



Department
for Business
Innovation & Skills

BIS RESEARCH PAPER NO. 112

**THE IMPACT OF UNIVERSITY DEGREES
ON THE LIFECYCLE OF EARNINGS:
SOME FURTHER ANALYSIS**

AUGUST 2013

RESEARCH

THE IMPACT OF UNIVERSITY DEGREES ON THE LIFECYCLE OF EARNINGS: SOME FURTHER ANALYSIS

Report prepared by Ian Walker and Yu Zhu on behalf of the Department for Business, Innovation and Skills

Authors: Professor Ian Walker (University of Lancaster) and Yu Zhu (University of Kent)

The views expressed in this report are that of the authors and not necessarily those of the Department for Business, Innovation and Skills or any other Government Department.

Department for Business, Innovation and Skills
1 Victoria Street
London, SW1H 0ET

BIS Research paper number 112

August 2013

Contents

Contents	3
Executive summary	5
Key Findings	6
1. Introduction	8
2. Methodology	11
2.1 Gross earnings for employees over the lifecycle.....	11
2.2 Cohort effects	13
2.3 Employment, Non-Employment, and Self-Employment	16
2.4 Treatment and Counterfactuals.....	17
2.5 Estimation.....	18
2.6 Simulating Net Incomes and Government Revenue	19
2.7 Dropouts.....	20
2.8 Subject of study.....	21
2.9 Summary.....	21
3. Literature	23
3.1 Subject of study.....	24
3.2 HEI effects.....	24
3.3 Postgraduate qualifications	25
3.4 Degree class effects.....	25
3.5 Part-time study	25
3.6 Study later in life.....	26
3.7 Summary.....	26
4. Data	27

4.1	Datasets	27
4.2	Summary statistics	28
4.3	Summary	35
5.	Statistical Analysis	36
5.1	LFS estimates for earnings and their robustness.....	36
5.2	LFS robustness	38
5.3	BHPS estimates and robustness.....	43
5.4	Two step method.....	46
5.5	Employment probability	48
6.	Simulated Net Present Values	50
6.1	Method	50
6.2	NPV results	51
6.3	Sensitivity Analysis.....	55
6.4	Summary.....	59
7.	Conclusions and Further Work	60
	Appendix – Further Results	62
	References.....	65

Executive summary

This report, commissioned by the Department for Business, Innovation and Skills (BIS), provides further evidence on the impact of higher education on lifetime *net* earnings. There are two important components to the research: the statistical modelling of the relationship between higher education and *gross* earnings (and employment) that is used to *predict* what non-graduate earnings (and employment) would have been if only they had been graduates; and then simulating the way in which taxes and the student loan scheme affect these predicted gross earnings to allow us to infer the relationship between higher education and *net* earnings.

While earlier BIS research has focussed on the *marginal* effect of qualifications our focus is to compare the earnings (and employment) of those individuals who have a first degree (and 2+ A-levels) with those with 2+ A-levels and no higher education degree – *irrespective of the subsequent qualifications of both sets of individuals*. We think of this as capturing the decision-making of young people thinking about the prospective futures.

While previous research has been based on the large Labour Force Survey (LFS) datasets, here we also exploit the British Household Panel Survey (BHPS) data - both in isolation and in conjunction with LFS. LFS has information on degree class and degree subject – the “major”. BHPS contains information on HEI type (Russell Group, etc), slightly better data on A-Level qualifications and some family educational background information. Neither dataset contains both sets of information, so exploiting both is helpful. Our estimates suggest larger impact of lifetime gross earnings than previous research has.

We simulate the predicted earnings (and employment status) of individuals in our data and then average these to show that the private benefit of a degree, in terms of lifetime earnings net of tax and loan repayments, is large - in the order of £168k (£252k) for men (women) on average. The social benefit to the government is also large (of the order of £264k (£318k) from men (women) graduates – far in excess of likely exchequer costs. These estimates are larger than previous ones: partly because of innovations in our estimation method which is less restrictive than previous studies; and partly because of innovations in the simulation method which allows for the variance around our central predictions.

Key Findings

- Our estimate of the return to a degree relative to 2+ A-levels but no degree, defined by the coefficient on this variable in a regression model of earnings, is **23% for men and 31% for women**. These are comparable to the findings of other research.
- We estimate that there are **very substantial effects of a degree on the net present value of the lifecycle of incomes**: our best estimates of the likely impact on discounted lifecycle net earnings of having a degree, relative to not having a degree, is **28% for men** (approximately £168k) and **53% for women** (approximately £252k) on average.¹ While our estimates of the average effect of a degree on earnings and the probability of employment suggest that previous estimates are robust to a number of criticisms, previous work fails to capture an important difference between graduates and non-graduates in the trend growth in real wages. It is this that counts for much of the difference between our estimates relative to previous research.
- **Our estimates for a good degree (first or upper second) are significantly larger than for lower degree classes** (by £76k for men and £85k for women, on average) suggesting a large return to student effort. Allowing for the effects to vary with family background suggest no significant differences.
- We attempt to identify dropouts by comparing age left education with age associated with highest qualification and find that **male HE dropouts earn approximately the same as individuals who never attended HE. For females we find that there is a small wage penalty – dropping out from HE is worse than never attending HE.** However, controlling for dropouts makes little difference to our conclusion about the graduate premium.²
- We find estimates of the effect broken down by degree subject that support earlier research - but we are not, with available data, able to test the robustness of these estimates. This is a matter of concern since there are grounds for thinking that some of the differences across subjects are a reflection of differences in the students rather than differences in curriculum – for example, the A-level admission requirements might be different across subjects. Research elsewhere has also been uninformative on this issue.

¹ These figures are simulations of lifecycles of earnings from a statistical model of gross earnings and are adjusted for tax and National Insurance liabilities, periods of non-participation, and for the effect of the loan scheme, and discounted at 3.5%.

² Our control group has 2+ A-levels. Relatively few such students go on to Further Education (FE). So we do not think that our definition of dropouts will be contaminated by FE dropouts.

- We extend earlier research to consider the effect of degree by type of higher education institution. While we do find significant differences in a simple specification, we find that these differences are smaller and become statistically insignificant when we include some basic family background controls that capture at least some of the differences across students. Thus, **our research does not suggest that there are large differences in returns across broad types of HEI**, controlling for background.
- One focus for our work is to untangle the effects of cohort (of birth) from age and calendar time. There are two reasons we are interested in doing this. First, over time, there has been a growth in the gap between graduate and non-graduate wages in many countries. Second, we are also interested in exploring cohort differences associated with the rapid and large expansion of HE in the late 1980s and early 1990s. We are particularly interested in the extent to which the expansion of HE that occurred then might have reduced the impact of a degree on earnings. We compare the earnings of young graduates and non-graduates pre and post the HE expansion but **find no significant differences in the graduate earnings differentials associated with the expansion of HE**.
- We extend earlier research on the net present value of the net impacts of a degree in a number of ways that turn out to be economically significant. In particular, **we simulate the impact of a degree on the whole distribution of earnings rather than just simulate the average effect**. This allows us to capture the extremes of the data - graduates who do better than average pay more higher rate tax than would otherwise be thought; and graduates who do worse than average repay less of their student loans than would otherwise be thought. Simulating the effect on the average predictions for graduate and non-graduate earnings distribution tends to underestimate the net present value associated with HE relative to what we find when we simulate each individual and then average these simulated effects.
- We find that the new higher tuition fees make a predictable difference to the present value relative to previous fee levels. However, **we also find that the recent changes to the loan system insulates students who come from low income households from much of the effect of higher tuition fees**.
- The research also suggests that the net present value of the additional tax payments made by graduates relative to non-graduates is much larger than earlier research had suggested. **HE is an important investment for the government as well as for students**.

1. Introduction

This report, commissioned by the Department for Business, Innovation and Skills (BIS), provides further evidence on the impact of higher education on lifetime *net* earnings. It attempts to extend, verify, and refine the work conducted in 2011 by Gavan Conlon and Pietro Patrignani of London Economics for BIS that was published as “The Returns to Higher Education Qualifications”, BIS Research Paper 45 (hereafter referred to as LE). One aspect of the extensions to earlier work is to investigate how higher education impacts across parental background, and how type of higher education attended affects lifetime earnings.

Like LE, we consider both the relationship between higher education and earnings, and with employment. There are two important components to the research: the statistical modelling of the relationship between higher education and *gross* earnings (and employment) that is used to *predict* what non-graduate earnings (and employment) would have been if only they had been graduates; and then simulating the way in which taxes and the student loan scheme affect these predicted gross earnings to allow us to infer the relationship between higher education and *net* earnings. LE’s evidence, while comprehensive, begged a number of questions. While LE provides a very detailed breakdown of the aggregate results by subject studied, it left unanswered questions, such as: the effect of higher education institution (HEI), the effect of obtaining a degree through part-time study, and the effect of obtaining a degree later in life as opposed to immediately after leaving schooling at the age of 18.

While we provide further evidence to LE in some of these areas, where data are available, we are also (deliberately) more restrictive in other ways. In the main, this reflects a philosophical difference between this work and LE. LE focusses on the *marginal* effect of qualifications. That is, they compare the earnings of individuals with a particular vector of qualifications with the earnings of individuals with a different vector. For example, they compare the earnings of individuals with 2+ A-levels and a first degree (but not subsequent higher education qualifications) with the earnings of individuals with 2+ A-levels but with no further qualifications above A-level, to yield the marginal effect of a first degree. And, they compare the earnings of individuals of those with a first degree and a Masters degree (but no further qualifications) with those with a first (Bachelor) degree alone, to yield the marginal effect of a Masters degree. And so on. In contrast, our focus is to compare the earnings (and employment) of those individuals who have a first degree³ (and 2+ A-levels⁴) with those with 2+ A-levels and no higher education degree – *irrespective of the subsequent qualifications of both sets of individuals*. We think of this as capturing the decision-making of young people thinking about the prospective futures. At the age of 18 (or later), entering higher education offers the option of further HE qualifications beyond a first degree. And, not entering higher education offers the *option* of further study, perhaps undertaken part-time and/or on-the-job, that might yield additional

³ Unlike LE we do not consider higher education qualifications that are at sub-degree level (they are in our “2+ A-level “control” group rather than in the “treatment” group) but the proportions are small so the effect on our results is insignificant. We treat Foundation Degrees the same as first degrees. They account for a very small proportion of graduates and excluding them, or treating them separately, make no effective difference to our analysis and findings.

⁴ Similar results are obtained when we consider degree holders without 2+ A-levels.

qualifications – which may be professional, vocational or technical. That is, our philosophy is to capture the option values associated with the decision made to attend higher education or not.

How the work reported here differs from LE is explained in more detail in Section 2. This section explains the methodology and, in particular, what is novel in what we do relative to previous research. Section 3 briefly reviews the finding of existing work – and more detail of UK research can be found in LE. Rather than repeat the review in LE, we focus on literature that relates to the extensions to the LE work that we try to implement. We also highlight some of the caveats that apply to research of this type and explain how one might examine the robustness of the results to these caveats. This section considers not only the statistical modelling but also the simulation analysis.


Section 4 explains the selection of our datasets, and describes the characteristics of the sample used in our subsequent analysis. While previous UK research has been based on the large Labour Force Survey (LFS) datasets, here we also exploit the British Household Panel Survey (BHPS) data - both in isolation and in conjunction with LFS. LFS has information on degree subject – the major. BHPS contains information on HEI type (Russell Group, etc). But neither dataset contains both sets of information. So exploiting both is helpful. Section 5 presents results for earnings conditional on being employed and for the probability of employment. In Section 6 we exploit these results to provide the simulated lifecycles and the implications for individuals and the exchequer. Finally, Section 7 concludes with recommendations for further research.

