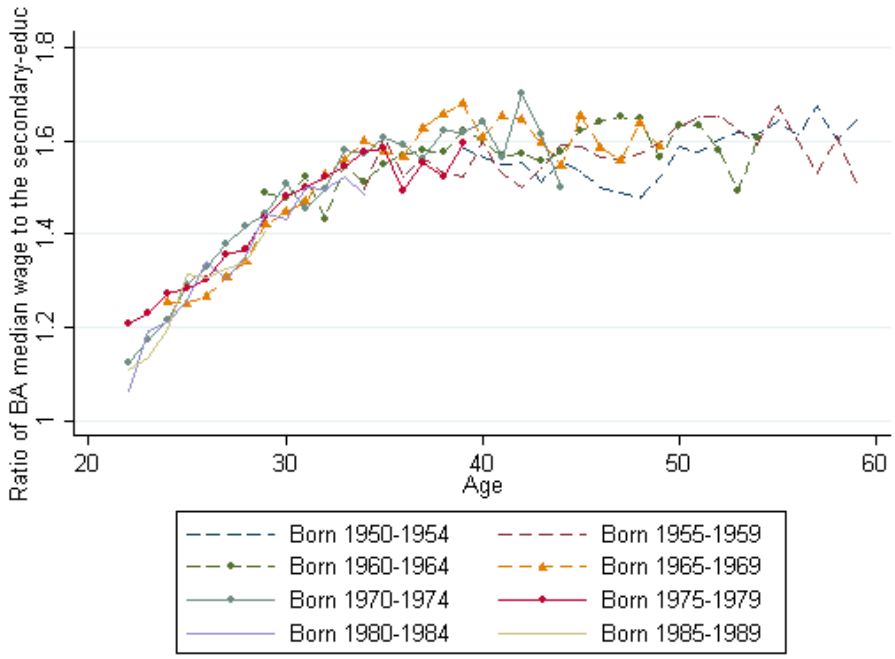

Figure 2: Ratio of BA median wage to that of high-school graduates by 5-year birth cohorts, UK

Wage is hourly. BA refers to individuals who have a first degree or a higher level qualification. High-school graduates refer to those who have at least GCSEs and do not have a first degree. Authors' calculation of the Labour Force Survey. All figures and tables in this paper are based on the UK Labour Force Survey unless otherwise noted.

In Figure 3, we put the wage and BA share movements together. Similar to what was done with BA proportions in Figure 1, we regress the median wage ratios on cohort dummies and an age polynomial. In Figure 3, we plot both the resulting wage cohort effects and the proportion BA cohort effects, normalizing both to zero for the 1965-69 cohort. For the pre-1965-69 cohorts, there is some increase in the wage ratio, fitting with results about the UK education differential before 1990 (Machin (2001)), combined with small movements in the proportion with a BA. 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 2 percent relative to the 1965-69 cohorts and is not statistically significant at the 5 percent level. Overall, the lack of movement in the wage differential is striking considering that the percentage with a Bachelors or higher degree at age 30 ranges from a low of 13% for

Figure 3: Cohort Effects for the BA-HS Wage Ratio and Proportion BA, UK
Wage is hourly. BA refers to individuals who have a first degree or a higher level qualification. High-school graduates refer to those who have at least GCSEs and do not have a first degree.

the 1960-64 birth cohort to a high of 37% for the 1980-85 cohort.[11] In the next section, we set out a model that accommodates this outcome as an endogenous response to the increase in educational attainment.

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 OÆLeary 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

[11] 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.

more recent data.

3 Accounting for Changes in Composition

One possible explanation for why such substantial increases in educational attainment were associated with little or no change in educational wage differentials 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 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.

3.1 Observable Characteristic Composition

Perhaps the most obvious compositional shift in terms of observable worker characteristics is related to the increase in female labour force participation. If the added female entrants with BA's are successively more able across cohorts (compared to the added females with a high school level education) then their entry could hide a decline in the education differential in prices per efficiency unit. However, even the most cursory glance at the data indicates that gender composition shifts are not a source of problems since the wage patterns are the same for males and females. In Figure 4 we plot the Proportion of BA's and the BA-HS median wage ratio for males and females separately (again, obtaining cohort effects from regressions including age polynomials). For both genders, we see the dramatic increase in BA's for the cohorts born between the mid-1960's and the mid-1970's, with the increase being larger for females. But, again, neither shows a change in the wage differential across cohorts and so a change in weighting between men and women would not alter the overall wage picture.

Immigration is another potential source of compositional change, with the proportion of UK workers born outside the UK doubling over the past two decades (Dustmann et al. (2013)). As immigrants are more likely to have university degrees, the large flows of

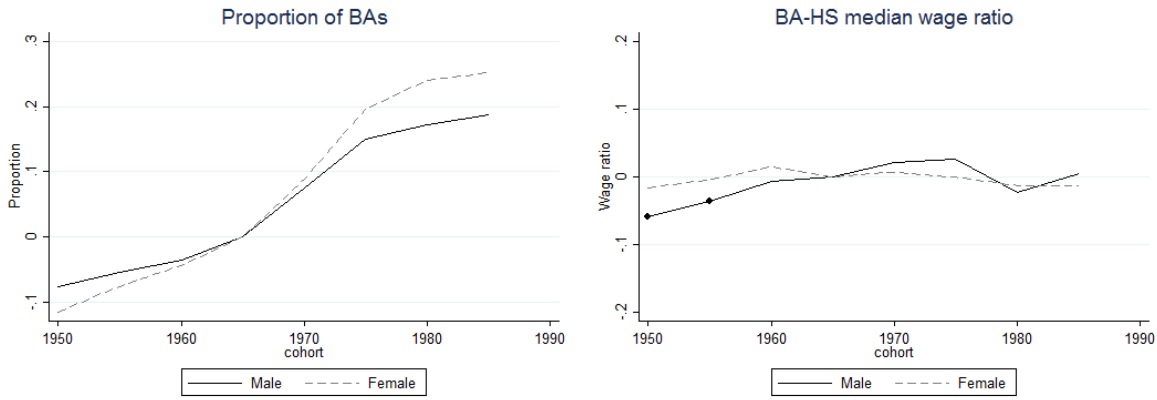

Figure 4: UK cohort effects on BA proportion and BA-to-HS wage differential by gender
The cohorts effects use the same specification as for Figure 3.

immigrants contribute directly to the aggregate increase in the share of BAs in the workforce and, again, differences in their ability levels could affect observed wage differentials. To check on whether this is an issue, in Figure 5 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, implying that composition changes related to immigration are not driving our main patterns.[12]

A third significant compositional shift among graduates is the increasing prevalence of postgraduate degrees. The dark, solid line in Figure 6 plots the proportion of people with a postgraduate degree conditional on having a university degree, normalized to zero for the 1965-69 cohort. The increase in that proportion across cohorts could generate an increase in the median wage for all university degree holders, which is the measure we have used so far. Importantly, though, the main increase in postgraduate degrees happens between the 1975-79 and 1980-84 cohorts - after the substantial increase in the proportion of people with any university degree - and thus is unlikely to impact the pattern of the lack of change in the education wage premium during the main expansion in overall university education. This, in fact, is what we find. To assess the impact of

[12] 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. In Appendix A.3, we present an exercise in the spirit of Dustmann et al. (2013) which shows that adjusting for immigrant effective education levels does not substantially alter our main pattern.

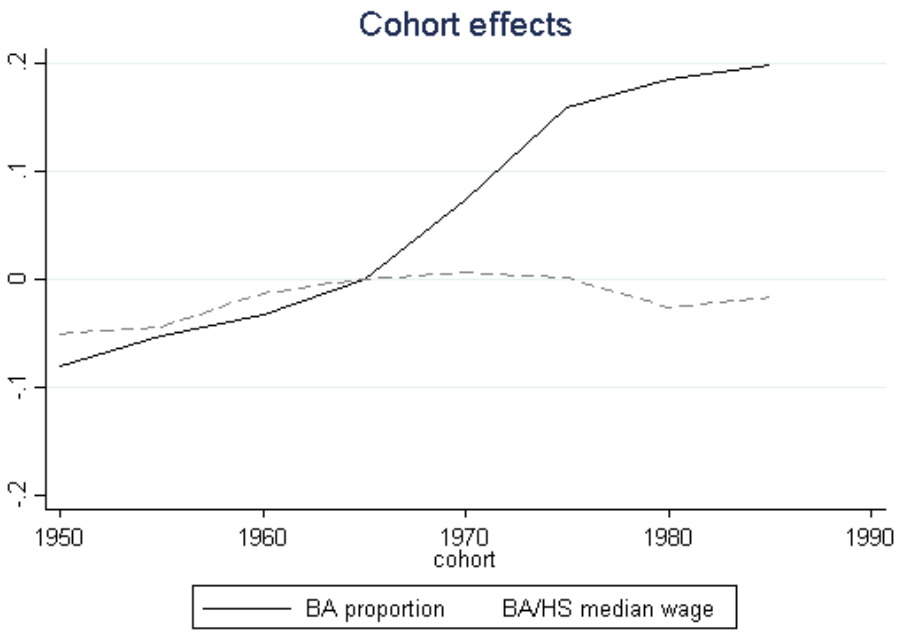

Figure 5: Cohort effects for the BA proportion and wage ratio, natives only
Cohort effects are normalized to zero at the 1965 cohort, as in Figure 3.

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. The wage ratio continues to show very little movement between the 1965-69 and 1975-79 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.

In Appendix A.3, we present a further exercise, breaking down the data into public versus private sector employment and wages. With no substantial change in the proportion of workers in the public sector and very similar movements of the BA proportion in the two sectors, public versus private sector composition also does not provide an explanation for our main patterns. However, as with focusing on the exact BA wages, working only with private sector wages implies a flat education differential for the biggest education expansion cohorts and a slight decline in the differential for the cohorts that followed. Overall, we conclude that shifts in composition with respect to observable worker characteristics cannot explain our main pattern of substantial education increases paired with an invariant education wage premium, but taking account of the shifts does reveal a decrease in the differential for later cohorts for which the education increases were

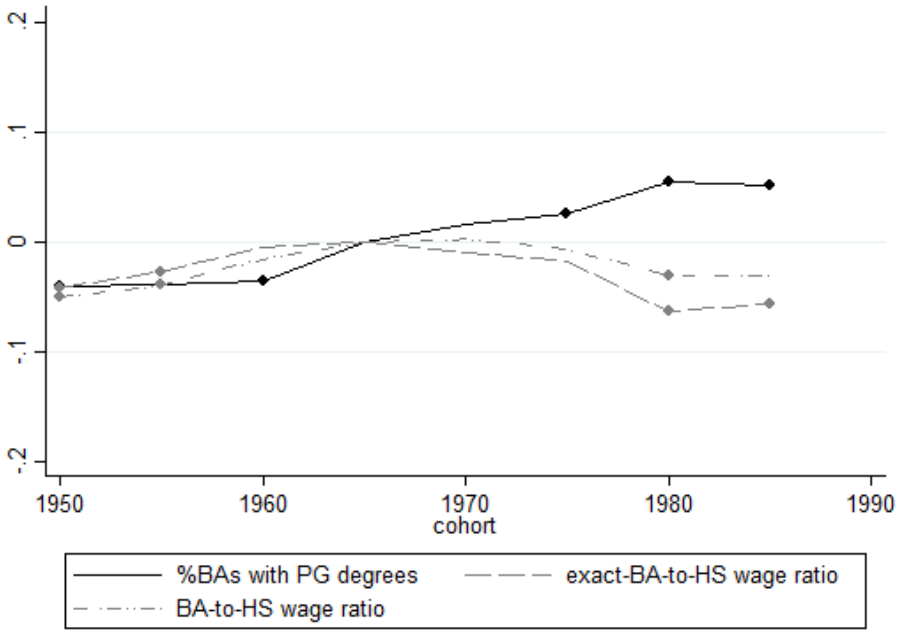

Figure 6: Cohort Effects for the Proportion of University Graduates with Advanced Degrees and the BA to HS Wage Ratio

The marked points are significantly different from zero.

smaller. We will return to this latter element of the data later in the paper.

3.2 Unobserved quality changes

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.

As higher education expands, it draws in pupils from a wider and 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 BA-HS wage differential is greater or smaller than the observed one.

The idea that BAs have a lower and wider range of quality after the higher education expansion has been advocated in OÆLeary 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 7. 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 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.

The exercise we just described is a somewhat indirect approach 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 provide a detailed discussion our bounding exercise in Appendix A.4, providing only a sketch of our approach here for brevity. Underlying this approach is an hierarchical model of ability. In this model, there is a single, unidimensional ability that is more productive the higher is an individual's education level. Under standard assumptions on

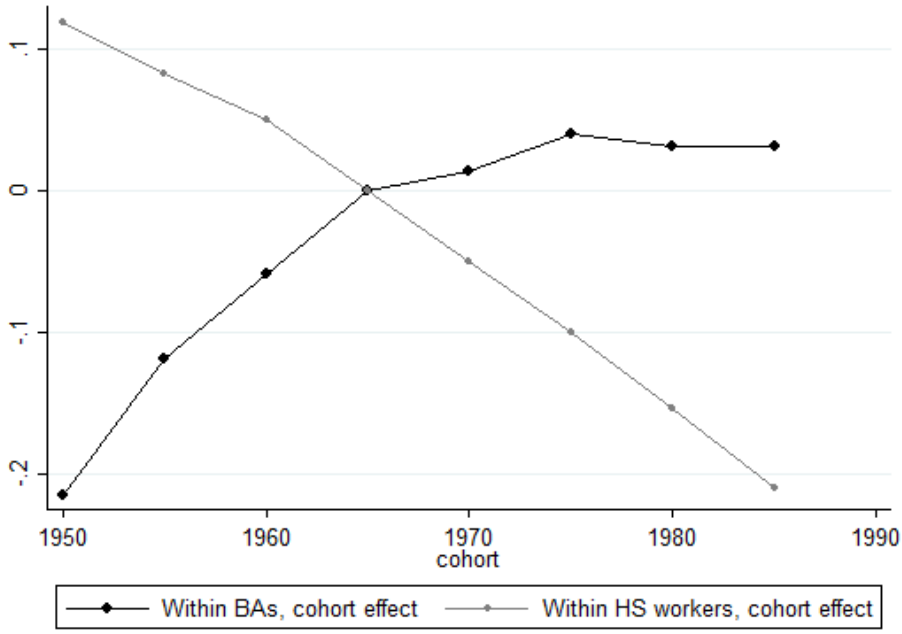

Figure 7: Ratio of median wage to the 10th percentile within education groups, UK

costs, higher ability individuals sort to higher levels of education. In this situation, there is a set of individuals (or, more properly, ability levels) who choose to go to university even in the pre-expansion period when universities were more costly to access. With the expansion of the university system these "university stayers" continue to get a higher education but they are joined by a set of "university joiners" who have been induced to enter university by the declining costs. Thus, the pre-expansion wage distribution for BA's consists only of university stayer wages while the post-expansion BA distribution includes both stayers and joiners. We have no way of identifying who is a stayer and who is a joiner in the post-expansion distribution, but by making extreme assumptions on which workers are joiners, we can construct extreme bounds on the mean or median wages for stayers. Comparing those bounded values to the mean or median wage for BA's before the expansion (who, remember, consist only of university stayers), we get bounds on movements in the mean or median university stayer wage. Since the stayers are a consistent group over time, these bounds reflect wage movements for a composition constant group.

We can make one of two extreme assumptions in order to form bounds. In the first, the 'joiners' are the lowest wage earners in the post-expansion cohort wage distribution. Thus, the 'stayers' wage distribution can be obtained by trimming from the lower tail

of the observed wage distribution the proportion by which the set of university educated workers has expanded between the two cohorts (where the proportion is expressed as a proportion of the post-expansion set of BA workers). At the other extreme, the 'joiners' would be better workers. But as Gottschalk et al. (2014) show, under a standard Roy model, the 'joiners' can be at best as good as the 'stayers'. If they were better then they would already have entered the BA sector. Thus, the other bound is the actual observed post-expansion distribution. Based on this, we can form one bound on wage changes relative to a base cohort for a composition constant group (the 'stayers') by trimming from the bottom of the distribution in expansion periods. The other bound is the actual observed change. Performing an analogous exercise with high school educated workers, we can form bounds on movements in the high school wage and on the BA-HS ratio.[13]

We repeat the sample trimming exercise for each cohort using the 1965-69 cohort as the base of comparison in each case. The resulting quality-adjusted wage differentials are reported in the left panel of Figure 8.[14] The right panel shows cohort effects derived from these profiles 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 and it did not decline between these cohorts, 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.

[13]In forming the ratio, use 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.

[14]Note that the line for the 1965-69 cohort is the actual data since it is used as the benchmark.

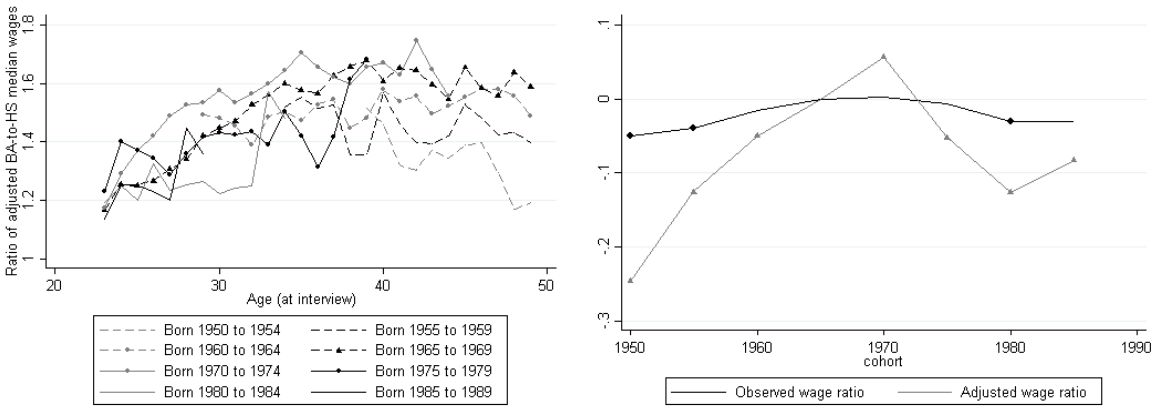

Figure 8: Bound on Changes in the UK Median BA-to-HS Wage Ratio Relative to 1965 Cohort

For each age and cohort, we form a bound on changes in median wages relative to 1965 as described in the text. The left panel shows wage differential by age profiles when each cohort is assumed to have the same BA proportion as the 1965-69 cohort. An expansion in the proportion BA relative to that cohort is assumed to occur with additions solely at the bottom of the BA wage distribution. Trimming BA and HS distributions under this assumption provides bounds on composition constant median wages. The right panel shows the cohort effects from this exercise along with the cohort effects from unadjusted wages, which form the other bound on cohort effect changes relative to the 1965-69 cohort. The marked data points are significantly different from zero at 5% level.

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. It does, though, reinforce our earlier observation that there may have been a decline in the ratio for the cohorts following the large educational expansion cohorts. It is worth pointing out, as well, that for selection to provide the explanation for why the observed age profiles match so closely across cohorts, selection would have to vary across cohorts and by age in a manner to just precisely offset declines in the wage ratio for a composition constant set of workers that would be caused by the increasing relative supply of BA workers. 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 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.[15]

[15] In the appendix, we present the results of an exercise in which we ask how much selection would be needed to make the relationship between UK relative education supplies and wage differentials match what one would predict based on a Card and Lemieux (2001) estimation exercise based on US trends.

4 A Model of Educational Changes, Technological Change and Decentralization

At this point, we have established that for the cohorts entering the labour market between the early 1990s and the early 2000s, the UK experienced a substantial increase in educational attainment but no change in the BA-HS wage differential. In this section, we set out a model that can capture these patterns as an outcome of the functioning of an economy undergoing a reduction in individual education costs in the presence of different technological options. As we discussed in section 2, the focus on firms choosing among different technologies is a conscious one that we see as fitting with our argument that the UK is a technological follower in this period.[16] The model is designed to reflect key insights in the recent literature on technological change: specifically, that the technological change has been accompanied by shifts toward decentralized organizational forms and that IT has replaced routine tasks. In our exposition, our emphasis will be on the decentralization element since this is the relatively more novel part of our model.

The idea that technological innovations or adoptions depend on relative wages or skill supply has been around for some time (Habakkuk (1962); 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

Doing that makes more precise that what one would need a whole set of offsetting selection changes of different sizes in order to make the experiences of the two economies compatible.

[16] Acemoglu and Zilibotti (2001) set out a model of endogenous invention with a Northern country where innovation happens and a Southern country that is a technological follower. In their model, on the balanced growth path for the North, an increase in the ratio of skilled to unskilled workers generates no change in the skilled to unskilled wage ratio. Thus, relative skill growth exactly offset by skill biased innovation could explain our main pattern. However, as we argued in Section 2, we believe there is good reason to see the US as the technological leader in this period and, so, it should be the US rather than the UK that should experience a flat skill wage ratio, which is not the case. Acemoglu and Zilibotti (2001) mention the skill wage ratio invariance result from their model but do not investigate it empirically.

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 fit with a model of endogenous technological adjustment (Lewis (2011); Dustmann and Glitz (2015)).

The general framework we wish to consider is one in which firms can choose to produce a single output either with a centralized (C) technology or a decentralized (D) technology. Having a single output is intended to emphasize the nature of these technologies as general purpose technologies that could be applied to the production of any product. Following Rosen (1978) and Borghans and ter Weel (2005), we will characterize production in engineering terms as having a Leontieff form in which a continuum of tasks, x, defined on the unit interval are required to produce an output.[17] The amount of each task required to produce one unit of output is given by the continuous function, $\alpha(x)$, $x \in [0,1]$. The tasks are performed by two types of workers: U (unskilled) and S (skilled). Total hours of work are inelastically supplied by each type of worker. Workers of each type are described by capacity functions, $\tau_l(x)$, which are continuous functions defined on [0,1] determining the amount of time a worker of type $l = U,S$ needs to produce the amount of task x required for one unit of output. Expressing capacities in this way is a bit awkward but will ease the writing of the cost functions later. Further, we assume that the tasks are ordered from least to most complex and that S workers have comparative advantage in more complex tasks, i.e., $\frac{\tau_S(x)}{\tau_U(x)}$ is decreasing in x.