In anticipation of our detailed results we find broadly similar results to LE when we confine ourselves to broadly similar specifications. In particular, the estimate of *average* effect of a degree on *gross* hourly earnings seems only slightly lower than that found by LE. Our results, for a specification that is similar to LE, suggests that a degree adds about 22% (33%) to male (female) hourly earnings (compared to LE estimates of 24% (31%)). To get a feel for what this implies in pure financial terms, note that 22% over and above the average annual male earnings for those with 2+ A-levels but no degree of approximately £30,000 p.a. is, over a working lifetime of 40 years, a little over one-quarter of a million pounds at current prices not discounted. Similarly for women, 33% of an average of approximately £21,000 pounds p.a. for 40 years is also a little over one-quarter of a million pounds. Of course, there are costs: tuition fees, additional subsistence, and the lost earnings during higher education. Moreover, the progressive income tax system will reduce this large apparent return. So, too, will discounting future incomes to make them comparable with the current costs incurred while studying. On the other hand, allowing for employment effects of higher education will make this return larger because, on average, graduates experience less unemployment and non-employment than non-graduates. All of these issues are factored into our more detailed work below.

When we apply the recent loan scheme and £9000 tuition fee we find, using some simple estimates of the gross earnings and employment differentials, that the discounted net present value of the differentials in expected lifetime net earnings are, on average, those shown in Figure 1. This figure compares our preferred (WZ) estimates with those in LE.

While many of the differences in specification that we explore make little difference to our (WZ) results relative to LE's work, there are some differences that are very important. Some of these differences make us believe that the *net* effects of higher education, allowing for all the relevant factors, are likely to be *substantially* larger than LE's calculations suggest and this is reflected in Figure 1.

Figure 1 Headline estimates of average discounted lifetime NPVs of graduate earnings differentials (£k): WZ vs LE



While we have searched for evidence that recent cohorts of graduates enjoy a smaller earnings premium over non-graduates, we have not found any statistically significant differences. Our suspicion is that a smaller graduate earnings premium *may* have been driven by the rapid and large expansion in the supply of graduates that occurred in the early 1990s. If this exceeded the expansion in the demand for graduates then we would expect that increasing competition for graduate jobs would lead to a fall in the graduate earnings premium. If this were the case then it might reasonably be regarded as a long term effect – until rising demand caught up with the higher HE participation rate. But a second possibility is that a lower earnings premium for recent graduates is a cyclical phenomenon arising from them entering a labour market that was suffering from recession. *If* this were the correct explanation then we might reasonably expect any low earnings premium for recent graduates to be a temporary phenomenon. Unfortunately, we do not yet have sufficient data on the earnings history of young graduates who graduated in the early 1990s and beyond to be able to say whether what we might observe is because HE participation had increased dramatically or because the recession resulted in unemployment rising from 6.9% to 10.7% between 1990 and 1993 which depressed the market for new graduates by more (or less) than other workers. This will only be resolved in time as we get more earnings history for the post-expansion graduates in the post-recession period.

Finally, we have the advantage over LE in that we can now provide simulations of the effects of new tuition, scholarship and loan arrangements because we now know what has been implemented. Here, we show, amongst other things, that studentships and the changes to the loan scheme provide considerable insulation, for those students from low income backgrounds, from the reductions in net returns associated with the large increase in fees.

2. Methodology

This section considers a wide variety of issues that arise in this literature. Prominent amongst these are: identifying lifecycle effects separately from time and cohort effects; selection into work; selection into higher education; and selection into subject of study.

2.1 Gross earnings for employees over the lifecycle

There is a long history of research that attempts to understand the distribution of earnings as the result of the investment decisions that individuals make in their human capital. The essence of the theory of human capital is that it is formed by investing – in education. If the capital market works efficiently then the return on an investment in a degree should be the same as the return on any other similar financial investment – call this rate r . If the proportionate gain in earnings associated with an investment in a degree is $(w_1 - w_0)/w_0$, where w_1 is the earnings⁵ for an individual with a degree ($D=1$) and w_0 is for the individual without ($D=0$), then this should differ from r only by some random amount e whose value on average would be 0. The proportional difference in earnings is (approximately) the difference in the log of earnings. Thus, the theory suggests that $\log w_1 - \log w_0 = rD + e$. And if $\log w_0$ depends linearly on a vector \mathbf{X} of other characteristics like sex, region, etc., so that $\log w_0 = \alpha + \mathbf{X}'\beta$, then we can write the log of earnings for any individual as

$$\log w = \alpha + \mathbf{X}'\beta + rD + e$$

where β is a vector of parameters associated with the variables \mathbf{X} , such that $\mathbf{X}'\beta$ is the sum of each of the products of each of the X 's and their corresponding β 's. Thus, α is the log wage of a default individual (with sex=0, region=0, etc) without a degree and $\alpha + r$ is the log wage of the default individual with a degree. That is, the earnings differential associated with a degree is (approximately) r . And this is the same for all individuals regardless of X because the effects difference out when we subtract the expected non-graduate earnings from the expected graduate earnings.⁶ The use of log earnings as the dependent variable is suggested by the theory and has the useful implication that the coefficients on the explanatory variables can be interpreted as percentage effects: thus, if the coefficient on D were 0.10 then we would infer that earnings for observations with $D=1$ were, on average, 10% higher than those for $D=0$. Similarly if the coefficient on sex (being male, say) were 0.2 then we would infer that, other things being equal, men would be 20% better paid than women on average.

This very simple (log) linear specification explains the variation in log earnings across individuals by the differences in their X 's, whether they have a degree or not, i.e. D , and a random component, e , that captures the unobserved factors that influence w_0 . The X vector is usually specified to include a measure of work experience to capture the accumulation of human capital that occurs on-the-job through training and learning by doing. This is often incorporated as a quadratic function of age. Age is a proxy for

⁵ We use hourly earnings throughout our empirical analysis.

⁶ Strictly speaking $r = \log w_1 - \log w_0 \approx (w_1 - w_0)/w_0$ and the approximation is fairly accurate for values of r below 0.1. Here we have larger estimates and we correct for the error in the approximation when we report the results.

experience (which is often not observed in many datasets) and the quadratic relationship derives from a presumption that learning on the job is subject to diminishing returns - the older one gets the less easily one learns. It is this (assumed) quadratic shape that drives the lifecycle of earnings.


Such an equation has become the workhorse model for empirical research in the determination of earnings. There are many estimates of such models and it is this basic framework that is used in the work by LE. However, the derivation of this relationship depends on a variety of assumptions made for convenience. The theory has strong implications that may not be borne out in practice. For example, this specification assumes that the effect of having a degree is the same for all values of X . In particular, the model implies that age earnings profiles are parallel in D - having a degree simply shifts the (log) earnings quadratic pattern vertically upwards. That is, the simple specification assumes that log earnings for graduates and non-graduates conforms to the left panel in Figure 2 – that age - log earnings profiles are parallel. This is a strong restriction and one that is usually rejected by the data when a model such as that in the right panel of Figure 2 is estimated. Indeed, it is common to find that log earnings differentials between graduates and non-graduates are narrower at an early age, as shown here. Similarly, the left hand side of Figure 2 assumes that the impact of any of the X 's, for example sex, is the same for those with $D=1$ as those with $D=0$. This is a strong restriction but one that is seldom tested. It is easily rejected here in favour of a specification where we allow all coefficients on the X 's to differ with D . Thus, a preferable model is one where the wages of graduates are allowed to be determined differently than the wages of non-graduates: in other words, by the *two* equations written as

$$\log w^D = \alpha^D + \mathbf{X}'\beta^D + e^D, \text{ where } D = 0, 1.$$

Thus the log wage differential associated with having a degree is now $(\alpha^1 + \mathbf{X}'\beta^1) - (\alpha^0 + \mathbf{X}'\beta^0)$. Since $(\alpha^1 - \alpha^0) = r$ this can be rewritten as $r + \mathbf{X}'(\beta^1 - \beta^0)$ which makes it very clear that the log earnings differential associated with a degree is now allowed to vary with X if β^1 and β^0 differ. Thus, a test of the general specification against the restricted specification, as in LE, is that $\beta^1 = \beta^0$ for all the coefficients in the model. This has typically been strongly rejected in the literature.

Thus, our methodology is an important extension of earlier work. However, it has the disadvantage that it does not yield a single *number* to summarise the effects of a degree. Rather there will be an estimate for each cell of the \mathbf{X} matrix.⁷ However, to arrive at an aggregate figure we only have to take the weighted average for each cell.

⁷ In practice, the number of cells is quite small – here, for each sex, we group the data into six regions, two ethnic groups, and two immigrant statuses. So there are 24 types of individuals, of each sex, in our data. We know the proportions of each type in the data so we can weight our estimates for each type by the proportion of each type to derive a single overall estimate of the impact of a degree, for each sex.


Figure 2 Possible specifications for the effect of HE on log earnings

2.2 Cohort effects

There have been considerable changes in the Higher Education (HE) system in the UK over the last 40 years. We have moved from an HE system that was provided for a small elite, free of tuition fees with mean-tested grants for subsistence, and so was highly subsidised. We have moved to a system where a large minority of any cohort attends HE, there are much smaller direct tuition subsidies and an extensive and sophisticated system of income contingent loans which are heavily subsidised, and grants in the form of income contingent bursaries and (academic) merit based scholarships. Differences across devolved governments have also emerged. Moreover, there have been large changes in the supply of graduates to the labour market - Figure 3 shows the dramatic increase in HE participation rates that occurred from the late 1980s until the mid 1990s. It seems likely that this dramatic growth has affected the distribution of wages – post expansion cohorts may well experience smaller degree differentials than pre-expansion.

A major concern for us is to capture the lifecycle variation in log earnings. However, we have a strong reason for suspecting that age affects earnings not just through work experience but also through year of birth. For example, this large expansion in the number of graduates in the labour market that began in the early 1990s may well have driven down the graduate earnings premium. This is a cohort effect: the earnings that you receive at any age depend independently on whether you graduated into a labour market where graduates were scarce (before the expansion) or when they were plentiful (afterwards).

Distinguishing between lifecycle, cohort, and time effects is important for our purposes. There is productivity growth over time and the literature has strongly suggested that this productivity growth is “skill biased” – that is, it is higher for skilled workers (graduates) than unskilled (non-graduates). There are changes in HE participation across cohorts that give rise to cohort effects in earnings that differ between graduates and non-graduates. Finally, there is the pure lifecycle growth in earnings that we need to account for in order to construct the lifecycle of earnings from earnings at a single point (or two points) in the lifecycle. All of these are related to each other.

Figure 3 HE Participation in the UK, 1960-2001⁸

If we were using just cross section data taken in a particular year then the age of each individual would be the survey year minus the year of birth. So year of birth and age would be perfectly collinear and it would not be possible to identify the age earnings profile except by assuming that there were no independent cohort effects driven by year of birth. This would be a very strong assumption. Repeated cross sections allow us to separate out any TWO of these effects assuming the third is zero. The LFS data provide cross sections from 1993 to 2010 which can be pooled together (adjusting wages for inflation) but even this would present us with a collinearity problem to an extent. As an alternative to this we can exploit the fact that LFS is a short panel – wages are recorded in waves 1 and 5, a gap of one year. Thus, we could exploit the wage growth recorded in the panel data to estimate the age earnings profile for workers.

There is a selection issue here – individuals could only be included in this analysis if they record earnings in both waves 1 and 5, while the alternative would only require that we

⁸ There have been two official measures: the Age Participation Index (API) for all Great Britain up to 2000, and the Higher Education Initial Participation Rate (HEIPR) for English domiciled students from 1999 onwards. The API is the percentage of each cohort currently undertaking higher education. The position was broadly stable over the 1970s and 1980s but increased quickly from about 15% for men and 12% for women in 1988 to 30% for men and women in 1994, where the rates stabilised. The great majority of UK students who attend higher education do so soon after completing high school at the age of 18 or 19, and the great majority study full-time. This study is typically for a three year first degree (Bachelor) course (health, and some other, courses are typically longer). The main driver of the expansion of higher education in this period was the increase in the full-time participation of 18-21 year-olds. It is this growth, from approximately 15 per cent in 1988 to 30 per cent in 1993, that is measured by the official API. In addition, there was some growth of mature students and of entrants taking alternative routes to higher education (HE). The HE expansion began in 1988, corresponding to the cohort born around 1969, and the expansion had ceased for men around 1994 corresponding to the 1976 birth cohort. The API series was discontinued in 2000 and replaced by the Higher Education Initial Participation Rate (HEIPR), for England, which counts the proportion of young people (17-30) who have had at least 6 months HE experience. The series are not consistent with each other but it seems likely that their trends will be quite similar.

observe wages in either wave. On the other hand, the usual cross section method makes the implicit assumption that age (minus the age at which you ceased education) and experience are one and the same thing, so it ignores spells of non-participation that might occur. This is a useful approach because it exploits the LFS data to its full. The LFS is a panel of addresses not people. So researchers need to drop observations in the panel where the wave 1 household members are not present in wave 5 – for example, the whole family may have moved. Thus, the dataset used for the first step is a subset of the dataset used for the second step.

On balance we prefer exploiting panel data to estimate age-earnings profiles because the estimates derive from the effects of one year of work experience on earnings. That is, we think of the age variable (and its square) as capturing the effects of work experience on earnings. For men, age is a good proxy for (unobserved) work experience but for women, especially women from early cohorts, it is not because of interrupted spells of work associated with childcare. Thus, estimating the lifecycle of earnings from the cross section data we would be likely to underestimate the effect of experience from the effect of age, for such women. But, by confining attention to individuals who are observed in successive years in a panel we can attribute the earnings growth to the effect of experience between successive interviews.⁹ Thus, the first step of the two step method allows us to estimate the way in which earnings evolve over the lifecycle.

To see, more precisely, how this **two-step method** works assume, for simplicity, that \mathbf{X} contains only age and age squared. Our model of earnings determination would then be

$$\log w_{it}^D = \alpha^D + \eta^D \cdot \text{Age}_{it} + \gamma^D \cdot \text{Age}_{it}^2 + e_{it}^D$$

where the i subscript indicates individual i and e_{it}^D captures the effect of all unobservable factors that we cannot control for using our data - for example, ability and ambition. Then we can use the panel data to form a wage *growth* equation obtained by taking the equation above, that shows wages at time t , and subtracting from it the same equation at $t-1$ to obtain

$$\Delta \log w_{it}^D = \eta^D + 2 \gamma^D \cdot \text{Age}_{it} + \Delta e_{it}^D$$

where Δ is the time difference operator (that is, $\Delta \log w_{it}^D = \log w_{it}^D - \log w_{it-1}^D$), to allow us to estimate γ^D and η^D from a wage growth equation in a first step. Notice that this equation is unaffected by any cohort differences IF the only effect of cohorts was captured by α^D in the first equation. It is possible to estimate the lifecycle parameters, η^D and γ^D , from panel data. These parameters can then be fixed in the second stage, where the equation for the level of wages, $\log w_{it}^D$, is estimated. This fixes the lifecycle component of wage variation and allows the cohort variation to be captured independently.

The assumption underlying this is that cohort effects do not affect the *shape* of the age earnings profile.¹⁰ But they can independently affect the height of the profile. So, once we have estimates of γ^D and η^D from the panel data we can impose these estimates on to the cross section data and estimate the remaining parameters, including any cohort effects. Thus, if α^D , which determines the height of the age-earnings profile, varies across

⁹ LE report that they estimate their log hourly earnings equation in five year age intervals. In addition they also report that they include age and age squared. This is more general than our specification – which effectively collapses the five intervals into one. We tested whether including age interval dummy variables were jointly significant when age and age squared were included and found that they were not.

¹⁰ This is a maintained hypothesis in the analysis that would be difficult to test directly with these data.

individuals in a way which depends on one's cohort (i.e. if we allow $\alpha_i^D = \alpha^D(\text{cohort}_i)$) then we can write log earnings as

$$\log w_{it}^D - \hat{\eta}^D \text{Age}_{it} - \hat{\gamma}^D \text{Age}_i^2 = \alpha^D(\text{cohort}_i) + e_i^D$$

where a “hat” above a coefficient indicates that it has been estimated (in the first step). We could then estimate how α^D depends on i 's cohort (defined by year of birth) by applying linear regression to this equation using the cross section data in a second step. As in the first step, we can estimate separate equations for degree holders ($D=1$) and others ($D=0$) in this second step. We refer to this procedure as our two-step model: in the first step we use the panel to estimate the lifecycle parameters, γ^D and η^D , for each group, $D=0,1$; and in the second step we estimate the remaining parameters (on the other X 's) including the cohort and calendar time¹¹ effects – again for each group.

In practice, we divide the data into three groups: the oldest cohort was born before August 1970 and so reached the usual age to attend university on or before the 1989 intake which is prior to the expansion period; the second cohort were born between September 1970 and August 1976 and so will have typically started university during the expansion period; while the third cohort were born after August 1976 and so attended university after 1995 when the rapid expansion was coming to an end.

It seems reasonable to suppose that this large and rapid expansion in the supply of graduates to the labour market will have forced down the earnings of graduates. However, Walker and Zhu (2008) were not able to find statistically significant effects of the expansion of higher education. It remains to be seen whether the arrival of more post expansion data since that earlier research will allow us to make more precise estimates.

2.3 Employment, Non-Employment, and Self-Employment

Naïve economic theory suggests that wages in competitive labour markets are determined by productivity alone. In such circumstances, market forces equilibrate demand and supply for all types of worker such that unemployment is eliminated quickly. However, there are many reasons for thinking that labour markets do not work very efficiently and that unemployment occurs as an equilibrium phenomenon. Thus, there may be a case for thinking that education affects not just wages but also the probability of being in employment. Therefore, like LE, we model the probability of being in employment similarly to log wages. However, just as in our approach to modelling earnings conditional on employment, here too we allow the importance of each determinant to differ with D . Of course, what we observe is not the *probability* that an individual is employed or not, but whether she is *actually* employed or not. The conventional way of modelling the probability of employment is to assume that the probability of being employed, i.e. $Prob(E=1)$, is a function of a vector of observable determinants, \mathbf{Z} , and by \mathbf{D} (and unobserved factors, captured by u). That is,

$$Prob(E=1) = \varphi^D + \mathbf{Z}'\delta^D + u^D$$

¹¹ We allow for an effect of calendar time on earnings to reflect the role of productivity increases that occur over time. Indeed, we allow for there to be a different rate of productivity increase for graduates and non-graduates. This reflects the widespread view that technical progress has been biased towards skilled labour – innovations that have occurred over time have raised the demand (and wages) of skilled workers relative

where again $D=0,1$. That is, we estimate separate equations for degree holders and others. In practice we assume that \mathbf{Z} contains the same drivers as in \mathbf{X} . We again include age and age squared in \mathbf{Z} in the first step, and we allow for cohort effects in the second step. If one assumes that u follows a Normal distribution then this model is known as a Probit model. The Probit model has the desirable property that it predicts that the dependent variable is between 0 and 1 and so the prediction can be interpreted as a probability.

Our estimates of the effect of higher education on earnings are clearly conditional on observing earnings. We only observe the earnings for employees - not for the unemployed, or for those out of the labour market, *and* nor do we observe earnings for the self-employed. It is not possible to sign the bias induced by estimating conditional on selection into employee status – there are arguments that work in both directions.¹² The only consolation is that the proportion of the workforce made up by the self-employed is reasonably small – of the order of 15%, and the unemployed are less than 10% typically, while the out-of the labour force group was large for women but has been getting substantially smaller. The size of the bias from using employees is therefore probably fairly modest except, perhaps, for older women. However, further research on the self-employed is clearly merited if better data on their incomes can be found.

2.4 Treatment and Counterfactuals

Our perspective in this work is to think of the decision faced by an 18-year old who is qualified to attend university. On the one hand she might choose to study for a degree, graduate and join the labour market, or may then study for a further degree and then join the labour market. She may obtain further qualifications of a vocational nature – say, achieve *chartered* status in some profession. On the other hand she might reject the higher education option and enter the job market immediately - which may then be followed by career progression, which might include gaining further qualifications of a vocational nature. Thus, our focus is to compare the labour market careers of people who could have accessed a university degree course and did so, with those who could but did not - irrespective of what happens to them after they make this decision. In other words, our analysis includes the value associated with the options that arise *after* the age of 18 – both for those who take a degree course and those that do not. This is unlike the analysis conducted in LE which estimates the effects of each qualification separately and so provides estimates of the effects of any qualification relative to any other: for example, a PhD relative to an undergraduate degree. The work reported here does not control for each qualification separately. Rather we report on the subsequent earnings of those that choose the HE route with those that do not. That is it includes the “option value”

to unskilled workers. In contrast, LE simply assumes that there is a common and fixed rate of productivity rise.

¹² The issue is the sign of the correlation between selection (into having positive earnings, or into employment) and unobservable characteristics. If unobserved “ability” is positively correlated with education (as seems likely) and is also positively correlated with being employed (perhaps because prospective employers can observe some of these qualities that we, as researchers, cannot), then the coefficient on education will pick up both the true direct effect of education on being employed and the indirect positive effect through unobserved ability – that is the coefficient will be biased upward. However, the bias will be downward if the correlation between these qualities and being employed is negative - perhaps because high ability people find it hard to credibly demonstrate their true potential productivity to prospective employers and so choose to employ themselves – i.e. become self-employed

associated with each route: graduates have the option of further academic and vocational qualifications which have a value in the labour market; while those that do not participate in HE may join the labour market and have the option of pursuing vocational qualifications up to and beyond degree level, which also have a value. By choosing not to control for further qualifications in our analysis we attribute the value of these options to the choices made at 18.¹³

Here, the *treatment* that we consider is entering higher education and obtaining at least a first degree¹⁴; while the non-treated state is not obtaining a first degree.¹⁵ Thus an appropriate control group would be similar individuals who did not enter higher education and obtain a degree. We follow the existing practice in the UK of using those who have 2+ A-levels as the control group; although, because we want to include the option value of further qualifications, we include in this those individuals with further qualifications but not obtained through higher education.¹⁶

2.5 Estimation

What should be included in Z and X is an open question. Here we are more parsimonious than the work of LE. We include only factors that we think are unlikely to be driven by having a degree. Thus age and age squared and ethnicity are all exogenous and are not affected by education (or anything else). We assume that immigrant status (arrived in the UK before age 5)¹⁷ is exogenous – because it was determined by one's parents. We assume that region of residence is also exogenous - although there may be some grounds for suspecting that having a degree leads you to be more likely to locate in London. Our reason for parsimony is that we do not wish to include variables that pick up the effect of having a degree – this would lead to estimated degree impacts to be lower than would otherwise be the case. For example, graduates are more likely to be married and there is some evidence that married men, less clearly so for women, earn more.

¹³ An individual may choose to leave education at 18 or before, work, and then return to education to obtain a degree later in life - perhaps after obtaining some qualifications through a spell of Further Education. Such individuals are in our treatment group but we do consider whether their earnings are the same as those that pursue Higher Education earlier in the lifecycle.