Rosen (1978) shows that based on such a specification, one can derive a production function defined over n_s and n_u (the number of hours of S and U labour used, respectively) in which the firm allocates a given amount of S and U to each task in order to maximize output. In particular, firms will allocate skill groups according to their comparative advantage in the sense that there will be a task ρ such that all tasks, $0 \leq x \leq \rho$ are assigned to U workers and, conversely, all tasks $\rho < x \leq 1$ are assigned to S workers. Further, ρ is declining in $\frac{n_s}{n_u}$. Thus, if the relative number of S workers is small then they will only be assigned to the most complex tasks and as that relative number grows, they will be moved progressively further down the list of tasks ranked by complexity. The marginal rate of technical substitution between S and U at given values n_u and n_s equals $\frac{\tau_S(\rho(n_u,n_s))}{\tau_U(\rho(n_u,n_s))}$, where we have written ρ as a function of n_u and n_s. Thus, profit

[17] This general form for production has become somewhat common in models of technological change, tasks, and polarization. For example, a variant of it is used in Acemoglu and Zilibotti (2001), and Acemoglu and Autor (2011) use this approach to provide a framework for interpreting existing research on tasks and technological change. Our model differs in the way we introduce decentralization and in our assumption that firms can choose between two such technologies.

maximizing firms will hire numbers of hours of U and S labour to equate the marginal rate of technical substitution to the wage ratio, $\frac{w_S}{w_U}$ (where, w_S and w_U are the skilled and unskilled hourly wages, respectively, and where we assume that the output price is 1), allocating those hours optimally according to comparative advantage over the tasks required to produce. The result is a production function that reflects the efficiencies from taking account of the comparative advantage of the two types of workers and which is, itself, not necessarily Leontieff in form. In this sense, the ultimate production function reflects more than just the engineering 'recipes' since it includes the optimal allocation of workers across the task combinations specified in the recipes.

As Rosen (1978) demonstrates, and as we draw in figure 1, in the case with two types of workers (our case) the unit output isoquant intercepts both axes. The intercept on the N_u (number of unskilled workers) axis equals $\int_0^1 \tau_U(x)dx$. As we move away from that intercept to the left, we begin to introduce S workers, replacing the U workers in the most complex tasks. Thus, the slope of the isoquant is given by $\frac{\tau_S(\rho(n_u,n_s))}{\tau_U(\rho(n_u,n_s))}$ and comparative advantage dictates the standard convex shape. The N_s intercept is given by $\int_0^1 \tau_S(x)dx$.

We will consider an economy with two possible 'recipes' or technological forms. The first is centralized and takes the form as set out above, where we will now write the technological requirements function as $\alpha^C(x)$ and the amount of time a worker needs to complete the number of tasks needed for a unit of output will also now be indexed with a superscript, C, i.e., $\tau_l^C(x)$. In order to match patterns in the data, we delineate management tasks from other tasks. In the centralized technology, management tasks are necessary in order to co-ordinate the other tasks and the producers of the other tasks just focus on production of their part of the process, leaving communication and co-ordination to the managers. We will arbitrarily denote tasks on the interval $[\theta,1]$ as management tasks. In the centralized technology, to keep the exposition simple, we will assume that the α and τ functions are continuous from above and below at θ.

The alternative technological form is decentralized. Caroli and Van Reenen (2001) describe modern organizational forms as being 'delayered' with 'some decision-making being transferred downstream.' Multi-tasking is also an important feature of this organizational form with the benefits that the firm becomes more flexible and managers have to spend less time monitoring and co-ordinating workers (Bloom et al. (2014)). Thus, rather than having workers performing physical tasks without regard to others and having a manager who co-ordinates the outcome, in a decentralized form, workers both produce and co-ordinate with other task producers. As a result, less of the pure management task is needed. All of this is made possible by (i.e., is complementary with)

IT technological change, which reduced the cost of diffuse information transfer.

We capture the differences in the decentralized form relative to the centralized form, first, by assuming that there is a lower requirement for the pure management tasks in the new form:

$$\alpha^D(x) = \lambda\alpha^C(x), \forall x \geq \theta \tag{1}$$

where, the D superscript denotes the decentralized technology, and $\lambda < 1$. For simplicity, we will assume that the requirements for the other tasks remain the same, i.e., $\alpha^D(x) = \alpha^C(x), \forall x < \theta$.

Following much of the literature on technical change and the labour market, we also assume that skilled workers are better at working with the new organizational form (Caroli and Van Reenen (2001); Bresnahan et al. (1999)). We represent this by assuming that skilled workers are perfect multi-taskers and can perform each of the non-management tasks in the same amount of time as before, performing the new, associated communications while they are doing them without extra effort (thanks to IT). For unskilled workers, performing each non-managerial task now requires more time since working with the new IT is more difficult for them. Further, skilled workers are able to take advantage of the new technology in management tasks while unskilled workers are not. Thus,

$$\tau_S^D(x) = \tau_S^C(x), \forall x < \theta \tag{2}$$

and

$$\tau_U^D(x) = \gamma\tau_U^C(x), \forall x < \theta \tag{3}$$

with $\tau_U^D(x) = \tau_U^C(x), \forall x \geq \theta$, $\tau_S^D(x) = \lambda\tau_S^C(x), \forall x \geq \theta$ and $\gamma > 1$. We view this specification as capturing the notion of decentralization in papers such as Lindbeck and Snower (1996); Caroli and Van Reenen (2001); Bresnahan et al. (1999), and Bloom et al. (2012) that it is an organizational form in which decision making and communications are spread throughout the firm rather than being done by a small cadre of managers. We could allow for decentralization forms in which communication and decision making are differentially allocated across tasks but elect for the simpler form in which they are essentially allocated evenly across the non-manager tasks for expository clarity.

The literature emphasizes that decentralization has been enabled by the advent of IT. Much of the recent work on IT and the labour market also emphasizes impacts of the new technology in replacing routine tasks that tend to lie in the middle of the wage distribution. Following Acemoglu and Autor (2011) and Borghans and ter Weel (2005),

we can model this effect by having the α values in middle tasks substantially reduced under the new (D) technology. Essentially, the idea is that IT capital performs those tasks and, thus, less labour is required in them. As described in Acemoglu and Autor (2011), the result will be a polarization in employment, with relatively more employment in low and high complexity jobs compared to those in the middle. However, this will not alter our main points about movements in educational wage differentials set out below. For that reason, we will not explicitly include the reductions in middle α's in our analysis for simplicity. We will, however, highlight implications from such reductions at relevant points in our discussion.

Given this setup, if there were only U workers in the economy then all firms would use the C technology since it would be cheaper at any given unskilled wage. Conversely, if there were only S workers in the economy, firms would use only the D technology. But we will start by assuming that the endowment of S and U workers in the economy is such that both technologies are in use (returning to the conditions under which that is true momentarily). We also assume that these are general purpose technologies that can be used for producing any good. Thus, to simplify, we assume both are used to produce a good which is the numeraire. Assuming free entry of firms, that implies two zero profit conditions:

$$1)1 = w_U \int_0^{\rho_C} \tau_U^C(x)dx + w_S \int_{\rho_C}^1 \tau_S^C(x)dx \tag{4}$$

$$2)1 = w_U \int_0^{\rho_D} \tau_U^D(x)dx + w_S \int_{\rho_D}^1 \tau_S^D(x)dx \tag{5}$$

where, w_U is the unskilled wage, w_S is the skilled wage, ρ_C is the task dividing the U from the S tasks for technology C and ρ_D is the threshold task for the D technology.

Several key points follow from these two equations. First, together they imply a factor price invariance result as in standard trade theory. Because ρ_C and ρ_D are determined by the equality of the wage ratio to the marginal rate of technical substitution (MRTS) in profit maximizing firms and the MRTS is given by $\frac{\tau_S(\rho)}{\tau_U(\rho)}$ (i.e., is technologically determined), everything on the right hand side of both equations can be written as functions of w_U and w_S. That, combined with the assumption that these are general purpose technologies and so are producing the same good with the same price, implies that we have two equations in two unknowns (w_S and w_U). We show the solution diagrammatically in Diagramme D1. The figure shows the unit output isoquants for the two technologies. The isoquant for the centralized technology intersects the number of unskilled workers (N_u) axis at $n_{u0}^C = \int_0^1 \tau_U^C(x)dx$, i.e., the total number of hours to produce one unit if only

unskilled workers are being used. Similarly, its N_s axis intercept is $n_{s0}^C = \int_0^1 \tau_S^C(x)dx$. The unit isoquant for the decentralized technology has a larger N_u intercept because of our assumption that unskilled workers take longer to do non-managerial tasks because of the requirement to communicate as well as produce but get no advantage in terms of the time they require to perform management tasks. In contrast, under the decentralized technology, skilled workers require no extra time to do non-production tasks and can take advantage of IT to spend less time on managerial tasks. The result is an isoquant with a larger N_u intercept, a smaller N_s intercept and a lower slope at all values of x than the C isoquant.[18] Given the continuity assumptions and the comparative advantage assumption, the isoquants will cross once. That, in turn, implies that there will be a single unit cost line that is just tangent to the two isoquants, i.e., a single pair of w_S and w_U values at which both technologies are in operation.

Diagramme D1: Wage Setting with Two Technologies

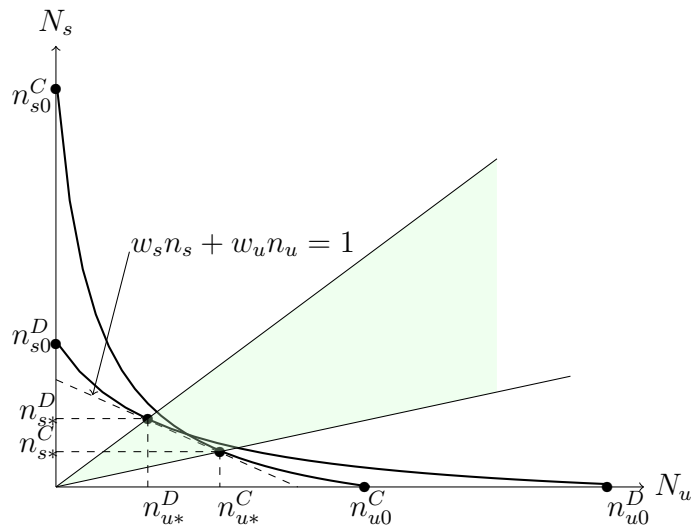

Diagramme D1 is, of course, a standard trade diagramme with two technologies instead of two sectors, and the same conclusions follow here as in the simple trade case. Our assumptions about the two technologies implies that the C technology will be relatively U intensive and in an equilibrium in which both technologies are used, n_{u*}^C and n_{s*}^C hours of unskilled and skilled work, respectively, will be used to produce a unit of the output with this technology. Similarly, n_{u*}^D and n_{s*}^D hours of unskilled and skilled work will be used with the D technology. Rays defined by $\frac{n_{s*}^C}{n_{u*}^C}$ and $\frac{n_{s*}^D}{n_{u*}^D}$ are drawn as diagonal lines in

[18] To make the exposition simpler, we assume that $\lambda \cdot \gamma = 1$. This implies that the isoquant is smooth at task θ. Without it, there would be a kink in the isoquant that would complicate the exposition but not the ultimate conclusions.

the figure. Those rays form the boundaries of the cone of diversification (the shaded area in Diagramme D1). As long as the ratio of skilled to unskilled hours in the economy falls within that cone, both technologies will be in use. If, instead, $\frac{N_s}{N_u} < \frac{n^C_{s*}}{n^C_{u*}}$ then only the C technology will be used. This is simple to see in the figure since on rays with lower slope than $\frac{n^C_{s*}}{n^C_{u*}}$, the cost line that is just tangent to the C isoquant will lie below the D isoquant, implying that it is less costly to produce just with C. Conversely, if $\frac{N_s}{N_u} > \frac{n^D_{s*}}{n^D_{u*}}$ then only the D technology will be used.