¹⁴ LFS categorises those with a Foundation Degree as having a degree. Foundation degrees were introduced from 2001. From 2004 the LFS records this separately. In the 2004 data and later the proportion getting a Foundation Degree is just 3%. The small size of this group implies that it will make little difference how we categorise them. We choose not to differentiate between these types of degree in our analysis – that is, we include those with Foundation Degrees in the treatment group. The alternative is to drop all of the data prior to 2004 which would have reduced the precision of all of our estimates. There is a small number of higher degree holders who do not record a first degree. We include them in our analysis and treat them as having a first degree (of the same subject as their higher degree). Excluding them would have had no discernible effect on any of our findings.

A HE diploma is often an entry qualification available at some HEIs and is often part-time study. This is a small group and we have chosen to include them in the control group if they have 2+ A-levels, and are otherwise dropped.

¹⁵ We discuss drop-outs (those that enter HE but do not obtain a degree) below.

¹⁶ For example, according to our definition, 11% of teaching professionals are non-graduates. This group are disproportionately female and are older than their graduate teacher counterparts.

¹⁷ We drop observations who arrived in the UK after age 5 on the grounds that their education may have started prior to arrival. Even if they have 2+ A-levels they will have important unobservable differences from non-immigrants – like language skills.

There is a worry that having a degree is correlated with *unobserved* determinants of earnings. That is, our estimate of the effects of a degree may be biased because having a degree is correlated with unobserved variables that also affect wages. For example, higher ability people are more likely to have a degree, and higher ability people earn more on average, irrespective of their degree status. Thus, we might be attributing the effects of unobserved ability onto the estimated effects of having a degree, and this biases the degree effect upwards. Thus, it might be reasonable to assume that simple regression estimates provide an upper bound to the true effect. We explore this important issue later in the literature review.

2.6 Simulating Net Incomes and Government Revenue

In addition to estimating the effect of a degree on earnings conditional on being in employment, our work also provides for an effect of a degree on the probability of being in employment. Together, this allows us to compute the expected earnings across the lifecycle. From these estimated lifecycles of earnings we can compute incomes net of income taxation (and NI contributions), we can also incorporate the expected loan repayments by applying the Student Loans Company (SLC) rules, and we can impute the additional VAT revenue associated with higher net incomes (from estimates of the expenditure patterns of households). This yields the net benefits (to the individual) in terms of consumption expenditure from following the degree lifecycle compared to the non-degree lifecycle.

There is a long history, pioneered by the Institute for Fiscal Studies, of computing net incomes from a given set of individual characteristics including gross incomes. The calculations required to do this are simple but tedious. It was commonly the case that such calculations were made just for an average individual, or for “representative” individuals. However, simulating the effects for an average individual will not, in general, yield the same answer as simulating the effects for every individual and then averaging. Thus, unlike LE who simulate an average individual, we simulate the net income consequences for every individual in our data and then average.

Indeed, we recognise that our estimates are just that – statistical estimates. The predictions from our estimates have an estimated variance, which we know from our estimation. We could take the point prediction from our estimates to generate simulated lifecycles. However, this would ignore the variance around our point prediction. This variance is large because the distribution of earnings is affected by many things, including unobserved factors such as luck and ability, and the small number of variables that are observable in our data account for a relatively small proportion of the variance in earnings (our results are typical in explaining only around 30% of the variance in wage rates). Failure to account for this variance would imply too few *extreme* observations because we would fail to account adequately for the actual variation in the data. Even if we are only concerned about aggregate effects we still need to take heed of the estimated variance in our gross earnings predictions because net earnings depend on gross in a nonlinear way. In particular, once earnings exceed some level a higher rate of income tax becomes payable. Another threshold affects the repayment of student loans. Thus, it is important that our methodology captures the uncertainty in our estimated predictions if we are to correctly estimate the aggregate effects.

Since our analysis forecasts the whole lifecycle of earnings for each individual, with and without a degree, we can compute loan repayments at each year of the lifecycle. Thus, we

allow for less than complete repayment of loans in our calculations without any *ad hoc* adjustment.

When computing earnings we take the average hourly wage rate and multiply by 2080 for men and 1820 for women to reflect the average number of hours worked per annum recorded in the LFS. A more sophisticated analysis might analyse the relationship between labour supply and education. For example, here we distinguish only between working (and being an employee) and not. A step beyond this would be to consider part-time and full-time employment. However, part-time work is usually worked during only part of the lifecycle and, while we can observe this in the data, in a model where individuals plan their future it would be the case that the lifecycle of labour supply would be jointly determined – for example, an individual who has a low taste for work would choose a low level of education and face low wages throughout the lifecycle. In effect, labour supply is the utilisation rate of human capital which has been generated by education and they are determined jointly. In the absence of a formal framework within which labour supply decisions can be modelled, jointly with education, we need to assume something about labour supply. We could assume that labour supply is fixed at the observed level; or that it is always full-time; or that it is the average that we observe across all individuals. For the moment we have chosen the latter option but this would be a useful extension if there were data that would support it.

An associated difficulty is what to assume about income when individuals are not employed. One option would be to assume that those not employed are unemployed and claiming Job Seeker's Allowance (JSA). Of course, JSA is only available for 6 months before means testing against household income and this would require that we estimate the determinants of partnership formation. Once means testing occurs, the transfers from the government will depend on housing costs and children as well as partner's income. The prospects for being able to construct a model that would admit this level of complexity with available data is very slim. Here we chose the simpler option - we assume that income when unemployed or out of the labour force is zero.¹⁸ Since we also drop the self-employed we are ignoring any of the effects of education on this group. Having computed net income from gross, government revenue is straightforward to compute as the difference between gross and net income.

2.7 Dropouts

Dropouts represent an important problem for our work, and that of LE. The LFS data do not tell us directly who entered higher education but failed to successfully complete.¹⁹ Thus, we cannot perfectly identify dropouts in the data. This could be a serious problem because we cannot then know how much dropouts earn and because they contaminate the control group in our data which consists of those 2+ A-levels who did not obtain a degree. In other words, we include in the control group those who entered HE but did not successfully graduate. In some institutions dropping out (after two years) can be

¹⁸ We explore the robustness of our results to allowing for JSA in Figure 17. There is a small difference for women who are much more likely to be out of employment than men.

¹⁹ HESA http://www.hesa.ac.uk/index.php?option=com_content&task=view&id=2064&Itemid=141 suggests that the proportion of 2006/7 entrants who get a degree is 77%. For post 2002 LFS data we find that approximately 22% of the control group have some HE experience but no degree. This constitutes 7% of the treatment group and this is only capturing about one third of dropouts. We examine the implications for lifecycle earnings later.

certificated in the form of a HE diploma in which case they could be identified.²⁰ But in many cases this is not applicable.

Our approach to this problem is to try to pin down dropping out by comparing age at which education ended and age completed highest qualification with the vector of qualifications recorded. Having A-levels but no degree with age of completed education of, say, 20 might be indicative of dropping out and we provide some estimates later that attempt to do this.²¹

2.8 Subject of study

LE provides extensive evidence of the extent to which the graduate earnings (and employment) differential relative to non-graduates varies across subject of study. Similar findings are reported in Walker and Zhu (2011). LFS identifies subject of study (of first degree, but not higher degree) by broad group so this is straightforward to do. There are two main difficulties in interpreting the estimates as subject “effects” – that is, differences induced by the differences in the curriculum. First, there is *differential* ability bias. Admission into medicine, law and other subjects is very competitive and graduates in these subjects are likely to be of higher average ability than in other subjects. One might argue that part of the differential that graduates in such subjects enjoy reflects their higher ability. On the other hand, one might argue that graduates in such subjects would have earned more than the average of non-graduates or the average of other graduates. So it is not clear which direction the bias would go in: if medical graduates would have made excellent butchers then the bias may even be downwards. Secondly, there is differential selection into *employment* across graduates from different subjects. The majority of medicine graduates that go into General Practice will become self-employed, and many hospital consultants will have large incomes from private practice. Similarly, most veterinary and dental graduates become self-employed, as well as many of the more successful law graduates. Thus, this censoring of the data is likely to bias the estimates of the returns to studying such subjects downwards.²²

2.9 Summary

This section has outlined the methodological issues that arise in considering how to model the impact of higher education on earnings. We have concentrated on those issues that seem most relevant to UK policy and draw attention to the differences between our approach and previous research.

²⁰ LE report (their Table 7) that males who are certificated with a Diploma have earnings that are not significantly higher than the rest of the control group who never entered HE. However, they also report that females who are so certificated have earnings that are not significantly different from the treatment group of those that graduate with a degree.

²¹ There are drop-outs that we would not capture with our methodology based on age completed *continuous* education. For example, individuals that have a break in education of *more than* one year after A-level and then dropped out of a higher level qualification would not be identified as a drop-out and would be included in the control group. We are not able to say what proportion of people would fall into this category. It is also possible that our method would capture some FE dropouts.

²² The LFS data record self-employment and while the overall rate for all male graduates is 12.5%, for medical and dental graduates it is 41%, for lawyers it is 34%, for architects it is 19%, and for Art and Design graduates it is 26%. The rates for women who have studied these subjects are also large relative to the average graduate self-employment rate for women of 7.5%. In simulations we drop medical and dental graduates – approximately 1.5% of all graduates.

We have been moved to take a more flexible approach to modelling than previous work. In particular, we relax the restriction that the effect of a degree is just to raise earnings across the board. We allow the degree differential to vary with the characteristics of individuals – in particular their age. That is, we allow the age-earnings profile to differ between graduates and non-graduates. Moreover, the expansion of the HE system in the late 1980s and early 1990s means that it is imperative for us to try to identify how the impact of HE may have changed across cohorts of students. Post expansion cohorts graduated into a more crowded labour market and we might expect the earnings premium to be lower for such cohorts. But one's cohort is not independent of age and calendar time so we adopt a two-step methodology that allows us to exploit the panel data that we have to provide an estimate of the age– earnings relationship independently of cohort. This then allows us to estimate the effect of cohort conditional on panel estimates of the age-earnings profile.

We also considered how to handle drop-outs in the data. We can use information on year at which education ends and compare that with age at which highest qualification was attained to give a sense of who, amongst the control group of 2+ A-levels but no degree might have continued education beyond 18 but not completed further qualifications. One reason why we might observe this is dropping-out.

Finally, we attempt to decompose our effects of degree by subject studied. While this is straightforward to do, the problem this raises is that there are unobserved ability differences across subjects. We are not able to separate out whether the impact of degrees of different subjects differ because of the curriculum differences or because of unobserved differences in ability that differs across subjects. This is not an easy problem to resolve even if we had better data so, for the moment, we recommend that little weight be given to subject specific returns. This problem is exacerbated with different rates of self-employment across subjects. The self-employed do not report earnings in LFS and are dropped from our analysis. Ideally, we would like to include them – not least because they are likely to have higher than average earnings so their omission could bias our estimates.

3. Literature

The workhorse model of earnings determination, and our generalisation of it which allows for cohort effects to be separated from age and calendar time effects, is conventionally estimated using linear regression methods.²³ This is true in LE, in our own earlier research (see Walker and Zhu (2008, 2011)), work by other researchers using LFS data²⁴, and here.

While there is a huge literature on the general issue of the impact of a degree on earnings²⁵, there is only a limited amount of previous research on the issues that we are most concerned with here. The distinctive issues that we are potentially concerned with are: differences in the impact of a degree across subjects studies (i.e. by “major”); the impact of post-graduate qualifications; the impact of degree quality (i.e. class of degree); the impact of HEI (or, at least, broad type of HEI); the impact of a degree obtained by part time study; the impact of a degree obtained later in life as opposed to straight after school; allowing for the effect of a degree to affect earnings differently across the lifecycle (i.e. not constraining the effect to be the same at all ages); and, finally, cohort effects.

One approach to dealing with the problem of differential unobserved ability between graduates and non-graduates is to uncover some exogenous variation in the probability of having a degree and use this to estimate the effect of a degree. It is unclear how one would do that with the available data.²⁶ A second methodology is provided by twins – who (arguably) are identical in terms of those unobservables that affect earnings. The twins researcher compares the earnings within (identical) twin pairs and so differences out the unobserved ability factor. Unfortunately, there is only one UK twins study (see Bonjour *et al* (2003)), which is based on a small sample of women. This work finds that the effects of (a year of) education is to raise wages by approximately 8% (which might, heroically, be extrapolated to approximately 25% for a three year degree). Their figure of 8% is identical

²³ While most analysis uses simple regression methods there is some research (see, for example, Walker and Zhu (2011)) that provides *quantile regression* (QR) estimates – regression estimates that are weighted to reflect different quantiles of the distribution of the error term. We also explore QR estimation and report simulation results that exploit these estimates.

²⁴ Most notably by O’Leary and Sloane (2004, 2005).

²⁵ Much of the literature considers education as a scalar measure – the number of years of education. Much of it is concerned with qualifications, usually just academic qualifications. A rather small proportion is concerned explicitly with just one qualification – a degree. However, there is a substantial literature on the “college” premium, predominantly in the US. There is no recent survey of the generic education literature but Harmon *et al* (2001) contains summaries for many European countries. Ashenfelter, Harmon and Oosterbeek (1999) provides a meta-analysis of a much wider literature, and Card (1999) provides a review of the issues which is illustrated by a number of key papers in the area. The most recent UK research, specifically on this topic, is by Walker and Zhu (2011) and this report builds on the methodology provided there.

²⁶ Existing estimates that use this method are suitable for identifying the effects of education at *lower* levels of education because the exogenous variation is provided by a policy change in the minimum school leaving age. See Harmon and Walker (1995). However, this identifies the effects of education on those individuals whose education is changed by the reform – in other words, for people who had wanted to leave school early. Thus, this literature is uninformative about the causal effects of *higher* education. While the reforms to the structure of fees may have affected the probability of getting a degree there are currently too few individuals in the LFS who are likely to have been affected by these changes and it is not practical to exploit this idea at this stage.

to what they obtain by applying least squares to the raw twins data, rather than looking at within-twin differences. This suggests that the extent of ability bias may be modest. A third methodology is to attempt to control for a rich set of observable characteristics. Blundell *et al* (2004) does this using the 1958 birth cohort in the NCDS data. This study also finds that least squares provides a relatively tight bound on the earnings effect of a degree. Thus, there are grounds, based on existing research, for thinking that least squares estimates are a reasonably good guide to the true causal effects of a first degree.

Moreover, we have some data (the BHPS) that contains at least a crude measure of ability – the number of A-level passes. We provide estimates that use this observed variable to try to control for unobserved ability and we find that the estimates of the effects of a degree on earnings is changed little. This again suggests that ability bias is probably not too problematic for our analysis.

Finally, while we have no way of controlling for selection into higher education, we can at least provide estimates that rely on *matching* the treated (graduates) and controls (non-graduates) on the basis of their observed characteristics. We use BHPS data which contain family background (father's education) and the number of A-level passes to weight the data to better ensure that we compare graduates with similar non-graduates. We do this using a "matching" method which estimates the probability (known in this literature as the propensity score) of being a graduate and compares the earnings of graduates and non-graduates with similar estimated propensity scores. The estimated results are quite similar to what we obtain with least squares, again suggesting that ability bias is not too problematic.

3.1 Subject of study

The existing literature on the effect of subject of study (referred to as "college major" in the US literature) is very thin (see LE, Sloane and O'Leary (2005), and Walker and Zhu (2011)) but the studies that do exist tend to report large differentials by major of study. Broadly speaking, differentials are higher for: science, technology, engineering and mathematics, medicine, dentistry and veterinarian studies (the STEM subjects); and for law, management and economics; this is compared to the Arts and Humanities. No studies, to our knowledge, make any attempt to deal with the complex selection issues associated with major choice that we expressed in the previous section. Nor do they allow for self-employment that is likely to be an important factor in some subjects. Nor do they allow for the impact of taxation or tuition fees. In addition, there are probably some non-pecuniary, as well as pecuniary, differentials and it is unclear that the labour market compensates for such non-pecuniary differentials by subject of study. In our view, while the existing results are interesting, we are a long way from being able to draw policy conclusions from them. It would be inappropriate, for example, to conclude that there was a shortage of one type of graduate relative to another.

3.2 HEI effects

There is a small literature on the impact of college quality (see Eide *et al* (1998) and Hoestra (2009) for the US, and Hussain *et al* (2009) for the UK). Hoestra (2009) is the most convincing study since it exploits a sharp discontinuity in admissions criteria to show that attending a "flagship" state university in the US increases earnings by about 20%. For the UK, in Hussain *et al* (2009) the effects of their proxy for HEI quality on earnings is statistically significant, but small compared to the overall return to higher education on average – they cite a one standard deviation in HEI quality results in a 6% earnings

difference. However, this study uses the DLHE surveys of graduates early in their careers when there is considerable noise in earnings differentials. It is unclear what the lifecycle effects would be from this data.

LE do not provide estimates of this effect because HEI is not recorded in LFS data. There is a worry that the omission of HEI type might (further) bias the estimates of effects of subject - because higher quality HEIs are likely to have a higher proportion of students studying the traditional high return subjects such as law and STEM.

In the UK Conlon and Chevalier (2003) consider the impact of attending a Russell Group institution and find small positive imprecise effects that were sensitive to the specification. We are not able to construct a quality proxy from our data so we follow Conlon and Chevalier (2003) in controlling for type of HEI: Russell Group, Pre-1992, Post-1992 and other – information which is available in the BHPS data. These results are reported later.

3.3 Postgraduate qualifications

The literature on the impact of postgraduate qualifications on earnings is similarly thin. A notable exception is Dolton *et al* (1990) for the UK but this uses a 1980 cohort of UK university graduates with earnings data observed just six years later so that they only identify qualification effects at a single, and early, point in the lifecycle – which we show below is a poor guide to lifecycle effects. Lindley and Machin (2011) use LFS data and estimate that the premium for a Masters (PhD) degree relative to a Bachelors degree rises from 8% (14%) in 1996 to 11% (24%) in 2009. LE also provides estimates – their average figures are approximately 9% for Masters and 15% for PhD.

3.4 Degree class effects

The literature on the quality of degree is particularly sparse. Surprisingly, there seems to be no research on US Grade Point Average for Bachelors degrees. Only LE provides any evidence for the UK. They suggest that a first (upper second, lower second) class degree returns approximately 33% (28%, 21%) more than the 2+ A-levels control group.

3.5 Part-time study

HESA administrative data suggest that approximately one-third of UK undergraduates study part-time. They tend to be substantially older than full-time undergraduates. They are disproportionately women. They are much more likely to have unconventional entry qualifications. They are much more likely to be studying for other undergraduate qualifications apart from a first degree. Their subject choices are different and they are concentrated in specific institutions (not least, the Open University). The only quantitative study of the effect of part-time vs full-time study on earnings, to our knowledge, is Callender *et al* (2011). This work uses the Longitudinal DLHE data collected forty-two months after graduation, merged with data for the same individuals just six months after graduation. The authors find, after controlling for other factors, that part-time study has only a small (positive) effect on the probability of having a high income after forty-two months. Regrettably neither LFS nor BHPS data record mode of study and our research is silent on this important issue.

3.6 Study later in life

LE reports a large fall in the earnings differential for graduates who graduate after the age of 25. There appears to be no other research in this spirit. Since this is recorded in LFS (from 2002) we can also explore this in our own specification and the results are reported later.

3.7 Summary

This brief review throws up surprising gaps in our knowledge. The reassuring message is that there is no clear indication that ability bias is large. Several methodologies suggest that conventional least squares estimation yields results that are surprisingly robust to the problem. However, this does not extend to the narrower question of the effect by degree subject since we do not have the data to test for such robustness in this context. So, for the moment, caution needs to be applied to existing estimates broken down by degree subject studied.

More encouraging is the literature on HEI quality. While some studies do suggest HEI type (or, more generally, HEI quality) has some effect these results seem to easily become statistically insignificant. Better data, as always, would help to be more definitive. Much more detail on earlier test scores, such as detailed A-level results, would give us more confidence.

4. Data

4.1 Datasets

There are several datasets that are helpful in this research area. The Labour Force Survey (LFS) is a large dataset that contains education information including: whether one has a first degree, subject, and class; earnings and hours of work; and a variety of characteristics including region of residence, ethnicity, and immigration status. It is first available from 1981 but earnings data first appear only in 1993, and it becomes a panel in 1997. Our data cover the period up to 2010. The BHPS is a panel of approximately 5000 households that have been followed since 1991 and data up to wave 18 are available. The original panel has been expanded in recent waves. BHPS includes data on earnings, hours, and characteristics, and education including an indicator for the type of HEI attended²⁷, but not subject studied or degree class. Further sources of data include the birth cohort studies - but these record earnings only infrequently during the lifetime²⁸; and the Destinations of Leavers from Higher Education (DLHE) dataset which is large and includes earnings, subject studied, and HEI, but covers graduates only and is recorded only at one point in time shortly after leaving HE (when there is a great deal of volatility in careers). Even the longitudinal version of this data that is available for a few cohorts records earnings just a little over three years after graduation.

We make extensive use of the large Quarterly Labour Force Survey (LFS) cross-section data pooled from 1993 to 2010. We also use panel LFS data for 1997 to 2010 derived from the cross sections - that is, matched across waves 1 and 5 by a unique identifier. We consider a parsimonious specification²⁹ to explain the variation in log wages across our sample. Briefly, our LFS sample construction is as follows. The LFS wage sample is based on employees aged 19-60 using Wave 5 of QLFS 1997-2010 inclusive and Wave 1 of 1993-1996. We include people who have a first degree (and/or a higher HE degree) together with those that left school, usually at the age of 18, with at least the minimum qualifications required for consideration for admission to university - two "A(dvanced)-level" qualifications³⁰. That is, we drop observations that would not have been able to gain

²⁷ Four types are identified: Russell Group, Pre-1992, Post-1992, and other (which are small institutions such as arts and music colleges). We pool other with Post-1992 because of the small size of this category. There are no significant effects of doing this on our results.

²⁸ We investigated the National Child Development Study – the 1958 birth cohort. However, we do not include the results in the report due to a much smaller sample size and inconsistent recording and updating of qualifications across sweeps which makes it not fully comparable to the LFS and BHPS in terms of sample composition. Moreover, it is also difficult to estimate the life cycle effects, given a lack of data points across the lifecycle for each individual.

²⁹ Estimating a more extensive specification that includes having a work-limiting health problem, union membership, and marital status makes little difference to our estimates.