What is of most interest to us is the implications for wage movements when there are increases in the amount of S relative to U labour. Given that equations (4) and (5) have a unique wage solution and are not functions of labour quantities, as long as both technologies are in use, changes in the amounts of S and U in the economy do not alter the individual wages or their ratio. This is the standard factor price invariance result from trade theory. Firms in the economy react to larger relative amounts of S labour not by increasing the amount they use with any one technology but by shifting toward the more S intensive technology (D). In fact, it is straightforward to show that a given increase in the ratio $\frac{N_s}{N_u}$ generates a more than proportionate increase in output from the D technology.[19]

Following from this, the empirical implications from the model are as follows. First, if the two technologies are available and the skilled to unskilled labour ratio, $\frac{N_s}{N_u}$, is in the cone of diversification then increases in the ratio of skilled to unskilled labour does not alter the wage ratio, $\frac{w_s}{w_u}$, or the individual wages, w_s and w_u. Second, if $\frac{N_s}{N_u}$ rises enough then eventually all firms will adopt the D technology and then subsequent increases in $\frac{N_s}{N_u}$ will generate decreases in $\frac{w_s}{w_u}$ as in the standard one technology case. Third, assume that there are unskilled managers in the C technology before the increase in the skills in the economy (as is the case in our sample period), i.e., $\rho_C > \theta$. In that case, the ratio of the number of unskilled managers to skilled managers will decline as $\frac{N_s}{N_u}$ increases. This happens because U workers form a larger fraction of managers under the C technology (indeed, they may not be managers at all in the D technology given the comparative advantage set up) while S workers form a larger fraction of managers under the D technology. As the number of skilled workers rises, there will be a disproportionate shift toward the D technology that will imply more S than U managers overall even

[19] To see this note that we can write the ratio of S to U hours employed in the economy as a weighted average of the ratios employed in the two technologies, i.e., $\frac{N_s}{N_u} = \phi_C \frac{n^C_{s*}}{n^C_{u*}} + (1-\phi_C)\frac{n^D_{s*}}{n^D_{u*}}$, where, ϕ_C is the fraction of output generated using the C technology. If the economy is in the cone of diversification, as $\frac{N_s}{N_u}$ increases, the two technology specific ratios do not change but ϕ_C decreases. In fact, ϕ_C must decrease more than proportionally to maintain the equality.

though the proportion of each type of manager will stay the same within each technology. Fourth, as $\frac{N_s}{N_u}$ increases, the proportion of S workers who are managers decreases. This is somewhat surprising given that the economy is shifting toward a more S intensive technology where more of the management positions are held by S workers. Indeed, if only the C technology were available then as more S workers became available, comparative advantage would dictate that they would take up more and more of the management jobs until they had squeezed out the U managers altogether. That is, the proportion of S workers who are managers would increase. However, there is actually a smaller proportion of S workers who are managers with the D technology (since all S workers are managers in the C technology if $\rho_C > \theta$) and so as the economy shifts toward the D technology the proportion of S workers who are managers will fall. This is a reflection of the fact that in the decentralized technology, where S workers can both produce and communicate at the same time, S workers are used farther down into the task structure than in the C technology in equilibrium. Fifth, as $\frac{N_s}{N_u}$ increases and the economy shifts toward the D technology, we should see more workers in all parts of the production structure making decisions and communicating not just to their managers.

It is interesting to compare these implications to those from a more standard model with exogenous technical change. To see those implications, consider a model in which one technology is in use at a time. Output, Y, is produced according to the Cobb-Douglas production function:

$$Y = M^{\alpha} L^{1-\alpha} \tag{6}$$

where, M is hours of managerial labour, L is hours of production labour, and α is a parameter. Each task is performed by a combination of skilled and unskilled labour, with the labour aggregated through CES functions:

$$M = [aS_M^{\sigma} + (1-a)U_M^{\sigma}]^{1/\sigma} \tag{7}$$

and

$$L = [bS_L^{\rho} + (1-b)U_L^{\rho}]^{1/\rho} \tag{8}$$

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_M is the amount of skilled labour in the managerial task; and U_L is the amount of unskilled labour in the basic labouring task. 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. Market clearing in the labour market corresponds to the total number of workers with each skill level in the economy being equal to the sum of the numbers employed in the various occupations and technologies:

$$S = S_L + S_M$$

and,

$$U = U_L + U_M$$

Workers of each skill type can choose freely whether to work as a manager or a labourer and so there will be one skilled wage, w_s and one unskilled wage w_u.

In this framework, a skill-biased technological change can be represented as an increase in a, i.e., an increase in the productivity of S workers as managers. This captures both that the technological change favours S workers and that it is related to management tasks. Note that we are assuming that the technological change arrives exogenously and alters the production function of firms without them choosing whether or not to adopt the new technology.

To understand the impact of this change note that, working from the firm's first order conditions, it is straightforward to show that the wage skill ratio is,

$$\begin{aligned} \frac{w_s}{w_u} &= \frac{a}{1-a}(\frac{S_M}{U_M})^{\sigma-1} \\ &= \frac{b}{1-b}(\frac{S_L}{U_L})^{\rho-1} \end{aligned} \tag{9}$$

Rearranging these expressions slightly, we get:

$$\frac{a}{1-a}\frac{S_M^{\sigma-1}}{S_L^{\rho-1}} = \frac{w_s}{w_u} = \frac{b}{1-b}\frac{U_M^{\sigma-1}}{U_L^{\rho-1}} \tag{10}$$

In the context of this model, in order to match the main data pattern of an increase in S accompanied by no change in $\frac{w_s}{w_u}$, we need an increase in a of just the right size so that the skill biased demand increase just balances the relative supply shift. We view it as somewhat implausible that there were an exogenous set of technological changes that

just balanced the supply shifts over an extended period of time, but we cannot reject that this could have occurred. Instead, we ask about the further implications of such changes if this were the mechanism driving our main data patterns. Examining (10), note that if a increases then the ratio of the number of skilled workers who are managers to the number who are labourers must also increase in order to match the unchanging wage ratio. This is the opposite of the implication from our endogenous technological choice model in which the expansion in S is accompanied by a decreasing proportion of S workers who are managers.

5 Evidence on Model Implications

5.1 Macro Evidence

We begin our investigation of the relevance of our model of choice between a decentralized and a centralized organizational form by examining the model implications for economy level movements in wages, employment, and occupational shares. The first implication of the model is that the substantial increase in the proportion of workers with a university degree should have no impact on either the skilled to unskilled wage ratio or skilled and unskilled wages individually. In section 2, we showed that the skilled wage differential did not change for the cohorts that experienced the largest increases in educational attainment and that this pattern cannot be explained as a result of compositional shifts in terms of observed or unobserved worker characteristics. Of course, we chose the form of our model in part because it fit that pattern and so cannot claim it as proof in favour of the model. But, remarkably, 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 Figure 9, we plot the cohort coefficients from our standard specification run separately for high school educated and BA educated log real wages[20], 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.

[20] We deflate the nominal wage series using the GDP deflator with 2012 = 1.

Figure 9: Cohort effects for log real median wages by education groups, UK
To obtain the cohort values, 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.

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.

The other implications of the model at the aggregate level have to do with occupational composition. In particular, as the relative number of workers with BA's increases, management roles should be increasingly taken over by BA educated workers. Thus, the model predicts that the proportion of managers who have a BA should increase across cohorts. In Figure 10, we plot the proportion of managers who have a BA by cohort and

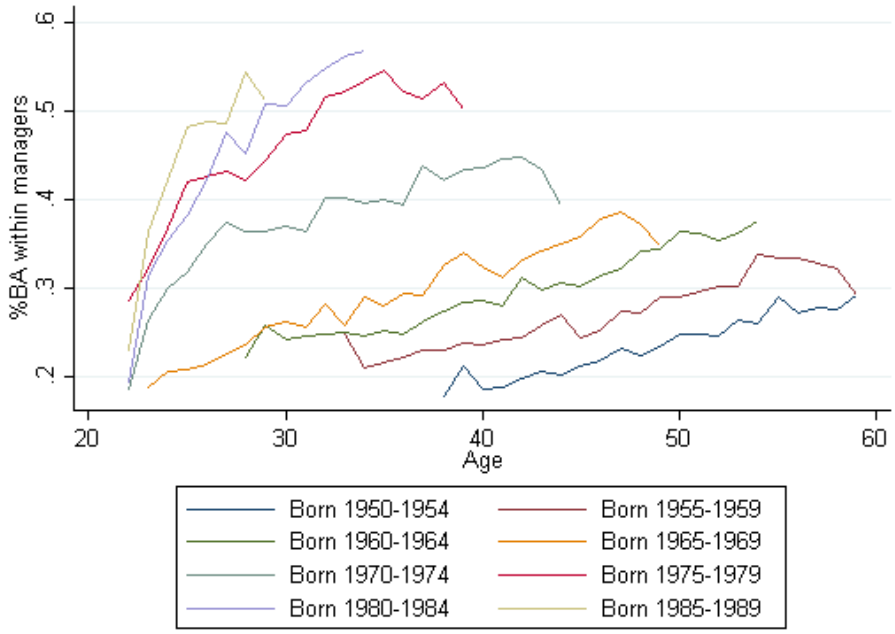

Figure 10: Proportion of Managers Who Have a University Degree, UK

age. Focussing on age 30, one can see a large shift in the direction predicted by the model: approximately 25% of managers had a BA in the 1965-69 cohort compared to nearly 50% in the 1975-79 cohort.[21] At the same time, the proportion of the BA educated workforce employed as managers should decline according to the model. In Figure 11, we plot the proportion of BA workers employed in management jobs. Concentrating again on age 30 in order to hold the position in the lifecycle constant, we see that 23% of BA workers in the 1965-69 cohort worked as managers compared to 21% in the 1975-79 cohort and 18% in the 1980-84 cohort. We argued earlier that this pattern in which an increasing proportion of BA workers in the workforce is accompanied by their taking on more of the management positions but the importance of management positions compared to other occupations falling for BA workers fits with our model but does not match the predictions of a standard model with an exogenous technological change favouring educated workers.

[21] The data underlying this figure and Figure 11 is from the LFS with managers corresponding to the first major group in the UK SOC2000, called 'Managers and senior officials'.

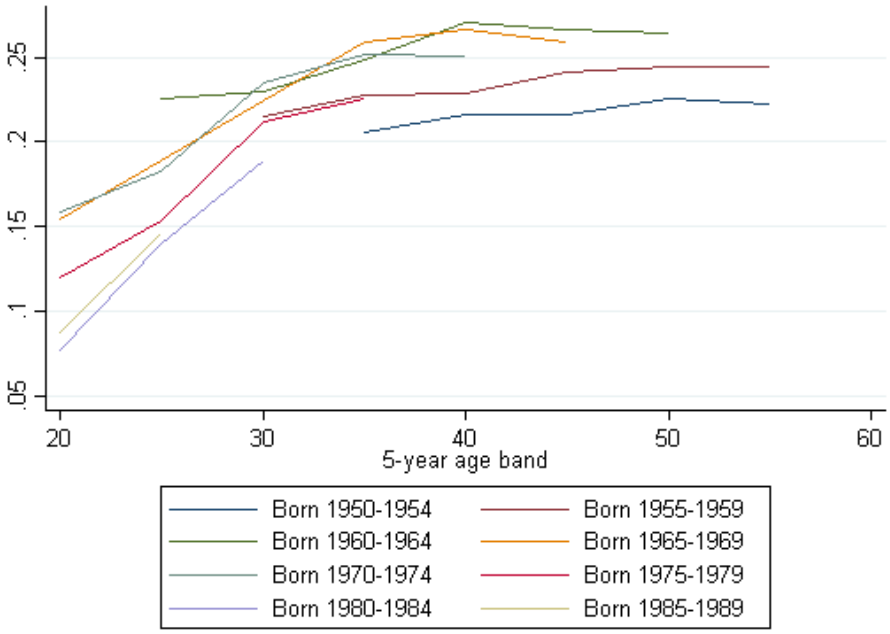

Figure 11: Proportion of University Degree Holders Who are Managers, UK

5.2 Micro Evidence

We turn, next, to using micro data to examine the main implication of the model: that firms in locations with larger increases in the relative number of educated workers make greater use of decentralized organizational forms. 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.[22] 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",

[22]The WERS surveys 25 employees per workplace. When there are fewer than 25 employees at the workplace, they are all given the questionnaire. the WERS is a representative survey and we incorporate its associated weight in all our calculations.

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 more influence. The index accounts for approximately 80% of the total covariance among the three questions. Finally, we normalize the influence index to have mean 0 and standard deviation 1 in the 3-wave-pooled sample.

Table 1 lists the overall mean and standard deviation of the influence index by WERS wave and education of employees. Across all firms, there has been a nearly 0.6 standard deviation increase in the mean influence index value between 1998 and 2011. Thus, there is a clear general trend toward decentralization of decision making. We examine differences between more and less educated workers in the lower panels of the table, presenting weighted averages with the proportion of workers at a firm in the particular education group as the weights. Doing that indicates that the increases in the index value were particularly large at lower educated firms. This makes sense since those are the firms that would most likely have used a centralized structure in the past and that, as a result, would have had the most leeway for adjustment.

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 influence index at the workplace level. "Local area" here refers to Travel To Work Areas (TTWA), which were developed to capture local labour markets using data on commuting flows in 1991.[23] There were around 300 such areas in the UK in the 1998 through 2011 period. 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.[24]

Table 2 reports the results from OLS 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 workplace. Given that our main variable of interest varies at the TTWA level, we cluster the standard errors

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

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

Table 1: Summary statistics of the influence index

Wave	Number of TTWAs	Number of workplaces	Mean influence index	Standard deviation
Influence index for all employees				
1998	204	1758	-0.32	0.92
2004	230	1657	0.024	0.96
2011	238	1917	0.27	1.025
Influence index for employees with degrees or above				
1998	190	1368	-0.018	0.79
2004	209	1272	0.001	0.83
2011	223	1557	0.13	0.80
Influence index for employees without degrees				
1998	203	1732	-0.34	0.88
2004	229	1620	0.024	0.93
2011	235	1823	0.25	1.047

Note: for each education group or for 'all employees', we first calculate 4 minus the first principle component of the three influence scores (ranged $1-4$). We then normalize that variable to have mean 0 and standard deviation 1 in the 3 wave pooled sample for the education group or for 'all employees'. Workplaces are weighted by the establishment's employment weight times the proportion of employees of that workplace in that education group. If a workplace has no employees of the labelled education group responding to the influence questions in the employee survey, the workplace is not counted in the sub-table for that labelled education group.
Authors' analysis of the UK Workplace Employment Relations Survey.

at that level. 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 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 0.09 standard deviation 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.

In the next set of columns, we check the robustness of this result across a series of specifications. In the 2nd column, we condition on the current HS proportion in the area, and the coefficient on the BA proportion changes very little. Thus, what matters for decentralization is the proportion of higher educated workers not more versus fewer high school drop-outs among the less educated. In the third column, we introduce controls for

Table 2: Workplace-level regressions of the influence index

	(1)	(2)	(3)	(4)	(5)
Current %BA	0.92***	0.95*	1.022***	1.01***	0.94***
in the TTWA	[0.17]	[0.50]	[0.16]	[0.16]	[0.30]
Wave=2004	0.30***	0.30***	0.28***	0.44***	0.30***
	[0.049]	[0.045]	[0.047]	[0.14]	[0.050]
Wave=2011	0.47***	0.47***	0.42***	0.77***	0.47***
	[0.062]	[0.066]	[0.064]	[0.22]	[0.064]
Current %HS people		0.047			
in TTWA		[0.70]			
Dummies for workplace size, organization size and industry (1digit)	No	No	Yes	Yes	Yes
Full interactions between industry (1digit) and wave	No	No	No	Yes	No
Including London	Yes	Yes	Yes	Yes	No
Observations	5,332	5,332	5,332	5,332	4663
R-squared	0.064	0.064	0.155	0.165	0.156

All regressions are at the workplace level, with standard errors clustered at the TTWA level. Each workplace is weighted by its employment weight. For both the BA proportion and the HS proportion at TTWA, the variable is the proportion of economically active people (working or unemployed) in that education group from the LFS in the two calendar years prior to the year of the dependent variable. For example, for workplaces observed in 2004, the BA and HS proportions are from the LFS 2003-04. Authors' analysis of the UK Workplace Employment Relations Survey.

industry, workplace size, and size of the organization.[25] 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 the fourth column, we further include interactions between industry and wave and the key estimate remains essentially unchanged. Finally, we are concerned that our results are being driven primarily by London as a potential outlier which contains a large number of observations and has both high education and high use of more modern technologies. In response, in column five we present a specification omitting London. The estimated effect of the BA proportion variable is very similar to the effect when we include London and thus,

[25]More 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+.

London is not driving our main results.

In Table 3 we report results with the dependent variable generated either only from the responses of the BA employees or only from the non-BA employees' responses. The specification includes industry, size, and year effects as in column (3) of Table 2 and we try both weights based just on establishment size and weights based on the size along with the proportion of workers in the specific education group. The results indicate that the positive correlation between BA proportion and employees' influence at workplace observed in the earlier specifications is not a mechanical result from a combination of BAs having more influence than non-BAs and an increasing proportion of BAs. In fact, the influence over work decisions reported by non-BA employees in their workplace is even more positively correlated with the local supply of BAs than for BA employees. Again, this fits with the idea that under the older, centralized organizational form, BA employees would have had managerial or quasi-managerial roles and, thus, some control over decision making. It is the non-BA's who will experience the greatest change in the shift to a decentralized workplace.

Table 3: Influence by education of employees

	Influence Reported By			
	BA employees	non-BA employees	BA employees	non-BA employees
Weighted by	Establishment employment		BA employment*	non-BA employment*
%BA in the TTWA	0.60***	0.86***	0.24	0.93***
	[0.21]	[0.24]	[0.20]	[0.29]
Wave=2004	-0.011	0.28***	0.0088	0.29***
	[0.047]	[0.040]	[0.041]	[0.0429]
Wave=2011	0.021	0.42***	0.092**	0.43***
	[0.052]	[0.052]	[0.045]	[0.059]
Observations	4,197	5,175	4,197	5,175
R-squared	0.08	0.121	0.062	0.143

The set of controls in each regression is the same as in column (3) in table 2.
* In the first two columns, each workplace is weighted by its employment weight. In the last two, the weight is multiplied by the proportion of employees in that education group. If a workplace does not have any employee in the education group responding to the influence questions in the employees' survey, it is omitted from the respective regression.
Source: Authors' analysis of the UK Workplace Employment Relations Survey.

Whether the estimated association between the local 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.

To address any remaining endogeneity, we adopt two approaches.[26] The first is to include the value of the dependent variable (the mean value of the influence index) in the first year for which we have it (1998).[27] 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 mean level of worker control over their workplace.[28] Thus, we compare two regions with the same initial level of use of decentralized organizational forms as a means of holding constant a general proclivity to use such forms for time-invariant reasons and ask whether the region that had a greater increase in the proportion of workers with a BA saw a larger proportion of firms increase the extent of their decentralization. The results without industry and firm size controls are given in column (1) of Table 4 and the results including those controls are given in column (2). The estimated effect of the proportion BA is again highly statistically significant and takes a value of about two-thirds of the comparable estimates in the first and third columns of Table 1. Thus, the proportion BA variable is picking up longer term differences in the extent of use of decentralized forms to a limited degree and not enough to overturn our conclusion that increases in the proportion BA induces a movement toward those forms. Interestingly, the historical use of decentralized forms itself has only a weak relationship with future use of those forms in a region.

Our second approach 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

[26] We implement these approaches using data aggregated to the TTWA. Estimation using data at the firm level with clustered standard errors yields very similar results.

[27] Since we have to drop the first year of our data, we are left with firm observations across only two years of data.

[28] Direct fixed effect estimators yield erratic and ill-defined coefficient estimates which we interpret as arising from the shortness of our panel.

Table 4: Influence Index Regressions: Initial Value and IV Estimates

	(1)	(2)	(3) IV	(4) IV
Current %BA in the TTWA	0.62***	0.65*	1.08***	0.49
	[0.20]	[0.23]	[0.37]	[0.40]
Wave=2004	-0.21***	-0.17***	0.28***	-0.19***
	[0.049]	[0.048]	[0.057]	[0.050]
Wave=2011			0.41***	
			[0.075]	
Influence Index in 1998	0.025	0.051		0.058
	[0.052]	[0.047]		[0.047]
Dummies for workplace size, organization size and industry (1digit)	No	Yes	Yes	Yes
Including London	Yes	Yes	Yes	Yes
Observations	390	390	672	390
R-squared	0.119	0.278		
First Stage F-stat			15.35	17.46
Associated p-value			0.000	0.000

All regressions are at the TTWA level, weighted by employment, with standard errors clustered at the TTWA level. The instruments in columns (3) and (4) are: the population share of the 1970-74 birth cohort; the population share of the 1975-79 birth cohort; the proportion of BA educated individuals in the parents' generation; and the proportion of GCSE/O-level holders in the parents' generation. All the instruments are measured at the TTWA level in 1995-6.
Authors' analysis of the UK Workplace Employment Relations Survey.

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. Thus, instead, 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.[29] 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

[29]The denominator for the proportions is the population born between 1940 and 1979.

predicted to have a more educated population later to the extent that people have some tendency to stay where they grew up. In addition, we use the educational composition of people in the generation who would likely be the parents of these cohorts (people born between 1945 and 1954). In particular, we construct the proportion of the parental generations who have a BA and the proportion who have a GCSE/O level, again measured in 1995-96. We also include the interaction of these parental education variables with the size of the 1970s birth cohorts in the area. The idea behind the instruments is that areas that one would predict to have a large increase in the proportion of BAs in their workforce between the early 1990s and the early 2000s are ones where there is a local baby boom in those generations and where the parents own education indicates that they would be interested in their children's education. For this set of instruments to be valid, we require that parents in the previous generation - and, in particular, more educated parents - did not have a tendency to have more children in areas which would later turn out to have more decentralized organizational structures. We also require that the parents did not locate in an area because it would undergo a shift toward a more decentralized organizational form several decades later, as part of a shift to a technology that did not even exist at the time at which most of them made their location choice. The fact that we control for industry and firm size effects in these regressions eliminates any concern that their location might have been related to persistent concentration in industries that would ultimately favour decentralization. We view the conditions under which this instrument set fails as very stringent. In particular, we find it hard to come up with situations in which differences in cohort sizes across areas are determined by the conditions that would affect the adoption of decentralized organizational forms decades later, especially after we control for industrial structure. The set of instruments are highly significant in the first stage, with p-values associated with the F-statistic for the test of their exclusion being effectively zero.

Column (3) in table 4 contains the results from our IV specification. The estimated coefficient on the proportion BA is 1.08, which is very similar to the value estimated with OLS in column (3) of table 1. This fits with our belief that endogeneity is not a substantial concern once we control for industry and firm size effects. In the final column of table 4, we present a specification in which we both use instruments and control for the initial, 1998 value of the influence index in the area. The estimate of the proportion BA effect in this specification remains positive, though it is approximately half the size of the estimated effect from using IV on its own and not statistically significantly different from zero at any standard significance level. We view this specification as pushing the data

quite hard, given that we have to drop the first year of data. As a result, the estimate of the main effect is poorly defined and we cannot reject the hypothesis that it is equal to the estimate from either the initial value specification or the simple IV specification on their own.