³⁰ A-level qualifications, usually in three subjects, are normally examined at the end of a two year post-compulsory spell of schooling. Grades in these qualifications are used as criteria for university entry. A minimum of two passing grade A-levels are a necessary but not sufficient condition for entry. Entry requirements differ considerably across institutions and subjects but there is a well-developed application system for matching students to courses that ought to ensure that a student with 2 A-level passes can find a place on some course at some institution. Although we are mainly interested in the return to having an undergraduate (i.e. Bachelor) degree we also include in our sample for analysis all individuals who also have

admission to university even if they had wanted to attend. Thus, we exclude anyone with a highest qualification below 2+ A-Levels. We also exclude full-time students for whom we have no earnings history yet; and immigrants who arrived in the UK after age 5 who would have received some of their education abroad. The treatment group comprises all graduates, who may also have higher degrees (approximately 20% do not have 2+ A-Levels). The control group includes anyone who has at least 2 A-Levels but no HE degrees. Over a quarter of the latter have some higher vocational qualifications. This definition of treatment and controls reflects our choice to model the option value of pursuing HE or not - either route might lead to higher qualifications beyond a degree or 2+ A-levels. To minimize the effect of outliers, we then drop observations in the top and bottom 1% of the real hourly wage distribution (indexed to April 2012 prices using monthly RPI) within each gender-highest qualification cell.

4.2 Summary statistics

The dependent variable is based on the LFS-derived variable "hourpay", which is defined as the ratio of usual earnings to usual hours (from main job) including paid overtime. For pre-1996 LFS, where this derived variable is not available, we compute the variable using the same formula using the same raw data. Similar results hold using the reported hourly wage rate in the data. However, only a small proportion of graduates report an hourly wage rate. From 1997Q3 onwards, LFS also collects wage information in W1 (as well as W5). This allows us to construct a short wage-panel to be used in the first step of the two-step estimation to provide estimates of the lifecycle of earnings. Table 1 gives brief summary statistics.

Figure 4 shows the distribution of log real hourly earnings by sex and graduate status in the raw LFS data. Rather than present a histogram, which would have a jagged shape, we "smooth" the raw data to yield what is known as a kernel density plot. The solid lines show the distributions of graduate earnings, while the dotted lines show the distributions for non-graduates. There is clearly higher variance and lower mean amongst non-graduates. Figure 5 shows how log hourly earnings vary with age. There is clearly a strong age pattern in earnings and the differential between graduates and non-graduates is smaller earlier in the lifecycle. Figure 6 shows how the employment probability varies with age. Here the differences by graduate status seem broadly similar across age. Graduate employment is higher than non-graduate across all ages. Of course, these graphs conceal cohort effects as well as true lifecycle effects. Figure 7 shows the distribution of the data across degree subject studied.

higher academic qualifications. The proportion of graduates who also have a *higher* degree in the UK is significant (24% of male and 18% of female graduates in the sample) and also shows a steady growth across our sample period (from 20% in 1993 to 28% in 2010 for males and from 12% in 1993 to 21% in 2010 for females). Controlling for any higher qualifications, as in LE, is likely to lead to a lower estimate of the return to first degrees. However, excluding this group or controlling for higher degree makes no effective difference to our conclusions.

Table 1 Summary statistics of LFS wage sample (N=119,921):

Variable	Mean	Std. Dev.
Log hourly earnings	2.79	0.51
Age (years)	38.2	10.2
Immigrant	2.45%	
Non-white	3.10%	
London	12.5%	
Southeast	23.2%	
England (outside Southeast)	50.2%	
Wales	4.3%	
Scotland	7.5%	
Northern Ireland	2.3%	
Male	51.0%	
Graduate	72.9%	
Time (1991=1)	11.35	5.05

Note: Earnings are reflat to April 2012 prices using the RPI. English regions outside the South-East are grouped together because we find no statistically significant differences in our results when we control for each region separately.

Figure 4 Probability Distribution of Log Real Hourly Earnings: by Degree

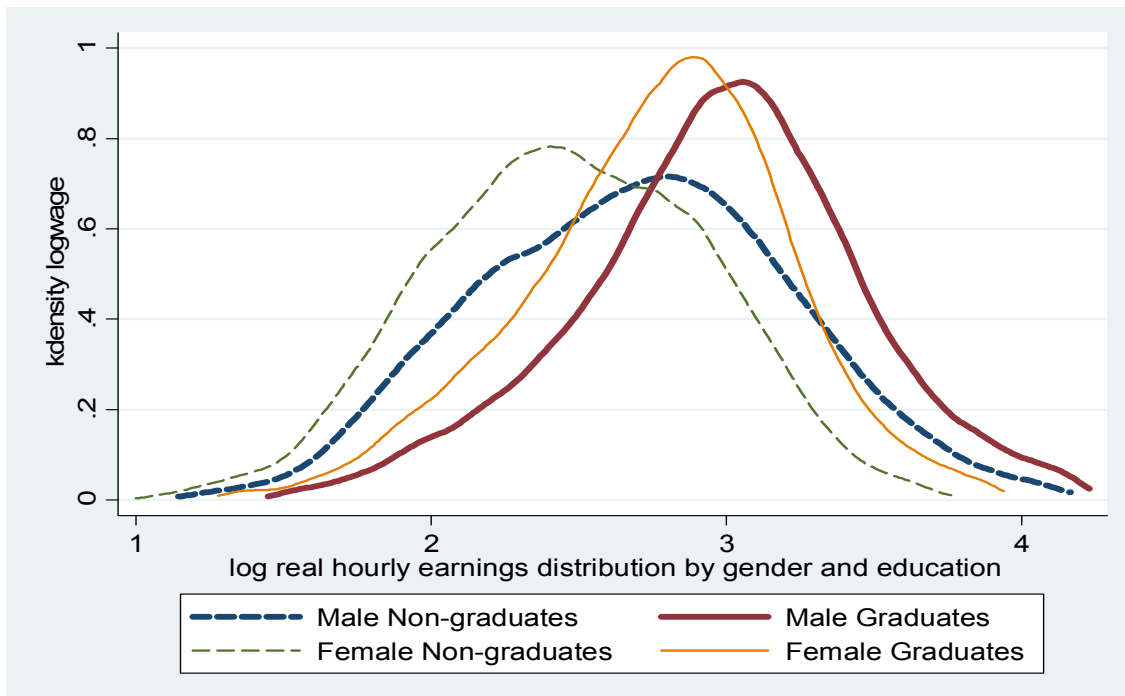


Figure 5 Log Hourly Earnings and Age: by Degree

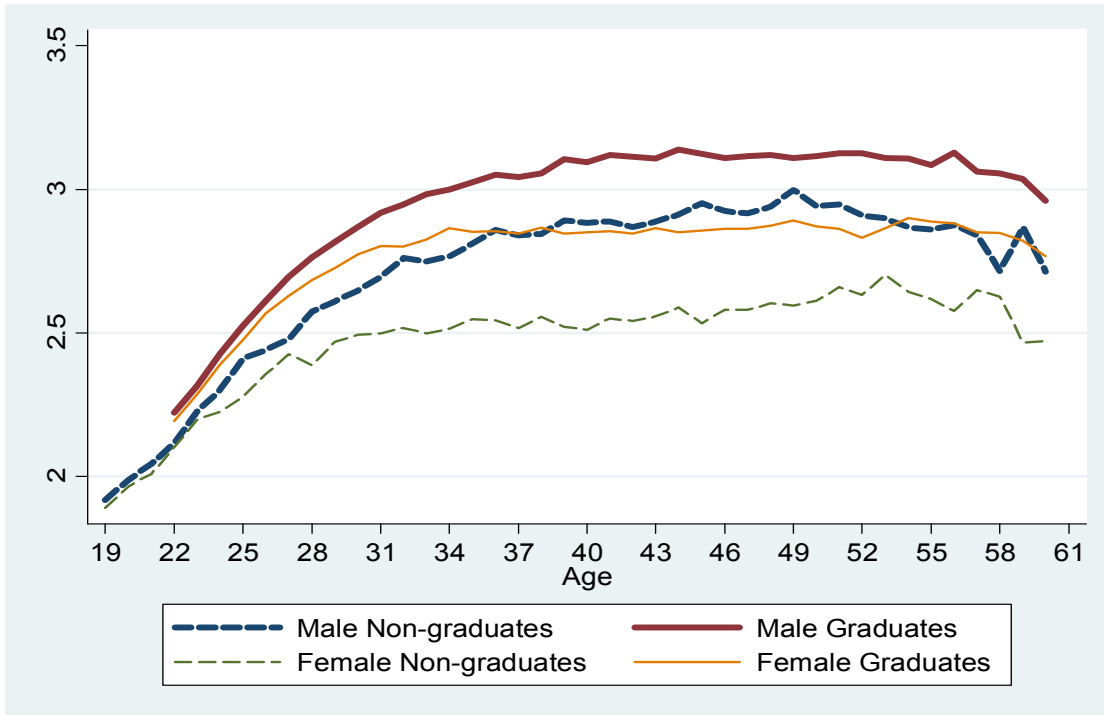


Figure 6 Employment Probability and Age: by Degree





Figure 7 Distribution of subject studied in HE (major)³¹: LFS with degree (%)



Note the high proportion with combined degrees. Unfortunately LFS data do not allow us to disaggregate this group consistently over the sample period because of changes in the classifications used across surveys.

Figure 8 shows the distribution of degree class for those with a degree (only available in the post 2003 LFS).

The proportion of our sample who have a degree can be broken down by birth cohort – Figure 9. This shows a large rise in this proportion for both men, from 21% for cohorts who were of university-going age (19) up to 1987 (i.e. birth cohorts up to 1968) to 32% by the

³¹ Our classification of subject studied is based on the Joint Academic Coding System (JACS), with minor modifications to allow for consistency of coding across different LFS survey years, and to merge similar subjects which also appear to have similar returns (in the case of linguistics and various language categories) in order to increase statistical power. However, when there is strong evidence of wage differentials within a JACS category, we split the sample whenever the cell sizes are sufficiently large (as in the cases of separating Nursing from Subjects Allied to Medicine, and Economics from Social Studies).

time the 1974 birth cohort reaches 19 in 1993 (approximately a 50% rise in the proportion of the flow) and an even larger rise for women, from 18% to 36% (a 100% rise in the proportion). These are huge increases over a period of just 6 years. The LFS data show that the expansion clearly starts in 1987 and ends in 1993 and this matches the trends in the official statistics in Figure 3, which are based on the population of college entrants, very well.

Figure 8 Distribution of Degree Class (%): LFS


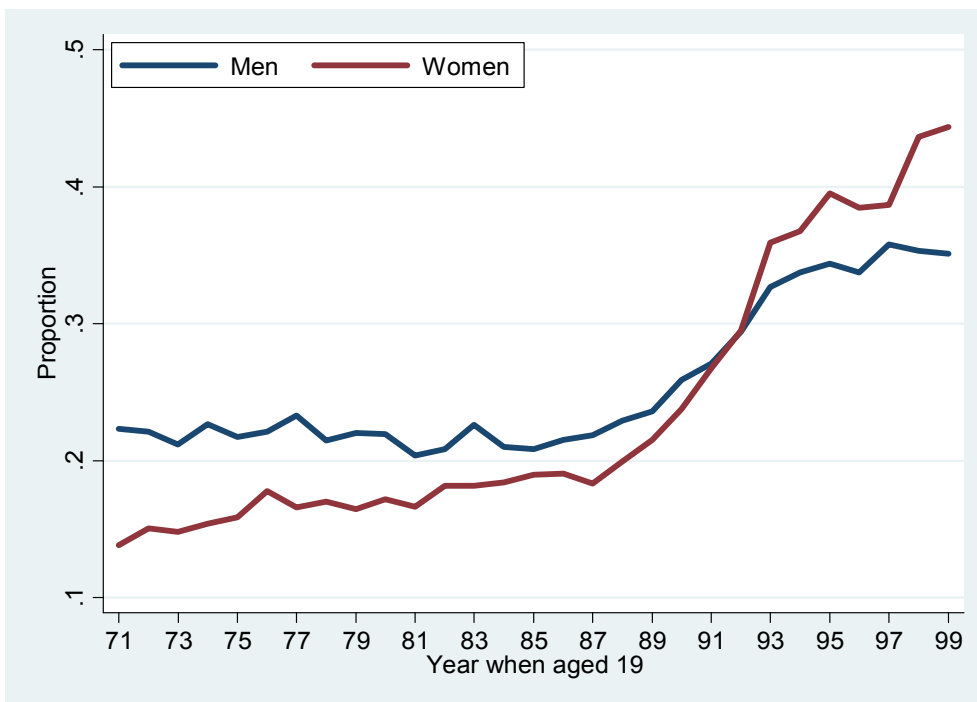



Figure 9 Proportion of birth cohorts with a first degree



We also use LFS to investigate the determinants of participation in employment. Figure 10 shows the raw LFS data proportions, which suggest that the probability of being in employment is somewhat lower for non-graduates (marked by the vertical lines) than graduates; it also varies across major.