Our overall conclusion from our estimates is that an increase in the proportion of the working age population with advanced education in a region causes firms in that region to increase their use of decentralized technologies, with the effect being on the order of a 10 percentage point increase in the percentage of the working age population with a BA generating a 0.1 standard deviation 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.

6 Alternative Explanations

Our model does not provide the only possible explanation for a pattern of large increases in education levels combined with unchanging educational wage differentials. In this section, we briefly consider two alternative explanations.

6.1 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 5). 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 6, 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.[30]

6.2 Exogenous Technological Change

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, one could explain the wage and education patterns by appealing to a standard story of Skill Biasted Technical Change (SBTC). We believe that such a story provides an unsatisfying explanation for the patterns we observe for three reasons. First, as we pointed out in section 4, a standard version of an SBTC model implies opposite movements in the proportion of BA's who become managers to our model and, as we show in section 5, the data fits with the implications from our model. Second, in Appendix A.2, we present an exercise in which we use the estimation framework from Card and Lemieux (2001) (CL henceforth) to examine wage and employment movements in the US and UK. We find that the canonical version of the model with a linear trend for the exogenous skill biased technical change does not fit the data for either country. This accords with conclusions drawn using US data in Card and DiNardo (2002), Beaudry and Green (2005), and Acemoglu and Autor (2011) among others. We find that we can only obtain theoretically sensible elasticities with US data for the period from 1993 to

[30] Dustmann 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 5: Share of BAs among 16-59-year-old employees by industry SIC1992, 1994-2014

	shares		from 1994 to 2014	
SIC92	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

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.

Table 6: 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

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.

2012[31] if we employ a 7th order polynomial in time to represent skill biased technical change. This, in itself, is an indictment of the model. But if we suspend disbelief and use the US based estimates of the exogenous skill biased demand trend and the elasticities in combination with UK relative educational supply shifts in our period, we predict a roughly 20 percentage point decline in the BA-HS wage ratio in the UK between 1950 and 1985 cohorts. In actuality, the wage ratio was no lower for the 1985 cohort than the 1950 one conditional on an estimated age profile. Moreover, the pattern of wage differential changes across cohorts is wrong: the predicted decreases are largest between the 1965 and 1975 cohorts when the education wage ratio is actually quite constant and smaller in the post 1980 cohorts when the actual wage differential starts to decline. Thus, a standard SBTC model with a common technological change across developed economies does not fit the UK data. One could make an exogenous SBTC model fit the data if the UK is allowed to have its own pattern of skill biased technology shifts. We obviously cannot reject this model since technological changes are defined in a way to fit the data. But those changes would have to be extremely variable across cohorts, speeding up to just the right extent when education supplies were at their largest and then slowing down more than the rise in education slowed thereafter. We view such a pattern as implausible, which is our third strike against the exogenous technical change model.

Of course, the model of exogenous SBTC embodied in CL is an older version which has been supplanted by models of technological change and polarization. While our model can incorporate polarization, to this point, we have not addressed the potential role of polarization in helping to understand our main data patterns. To look further into the role of polarization in the UK wage and employment structure, in Table 7 for 30-34 year olds, 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 (in the last column), there is either negative or very small growth in employ-

[31] These years were chosen to match our UK data.

Table 7: 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

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.

ment 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 SBTC theory will provide a useful lens through which to understand the specific wage and employment patterns we are examining.

7 Conclusion

In this paper, we highlight 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. The combination of increased educational supply and a lack of movement in the educational wage differential necessarily implies a skill biased demand shift across cohorts. Importantly, to generate a flat profile of the education wage differential across several cohorts requires the skill biased demand shifts to exactly offset increases in educational supply across those cohorts. Since the

supply increases varied substantially across cohorts, an explanation for this pattern that relies on an exogenous set of technological changes that just happen to be the right size to yield no relative wage changes appears to us to be implausible.

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 Hecksher-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 employment 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. This plus the invariance in the educational wage differential points toward models of endogenous technological choice in which firm choices in response to the relative factor supply change induce relative demand shifts. These models have two broad forms: ones in which changes in the relative supplies of factors of production induce invention of technologies that are intensive in the expanding factors or ones in which the supply change induces firms to choose skill biased technologies from among the set of existing technologies. We believe that models of induced invention may be relevant for the US in recent decades since it was in a position to be a technological leader in skill biased technologies by virtue of having a much more educated work force than other developed economies at the dawn of the computer era, Beaudry et al. (2006). In contrast, the UK underwent its educational expansion much later and, as a result, we believe it is plausible that it was a technological follower for this type of technology - following an induced technological adoption model rather than one of induced innovation.

More explicitly, we argue for a model for the UK 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. But the model also has further implications; most notably that the proportion of managers who have a BA should increase but the proportion of BA's who work as managers should decrease as the decentralized technology spreads. The latter is the opposite of the prediction from a standard skill biased demand model built around a nested CES production function. In addition, 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. We develop an instrumental variable strategy in which we instrument for area specific educational changes using differences in fertility in previous decades and parental education. We believe these instruments are very likely to be valid, based as they are on an assumption that parental decisions on fertility in the 1960s and 1970s did not arise from predictions of decentralized technologies coming to their areas in the 1990s and after. Again, it is important that we control for industry in all our specifications, implying that parents would have to make their guess about future technology use independently of the local industrial structure for our instrument to be invalid. The IV results indicate that increases in the education level in a local economy have a causal impact on the adoption of decentralized organizational forms by firms in that economy.

The key point we see as arising from this exercise is that the effects of technological change are not one size fits all. There are good reasons to believe the US has been a technological leader and there has been considerable study of the interactions of technological change and educational supply shifts in the US. The interesting question then becomes, can the experience of the US be generalized to other countries? The UK provides an interesting case study to examine this question. Its large expansion in education happened quickly and well after the main expansion for the US. Because of that, we believe that the UK provides evidence on what happens to technological followers as their conditions shift toward favouring the technologies that the leader has developed. We argue that during the transition period for a follower economy, one could observe no real impact on skilled wage differentials even though the economy was being substantially transformed. Our evidence lines up with this interpretation. We believe this calls into question approaches in which technological change effects are identified from commonalities in wage and employment movements across countries, with remaining differences assigned to differences

in institutions and difference in supply shifts. This does not mean that there are no commonalities across economies and that we should devolve to studying each economy in isolation. Instead, we view our results as indicating the need for a broader view of the impact of technological change - one which emphasizes the role of differences in movements in relative factor supplies in determining the point in the lifecycle of a technology at which each economy adopts it.

References

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

Acemoglu, D. (2007). Equilibrium bias of technology. *Econometrica 75*(5), 1371–1409.

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

Acemoglu, D. and F. Zilibotti (2001). Productivity differences. *The Quarterly Journal of Economics 116*(2), 563–606.

Antoncyzk, D., T. DeLeire, and B. Fitzenberger (2010). Polarization and rising wage inequality: Comparing the us and germany. Technical report, ZEW.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *The American Economic Review 103*(5), 1553–1597.

Beaudry, P., M. Doms, and E. Lewis (2006). Endogenous skill bias in technology adoption: City-level evidence from the it revolution. Technical report, National Bureau of Economic Research.

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

Beaudry, P. and D. A. Green (2003). Wages and employment in the united states and germany: What explains the differences? *The American Economic Review 93*(3), 573–602.

Beaudry, P. and D. A. Green (2005). Changes in us wages, 19762000: Ongoing skill bias or major technological change? *Journal of Labor Economics 23*(3), 609–648.

Bloom, N., R. Lemos, R. Sadun, D. Scur, and J. Van Reenan (2014). The new empirical economics of management. Technical report, National Bureau of Economic Research.

Bloom, N., R. Sadun, and J. Van Reenan (2012). Americans do it better: Us multinationals and the productivity miracle. *American Economic Review 102*(1), 167–201.

Blundell, R., A. Gosling, H. Ichimura, and C. Meghir (2007). Changes in the distribution of male and female wages accounting for employment composition using bounds. *Econometrica 75*(2), 323–363.

Borghans, L. and B. ter Weel (2005). The divison of labor, work organization, and technological change. Technical Report 1709, IZA.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (1999). Information technology, workplace organization and the demand for skilled labor: Firm-level evidence. Technical report, National Bureau of Economic Research.

Card, D. and J. E. DiNardo (2002). Skill biased technological change and rising wage inequality: some problems and puzzles. Technical report, National Bureau of Economic Research.

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

Carneiro, P., K. Liu, and K. G. Salvanes (2014). How does labor demand react to changes in the supply of skills? evidence from schooling reforms.

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

Carpentier, V. (2004). Historical statistics on the funding and development of the uk university system, 1920-2002.

Carpentier, V. (2006). Funding in higher education and economic growth in france and the united kingdom, 1921-2003. *Higher Education Management and Policy 18*(3), 1–26.

Chen, U. (2013). The impact of the expansion in post-secondary education on the university wage premium in taiwan. Technical report, University of Michigan.

Chevalier, A., C. Harmon, I. Walker, and Y. Zhu (2004). Does education raise productivity, or just reflect it?*. *The Economic Journal 114*(499), F499–F517.

Devereux, P. J. and W. Fan (2011). Earnings returns to the british education expansion. *Economics of Education Review 30*(6), 1153–1166.

Dustmann, C., T. Frattini, and I. P. Preston (2013). The effect of immigration along the distribution of wages. *The Review of Economic Studies 80*(1), 145–173.

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

Fortin, N. (2006). Higher-educaton policies and the college wage premium: cross-state evidence from the 1990s. *American Economic Review 96*(4), 959–987.

Goldin, C. and L. F. Katz (2008). *The Race Between Education and Technology*. Belknap Harvard University Press.

Gottschalk, P., D. A. Green, and B. M. Sand (2014). Taking selection to task: Bounds on occupational task prices for the us: 1984-2013. Technical report, University of British Columbia.

Green, F. and Y. Zhu (2010). Overqualification, job dissatisfaction, and increasing dispersion in the returns to graduate education. *Oxford Economic Papers 62*(4), 740–763.

Habakkuk, H. J. (1962). *American and British Technology in the Nineteenth Century: The search for labour saving inventions*. Cambridge University Press.

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

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *The Review of Economic Studies 76*(3), 1071–1102.

Lemieux, T. (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *The American Economic Review*, 461–498.

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

Lindbeck, A. and D. J. Snower (1996). Reorganization of firms and labor market inequality. *The American Economic Review 86*(2), 315–21.

Machin, S. (2001). 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.

Machin, S. and S. McNally (2007). Tertiary education systems and labour markets. *Tertiary Review*, 1–57.

Machin, S. and A. Vignoles (2006). *Education Policy in the UK. CEE DP 57.* ERIC.

Manski, C. F. (1994). The selection problem. In *Advances in Econometrics, Sixth World Congress*, Volume 1, pp. 143–70.

McIntosh, S. (2006). Further analysis of the returns to academic and vocational qualifications*. *Oxford bulletin of economics and statistics 68*(2), 225–251.

Milgrom, P. and J. Roberts (1990). The economics of modern manufacturing: Technology, strategy, and organization. *The American Economic Review*, 511–528.

OÆLeary, N. C. and P. J. Sloane (2005). The return to a university education in great britain. *National Institute Economic Review 193*(1), 75–89.

Rosen, S. (1978). Substitution and division of labor. *Economica 45*(179), 235–250.

Walker, I. and Y. Zhu (2005). The college wage premium, overeducation, and the expansion of higher education in the uk. *IZA Discussion Paper No. 1627.*

Walker, I. and Y. Zhu (2008). The college wage premium and the expansion of higher education in the uk. *The Scandinavian Journal of Economics 110*(4), 695–709.

A Appendix (for online publication)

A.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.[32] 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 deflator[33]. We are worried that student wages may include distortions related, for example, to co-op programmes and so drop all individuals

[32]For 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.

[33]Source: OECD

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 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).

A.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.[34] 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

[34]See 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 8: 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 $\widehat{1/\sigma_E}$
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

This assumes θ_t is linear in t. The sample excludes 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 elasticitity 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.
Authors' calculation of the UK Labour Force Survey and the Current Population Survey.

UK separately and for various subsamples defined by gender and wage measure (weekly versus hourly). Table 8 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.

Next, we allow the demand shift to follow a 7th-order polynomial in year. The estimation on US data of hourly wage gives us $\hat{\sigma^A} = 9.5, \hat{sigma^E} = 3.1$ and demand shifters that increase gradually over time as plotted in Figure 12. The same regression

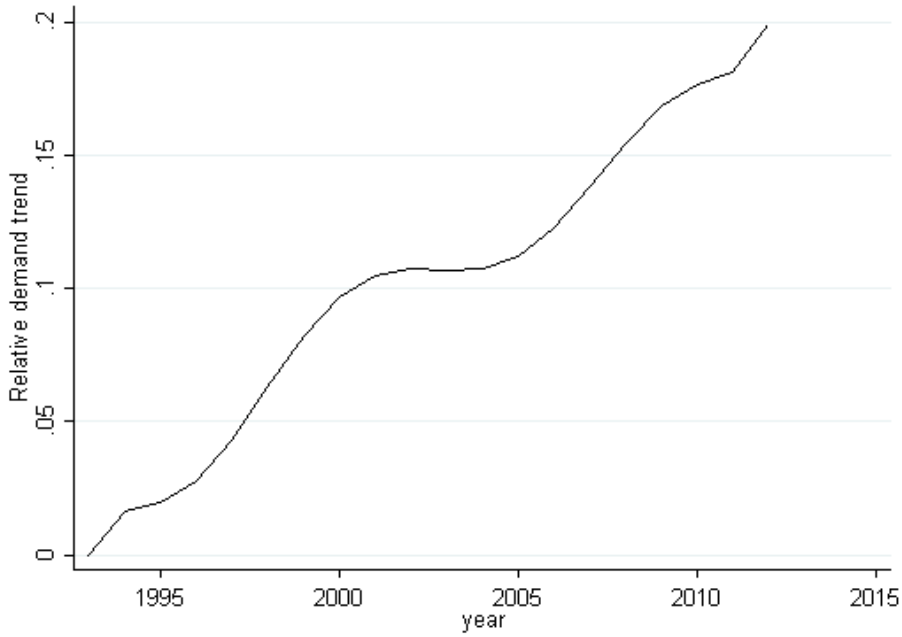

Figure 12: Estimated relative demand trend θ_t from US data
Authors' calculation of the UK Labour Force Survey and the Current Population Survey.

Table 9: Estimated labour substitution elasticities with non-linear demand trend

country	$\widehat{\sigma_A}$	$\widehat{(-1/\sigma_A)}$	s.e. of$\widehat{(-1/\sigma_A)}$	$\widehat{\sigma_E}$	$\widehat{(-1/\sigma_E)}$	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

θ_t is assumed to follow a 7th-order polynomial of t. The real wage measure is hourly wage. Both genders are included.
Authors' calculation of the UK Labour Force Survey and the Current Population Survey.

for the UK generates rather imprecise estiamtes as reported in Table 9. 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.

Next, we 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 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 joiners would have to earn above-the-median wages of the university stayers 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.[35] Comparing the predicted wages with the observed wages in each age group by year, we back out the proportion of joiners that must have earned more than the latent BA median in order to rationalize the difference.[36]

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.[37] 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.[38] 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

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

[36]We 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.

[37]In 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.

[38]The 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.

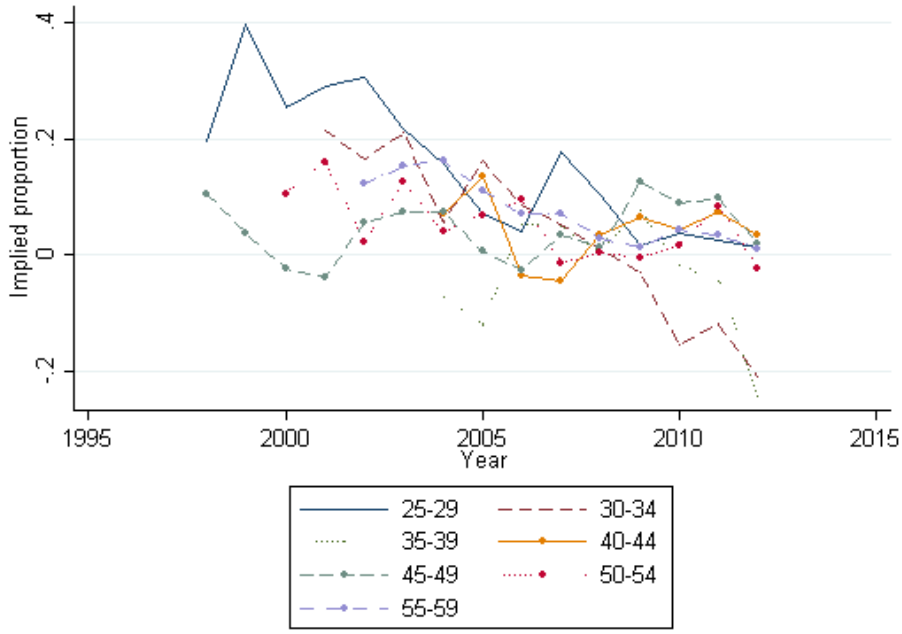

Figure 13: implied proportion of switchers ending above the latent BA median

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 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.

Similar to Figure 1, 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 14. 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)).

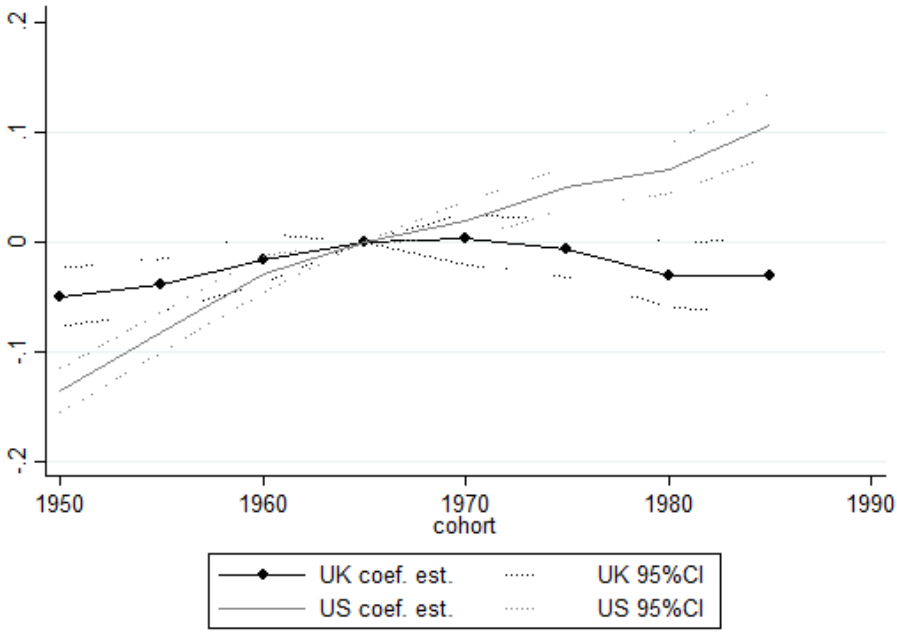

Figure 14: Ratio of BA median wage to that of high-school graduates, cohort effects
Wage is hourly. The cohort effect is normalized to 0 for the 1965-69 birth cohort. Source: Authors' calculation of the Labour Force Survey and Current Population Survey.

A.3 Observable compositional changes

In this appendix, we present added investigations into compositional change effects. The first relates to 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 15 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 overall proportion with a BA between the 1965-69 and 1975-79 cohorts, in particularly, 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 15 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.

The other observable composition dimension we investigated was between public and private sectors. Public sector employees are, on average, better educated and, with wages largely protected from direct market forces, we might expect wage differentials within the public sector to be more rigid. Given that, an expansion in the public sector might partly explain the patterns we have described. That possibility, though, falls short in two ways with respect to employment numbers. First, the proportion of workers in the

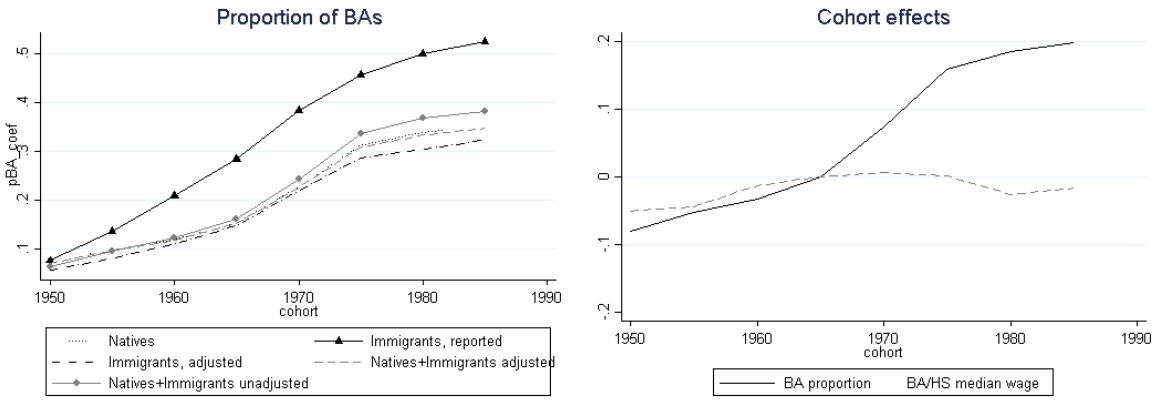

Figure 15: Cohort effects for the BA proportion and wage ratio, natives and immigrants
The first sub-figure is based on the same specification and normalization as in Figure 1. In the second sub-figure, the cohort effects are normalized to zero at the 1965 cohort, as in Figure 14.

public sector does not change substantially over our data period. Second, the growth in the proportion of workers with a BA is very similar between the private and public sectors.

The public-private sector dimension of movements in wage differentials is a bit more nuanced. In Figure 16, we plot the cohort effects for the median wage ratio constructed as before but only for private sector employees. As in Figure 14, which was for private and public sector employees combined, there is little change in the wage 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 wage differential across cohorts while the post-1979 cohorts show a more substantial decrease than in Figure 14. 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 drive the pattern in Figure 16 is a negative time effect in the wage differential affecting all cohorts in the same way in the last few years of our sample period. In Figure 17, 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

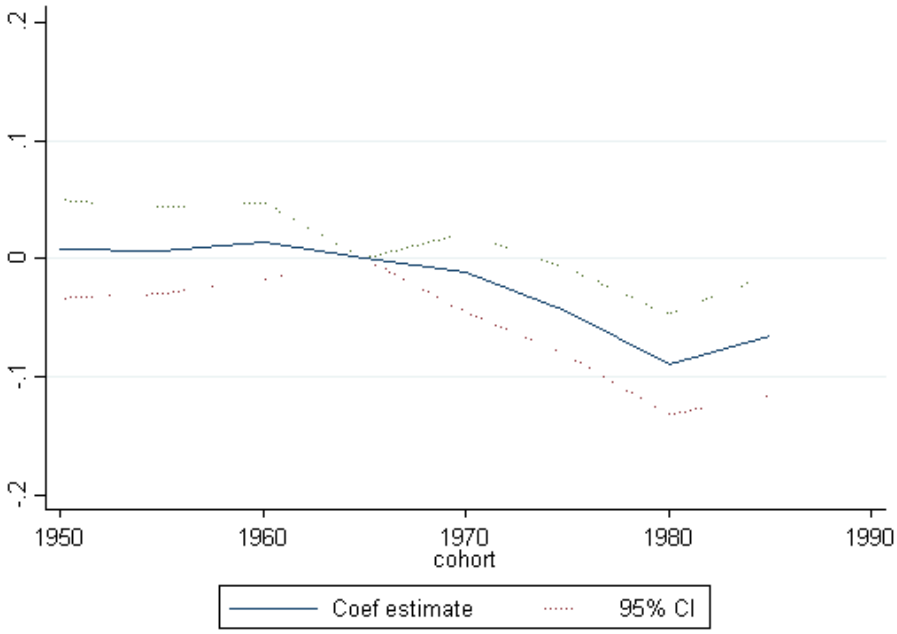

Figure 16: Cohort effects for the BA-to-HS wage differential, UK private-sector only
Cohort effects are derived in the same way as in Figure 14.

time effect disproportionately drags down more recent cohort effects. In fact, when we re-plot of Figure 16 using data only up to 2012 in Figure 18 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.