Figure 10 Proportions employed by major: LFS



We make similar sample selections from BHPS as we have with LFS and Table 2 shows some summary statistics of both datasets. The BHPS wage variable is constructed in a similar way to LFS. While BHPS does not have subject major and degree class, it does have number of A-Levels which enables us to select a sample similar to the LFS. The two datasets are broadly comparable but the BHPS only records highest qualification. Thus, those with 2+ A-levels who also have high vocational qualifications (for example, chartered professional status) would be coded as being in the degree+ category. Nonetheless BHPS contains information on type of university/college attended: Russell Group, other pre-1992


(which we refer to as “Old”), post-1992 and other (which we group together and refer to as “New”).³² This is graphed in Figure 11.

Table 2 LFS and BHPS employment datasets: Means

	LFS	BHPS
Year born	1962.5	1963.6
Age (years)	38.9	37.0
% male	51.1	49.7
% employees	77.7	76.9
% self-employed	10.0	9.1
% unemployed	2.9	2.7
% inactive	9.4	11.4
% part-time	20.7	14.6
Survey years	1993-2010	1991-2008
% with 2+ A Level and no HE	28.9	26.4
% with degree or higher	71.1	73.5
% with higher degree	14.6	13.9
No. of observations	196,218	29,981

Note: BHPS observations are for the pooled data. Definitions of variables are consistent across surveys.

Figure 11 Distribution of HEI type: BHPS (%)



Note: BHPS wave 13 data.

³² The Russell Group of HEIs was formed in 1994. But graduates in BHPS were asked about this in 2003 so we assume that they were able to identify whether they went to a university that was or became a Russell Group institution up to 2003. Note that the Russell Group was extended recently and this will not be captured in our definition. Post-1992 institutions are the ex-polytechnics. Pre-92 are those outside the Russell Group but not ex-polytechnics. Other institutions are small specialist institutions such as art schools and music conservatories which we group with the Post-1992 institutions.

4.3 Summary

We confine our attention to LFS and BHPS datasets. The LFS dataset is large and long – for some of our analysis we can use data as far back as 1993. The LFS data confirm our basic knowledge of wage distributions. Wages are approximately (log) Normally distributed; men earn more than women; graduates earn more than non-graduates; and wages rise across the lifecycle but start to fall from the 1950s. The BHPS, although smaller in size, is quite consistent with the LFS data. It has some advantages – it is a very long panel of almost 20 years; and it contains some additional useful information such as the number of A-levels obtained, and the HEI type attended.

5. Statistical Analysis

5.1 LFS estimates for earnings and their robustness

We conduct a variety of analyses – from the simplest possible model to more complex flexible specifications that allow for a wider range of variation. Our simplest, baseline, specification is similar to the statistical exercise pursued in LE who assume that the effect of a degree is independent of age and all other variables and there are no cohort effects. Unlike LE we do not control for higher degrees (nor do we control for vocational qualifications beyond 2+ A-levels for those that did not pursue higher education). We also, for the moment, choose not to include a variety of additional variables that we feel are likely correlated with having a degree.

In Table 3, the coefficient on first degree of 0.205 (men) and 0.268 (women) are our headline figures corresponding to LE's reported 0.21 and 0.26 respectively. Higher effects for women are common in this literature and are likely to be a reflection of greater wage discrimination for women with low education relative to higher educated women. Thus, for similar specifications we obtain very similar estimates as previous work. We also find shallower age earning profiles for women relative to men – a common result in such cross section modelling arising from the weaker correspondence between age and work experience for women. The remaining estimates are also conventional: some strong regional effects on wages, a negative impact of being non-white, and significant real earnings growth over time of around 0.4% per annum. These coefficients are log wage differences by degree. These are approximately equal to percentage differences – but the approximation is only good for low values of these coefficients (less than 0.1). To turn these coefficients into true percentage effects one needs to subject them to the transformation $100(e^{\text{coeff}} - 1)$. Thus, the 0.205 log male wage difference that is associated with a degree corresponds to a 23% wage difference, and 0.268 for women corresponds to a 31% wage difference.

Table 3 **Baseline LFS Specification:**
Dependent variable log real hourly earnings

	Men	Women
First degree	0.205***	0.268***
Age	0.124***	0.086***
Age squared	-0.00134***	-0.00094***
Immigrant	0.008	-0.016
Non-white	-0.064***	-0.048***
Time	0.004***	0.004***
N	61122	58799
R ²	0.298	0.227

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

In Table 4 we disaggregate the degree indicator in Table 3 into a good degree (Upper second class or first) or not (Lower second class or below). The latter raises wages by approximately 17% for men, relative to the control group of 2+ A-level and no degree, and 23% for women (the same as LE). A good degree adds approximately a further 8% for men and 7% for women. These are large proportionate effects. Recall that all of these least squares results of degree effects are best interpreted as upper bounds because of potential ability bias. Other results are effectively the same as in Table 3.


In Figure 12 we break down the simple estimates by degree subject. The bars indicate the log wage differentials between graduates in each subject relative to those with 2+ A-levels (regardless of the subjects and grades of those A-levels – these are not recorded in either LFS or BHPS). These are comparable to the results in LE – although the categorisation of major differs slightly.³³ It is clear that the effects of studying different subjects can differ markedly, especially for men. Some subjects offer very high returns, such as Medicine Related, Economics, Law, and Business and Management, while other subjects provide only low returns for men, such as Languages and Linguistics, Art and Design, History and Philosophy, and Mass Communications. Women have, on average, somewhat larger returns, as in our previous estimates and is typical in the literature, and the differences across subjects are smaller.

Table 4 Baseline LFS specification with degree class

	Men	Women
Lower second degree or below	0.169***	0.226***
Upper second degree or above	0.248***	0.291***
Age	0.122***	0.080***
Age squared	-0.00129***	-0.00086***
Immigrant	0.003	0.024
Non-white	-0.072***	-0.037***
Time	-0.001	-0.005**
N	12580	14547
R ²	0.304	0.208

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** 5% level, *** and 1% level.

³³ Our classification of major is based on JACS, with minor modifications to allow for consistency of coding across different LFS survey years, or to merge similar subjects which appear to have similar returns (in the case of linguistics and various language categories) in order to increase statistical power. However, when there is strong evidence of wage differentials within a JACS category, we split the sample whenever the cell sizes are sufficiently large (as is the case when separating Nursing from Subjects Allied to Medicine, and Economics from Social Studies). Medical related includes pharmacists, medical practitioners (SOC 2010 code 2211 which include acupuncturists and various other specialists but not GPs), therapists of various types, and radiographers. We assume that all courses are 3 year duration.

Figure 12 Baseline LFS estimates by major


Note * Medical related excludes medical and dental graduates who exhibit a very high degree of self-employment for whom we do not have income data.

Source: Author estimates of WZ baseline model but with disaggregated degree subject.

5.2 LFS robustness

Our analysis above differs from LE in a number of ways. Thus, in this sub-section we examine the robustness of our results to the differences. While we prefer our specification, labelled WZ in Figure 13, there are arguments for and against the LE specification. But there is no specification that is preferable to another on purely theoretical grounds. Figure 13 provides a graphical summary of the coefficients from various specifications. First, we use the LE specification. Then we select our desired sample: only observations aged 19+ for non-graduates and 22+ for graduates and, in both cases, select only those aged less than 61 (for both males and females). The latter selection is motivated by the sharp fall in employment rates from age 60 onwards. The former selection is motivated by the large proportion of individuals in the control (treatment) group who have no earnings at age 18(21).³⁴ We also use quite a different set of control variables: we drop variables that

³⁴ This is much more of a problem in BHPS, where the earnings reported is for the previous year, than it is for LFS so changing the age selection helps us better compare the BHPS and LFS estimates. It is unclear what direction this would change the results – it is an empirical question.

Figure 13 Robustness of baseline results to specification and sample

Note: Author estimates using LFS data

control for marital status, dependent children, full vs part-time, temporary vs permanent contract, public vs private sector, workplace size, year dummies (we include a year trend instead to capture real wage growth induced by technical change), and (unspecified) interaction terms. Our rationale for doing so is that each of these variables (except the last) represents a choice by the individual. And those choices will be correlated with education. Including them or not may therefore affect the estimate of the variable of primary interest – depending on the sign of the correlation and whether that variable is positively or negatively correlated with earnings. For example, graduates are more likely to work in the public sector and public sector workers are paid more than private sector workers, on average. So including a public sector control variable is likely to rob the degree variable of some of its impact on earnings. This is our WZ baseline specification. The results are almost identical for men and just slightly (and insignificantly) lower for women.

Third, we add post-graduate qualifications as in LE. Since graduates with post-graduate qualifications earn, on average, more than those without this tends to reduce the effect of a degree because we have lost the option value of being able to undertake postgraduate study. However, the effect is small. Next we use the LE definitions of treatment and control groups. LE presents various estimates. When they are considering the marginal impact of postgraduate qualifications the control group is having an undergraduate degree. When considering the marginal impact of an undergraduate degree the control group is 2+ A-levels and the treatment is just undergraduate degree because separate controls are included for postgraduate. In all cases individuals who have vocational qualifications above level 3 (equivalent to A-level) are excluded. This implies that the more successful controls are excluded throughout. This seems to have a large effect for women, but not for men. Finally, we restore the restrictions of the LE sample. The results remain the same as the previous case.

Thus, apart from how we define our control group – where we include those with higher qualifications that were not obtained in higher education, the results are robust. Despite our misgivings, the LE results are similar to ours apart from this one (albeit) important exception – and this exception seems justified by our philosophical approach.

We also differ from LE in that we disaggregate degree subject in slightly different ways. Here we are motivated by considerations of robustness – if we slice the data very thinly then we will obtain results that are unlikely to be statistically significant. For example, we group languages and linguistics to ensure a large enough sample for men, we group engineering and technology to ensure a large enough sample of women - but we are able to separate economics from business and management. We have experimented with combined degree subjects but failed to establish significant differences between types of combinations and have grouped them into a single combined degree. The majority of combined degrees seem to combine similar subjects and the results that we obtain (using the post 1997 data where we can make these distinctions) suggests that the combinations reflect our estimates of the constituent parts.

One important deficiency in our statistical estimation work is our inability to model labour supply and education simultaneously. We were concerned that part-time workers might have a different attitude to work and, since work is the utilisation of human capital, this would be reflected in the estimates of the effect of degree. While we cannot construct a complete model of labour supply and education we can at least check whether part-time workers have different returns to a degree than full-time workers. Table 5 reports the results: first for our baseline (defined as FT and no degree) without controlling for part-time work; then controlling for part-time work; then allowing for an interaction between part-time work and a degree that allows the returns to a degree to differ between PT and FT workers. For men there is a large negative effect of being a part-time worker (but only 4% of graduate men and 5% of non-graduate men are PT). And there is no significant interaction effect – part-time male workers are estimated to have the same returns to a degree as full-time. For women, there is a large negative effect of PT work (and 25% of graduate women and 35% of non-graduate women are PT) on hourly earnings. PT female workers have a significantly (6%) *higher* return to a degree. That is, the PT penalty is larger for female non-graduates than for female graduates.