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 19, 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

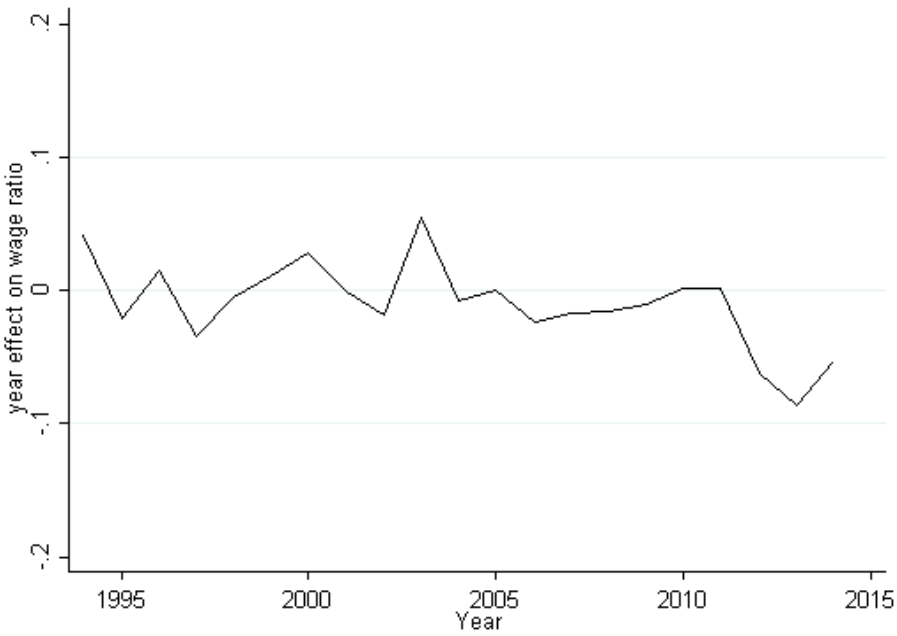

Figure 17: Time effect in BA-to-HS wage differentials, UK private-sector only
As normalization, the average time effect of the period till 2010 is set to zero.

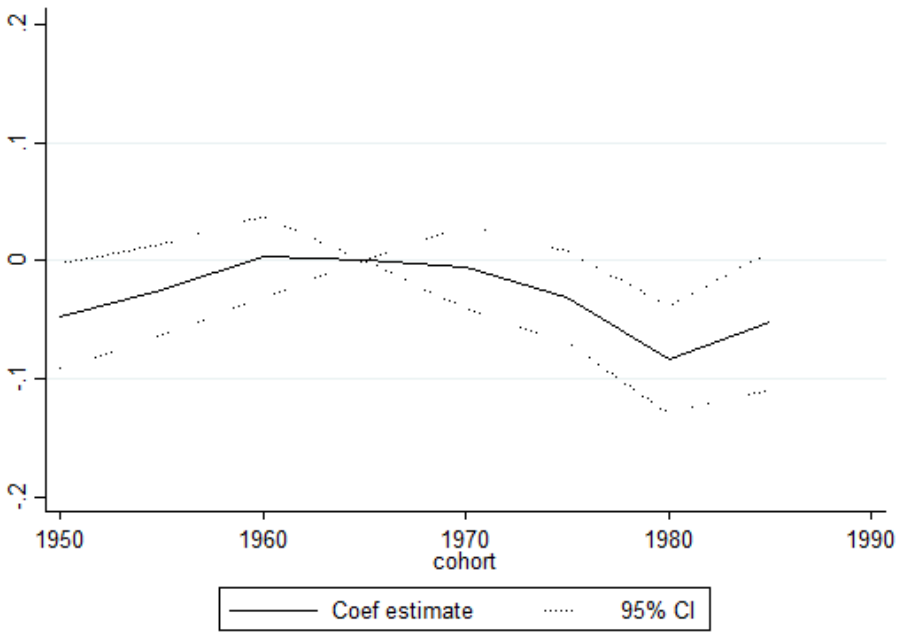

Figure 18: Cohort effects for the BA-to-HS wage differential, UK private-sector only and till 2012

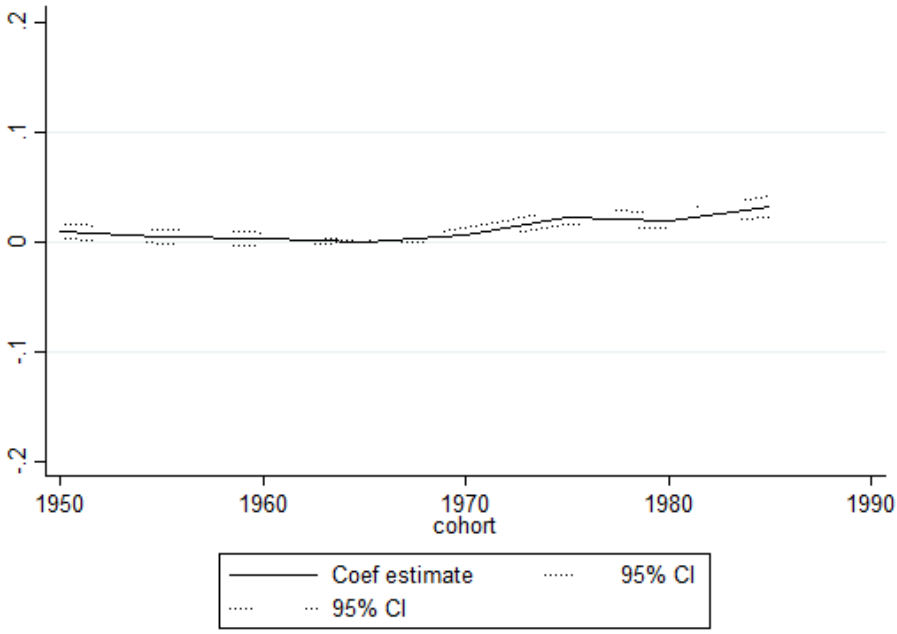

Figure 19: Difference in employment rate between BAs and HS workers, UK cohort effects

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.

A.4 Unobservable compositional changes: bounds

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:

$$\ln w_{ict} = \sum_{j=1}^{3} D_{ijt}\beta_{cj} + \sum_{j=1}^{3} D_{ij} f_{cj}(age_{it}) + \sum_{j=1}^{3} D_{ij}\lambda_j \eta_i + \epsilon_{ict} \tag{11}$$

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 are interested in the price per efficiency unit of workers with a given type of education ($\beta_{cj} + f_{cj}(age_{it})$ in (11)). This is unobservable because we do not observe the median wage for a composition constant group, Below we will adopt some assumptions and bounds on the composition-constant median wage for each education group.

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 hierachical 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 hierachical 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".[39] Obviously, the joiners' ability distribution lies above that of those who remain in the HS group in cohort c+1.

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

[39] Calling 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.

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 extreme bounds 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. The next two pages explain mathematically why the trimming method and the observed median are THE upper and lower bounds under the hierachical model. Differencing these 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.

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":

$$\begin{aligned}\Pr(\ln W_{uc+1} < w | \eta > A_{uhc+1}) &= p_{uc+1}\Pr(\ln W_{uc+1} < w | \eta > A_{uhc}) \\ &+ (1 - p_{uc+1})\Pr(\ln W_{uc+1} < w | A_{uhc} \geq \eta > A_{uhc+1}), \forall w \qquad (12)\end{aligned}$$

where, p_{uc+1} is the proportion of the university educated in cohort c+1 who are stayers. Equation (12) 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,

$$\tilde{w}_{uc+1} = \beta_{c+1u} + f_{c+1u}(age_{it+1}) + \lambda_u med(\eta_i + \epsilon_{ic+1t+1}|\eta_i > A_{uhc}) \tag{13}$$

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,

$$\tilde{w}_{uc+1} - med(\ln W_{uct}|\eta_i > A_{uhc}) = \beta_{c+1u} + f_{c+1u}(age^*) - \beta_{cu} - f_{cu}(age^*) \tag{14}$$

That is, by comparing wage movements for people with the same set of $\eta's$ (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 (12), 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 (12) ($\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1}|\eta > A_{uhc})$) 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,

$$\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1}|\eta > A_{uhc+1}) = \frac{1}{2}p_{uc+1} + (1 - p_{uc+1}), \tag{15}$$

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,

$$\Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | A_{uhc} \geq \eta > A_{uhc+1}) \geq \Pr(\ln W_{uc+1} < \tilde{w}_{uc+1} | \eta > A_{uhc}) \tag{16}$$

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 (12) implies that the right hand side of (12) just equals 0.5. That is, the other bound is the c+1 median itself.

Meanwhile, 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, 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 2.

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 resulting quality-adjusted wage differentials are reported in the left panel of Figure 20. The second panel 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,

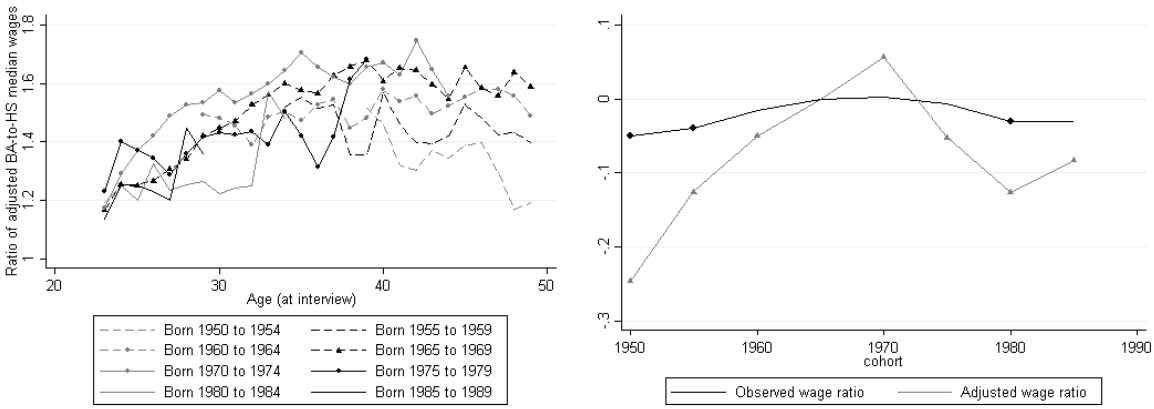

Figure 20: UK Median BA-to-HS wage ratio, adjusted to the education split of 1965 cohort

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.

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.