Table 6 considers the effect of obtaining a qualification later in life (age 27+ for graduates and 20+ for non-graduates)³⁵. 28% of male graduates and 31% of women graduates were late qualifying in this sense. Our results suggest that late male qualifiers, either late graduates or late qualifiers for other qualifications, tend to have lower hourly earnings (9% lower) but there is no effect of being late for females. Columns 2 and 4 separate out the late graduation effect from the late effect *for other qualifications*. Late qualifying women earn 11 log points more (about 13%) than the early qualifying women do; but late qualifying female *graduates* earn 4.5 (0.110-0.155) log points (i.e. about 4%) *less* than *graduates* who qualify at the normal age. For men, both effects are negative, with a penalty of 8 log points for late qualifiers and a *further* 3 log points for late qualifying graduates.

³⁵ This is therefore based on a post 2002 sample when age of qualification first became available. So the results here are not directly comparable with Table 5, although the differences are small.

Table 5: Baseline LFS specification and part-time working

	Men	Women
FT work * Degree	0.203***	0.227***
PT work	-0.335***	-0.235***
PT work *Degree	0.014	0.057***
	61122	59799
R ²	0.42	0.41

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

Table 6: Baseline LFS specification and late qualification

	Men		Women	
Graduate	0.213***	0.235***	0.277***	0.340***
Late qualification	-0.090***	-0.076***	0.001	0.110***
Late graduation		-0.033***		-0.155***
N	27848	27848	29840	29840
R ²	0.43	0.43	0.42	0.42

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

Table 7 is concerned with cohort effects. In this simple baseline specification we do not exploit the panel element of the LFS data. Thus it is based only on cross section data and in a cross section age and cohort (defined by date of birth) are perfectly collinear so it would not be possible to estimate the effects of age and cohort together. So, we select samples of young workers in each cohort. Since each cohort has the same average age we do not need to control for age and age squared in this analysis. We select workers aged 25-29 and 30-34 for two cohorts – those born up to August 1970 (the pre-expansion cohort), and those born afterwards (the expansion and post-expansion cohorts). In the simplest specification (columns 1 and 3) there is a significant graduate premium (0.203 log points for males and 0.277 for females), a significant (5-year) age or experience effect (0.242 log points for men and 0.174 for women) and a post-expansion *decrease* in wages (of 0.017 log points for men and 0.014 for women). But this specification is our highly restrictive model where the effect of post-expansion is assumed to be the same for graduates and non-graduates. Columns 2 and 4 correspond to a more general specification which allows the effects to differ by graduate status (degree or not), age interval (25-29 or 30-34), and cohort (expansion or not). The effect of age is now bigger for graduates than non-graduates consistent with our more general specifications reported in Table 8 of Section 5.4 below – being 30-34 rather than 25-29 for males (females) raises log wages by 0.173 (0.088) for non-graduates and $0.173+0.040=0.213$ ($0.088+0.083=0.171$) for graduates, both prior to the expansion.

The effect of the expansion (as opposed to being born before 1970) raises wages by 0.037 (0.007) log points for young males (females) more for graduates than for non-graduates. The triple interaction term, that appears as the last variable in Table 7, captures this same effect, but for the older age group. Both of these latter effects are statistically insignificant at conventional levels, and so the *difference* between these effects is also insignificant. That is, the post expansion older age group seems *not* to have had a significant rise in wages relative to the younger age cohort irrespective of their degree status. Thus, we

have no reason to think that young graduates in the post expansion period earned a lower graduate wage premium than young graduates did in the pre-expansion period – which is consistent with the results in Walker and Zhu (2010).

**Table 7: Baseline LFS specifications:
by pre-expansion cohort vs post expansion**

	Men		Women	
Degree	0.203***	0.161***	0.277***	0.238***
Age 30-34	0.242***	0.173***	0.174***	0.088***
Post expansion	-0.017***	-0.082***	-0.014**	-0.040*
Degree * Age 30-34		0.040		0.083***
Degree * Post expansion		0.037		0.007
Age 30-34 * post expansion		0.055*		0.061*
Degree * Age 30-34 * Post expansion		0.010		-0.034
N	18429	18429	19596	19596
R ²	0.114	0.116	0.123	0.125

Note: Omitted category is age 25-29, pre-expansion, non-graduate. * indicates statistically significant at 10% level, ** indicates statistically significant, at 5% level, *** indicates statistically significant at 1% level.

Our estimates suggest that the expansion of the higher education system has not had a deleterious effect on the returns to higher education. However, the HE expansion occurred at the same time as the economy went into a deep recession with the unemployment rate rising by over 50% in the space of just three years. In the face of the collinearity between the state of the labour market and the rise in the supply of graduates it would be difficult to discriminate between these two hypotheses. There is a literature on the effects of entering the labour market during a recession (see, for example, Oreopoulos *et al* (2006) for Canada, and Lang (2010) for the US) which suggests that such shocks can have large and quite persistent effects. Of course, for a given cohort graduates enter the labour market three years (or more) later than non-graduates. Indeed, part of the expansion in HE might have been due to individuals using HE to postpone entry into an unfavourable labour market.

In an attempt to separate out labour market effects from HE expansion we attempted to model the impact that the level of unemployment that reigns when one enters the labour market on the earnings for graduates and non-graduates. However, the unemployment data are not sufficiently refined and our results are too imprecise to be definitive.³⁶ Thus, we are not able to point to clear evidence that the HE expansion has reduced the returns to HE.

Finally, we examined the effects of dropping out. We defined a dropout as someone who has a positive difference between the age at which full-time education was completed and the age at which his or her highest qualification was obtained. Thus, a HE dropout is

³⁶ When we estimated our baseline model separately by cohort (pre expansion vs expansion and post-expansion) but when we included the level of regional unemployment on joining the labour market at 18 or 21 we found that the effect of a degree on wages rose (insignificantly) for men in the latter period relative to pre-expansion, although it did fall slightly for women. Unfortunately, the youth unemployment data do not go back far enough to be useful and our estimates rely on the ILO definition of unemployment for 16-64 year olds rather than youth unemployment. Graduate unemployment data are similarly unavailable for the period.

someone who completed their education at, say, 20 but did not obtain a degree since their highest qualification was coded as A-levels that were obtained at, say, age 18. When we do not control for HE dropouts (i.e. we leave dropouts in the control group with those who have 2+ A-levels but left education at the time that they obtained those A-levels) we find that the earnings effect of a bachelor degree is 24% for males and 32% for females. When we include control variables to capture dropouts (that's is we remove dropouts from the control group) we find that the earnings effect is 26% for men and 30% for women. The effect of being a HE dropout is estimated to be -2% for men, which is statistically insignificant, and -6% for women, which is statistically significant. Thus, it appears that HE dropouts actually earn less than those who never entered HE (the 2+ A-levels control group). However, failing to control for dropouts makes very little difference to our estimated effect of a bachelor degree.

5.3 BHPS estimates and robustness

The BHPS offers several opportunities relative to the LFS data.³⁷ Firstly, it contains information on family background (in particular, parental education) and on the number of A-level passes. Both variables are likely to be associated with unobserved ability and using these variables to try to control for ability might help to *tighten* the upper bound on the estimated earnings effect of HE that least squares regression is likely to provide – in other words, reduce the extent of bias. While, unfortunately, A-level grades are not available this does not change the point that the number is likely to be correlated with ability. Thus, our primary concern with the BHPS data is whether we find lower effects of higher education on earnings than we find with LFS where we cannot make any attempt to control for unobserved ability.


In Figure 14 we present five sets of BHPS estimates: our baseline specification, as in LFS, of the effect of a degree on hourly earnings – which give similar results for men as using the LFS data, but larger for women; then we add father's education (we include dummy variables for father having some qualifications, further education qualifications, degree level qualification and missing qualifications, with no qualifications as the baseline)³⁸ – and the results are the same.³⁹ Then we add the number of A-levels (dummy variables for missing, 0, 1, 2, 3 and 4+ A-levels) – and the results are virtually the same. Then we add both – and the results remain the same. And, finally, we drop those observations where the number of A-levels is missing – again the results are the same. The robustness of these results is consistent with the view that any ability bias in the baseline specification is likely to be small.⁴⁰

³⁷ Note that we are using the BHPS data as a cross-section because once education is finished it becomes fixed. Our analysis is conducted on all waves to capture the lifecycle variation, but HEI type information is only in wave 13 and we assume that HEI type has the same impact at every point in the lifecycle.

³⁸ We use father's education because mother's education is missing more frequently and because mother's and father's education are highly correlated.

³⁹ An interaction between controls for father's education and whether one has a degree also proved insignificant – suggesting that a degree has the same impact on earnings, on average, irrespective of family background.


⁴⁰ Altonji *et al* (2005) provide a method for estimating upper and lower bounds associated with an inability to control for unobserved ability. See Krauth (2011) for an application. Our implementation of this method, reported in Appendix Table A1, suggests that the bounds for women are likely to be quite tight – again indicating that unobserved ability may not be inducing much bias.

Figure 14 Robustness of BHPS baseline results

Note: Author estimates using BHPS data.


In Figures 15 (for men) and 16 (for women) we go back to the baseline BHPS specification and consider the role of HEI type. We collapse the very small group that are categorised as “other” HEI type with the “New” universities (Post-92 group), and so only consider this merged group against the Russell Group, and the “Old” universities (the Pre-92 group). The differences by HEI type in the baseline specification are quite large. Men (women) who graduate from an “Old” HEI earn 12% (7%) more than a similar graduate from a “New” HEI, and a Russell Group graduate earns a further 4% (2%). New universities seem to add significantly less value than the others – although the 20% premium relative to the control group is still highly statistically significant.

Figure 15 Robustness of HEI type results to specification and sample: MEN



Note: Author estimates using BHPS data

Figure 16 Robustness of HEI type results to specification and sample: WOMEN



Note: Author estimates using BHPS data

However, the additional premia for Russell and Old in the baseline estimates could be due to selection effects – students at these institutions could simply be better. These results do not tell us what would happen to the earnings of someone who had attended a New University if they had wanted to attend and been admitted to another type of HEI.

Successive blocks in Figures 15 and 16 attempt to provide some control for selection into HEI type. We add father's education, defined as either graduate, 2+ A-levels, or less; we add the number of A-levels, defined as 2, 3 or 4+; and we then add both controls; and finally we add a dummy variable that records whether information on the number of A-levels is missing. We find that the estimated differentials across HEI types do tend to fall as we add more controls that seem likely to capture the effects of selection by ability. Even though our controls for ability are imperfect, the differences in the final block are now not statistically significant, suggesting that there are no statistically significant effects of HEI type.⁴¹

5.4 Two step method

Two important weaknesses of the simple analyses so far are that: they impose the assumption that the impact of higher education is independent of age, and other characteristics; and they use only (pooled) cross section data which makes it problematic to estimate the effects of age on earnings over the lifecycle separately from the effect of being born in a particular cohort, and separately from the effect of calendar time. Clearly birth cohort is survey year minus age, and calendar time is cohort plus age. Thus, there is a high degree of collinearity between effects that we wish to capture. We wish to identify lifecycle effects so that we can compute the present value of lifecycle earnings; we wish to identify cohort effects because of our suspicion that the expansion of HE will have depressed the effects of HE on earnings; and we wish to identify the effect of calendar time to capture the productivity gains associated with workers of different skill levels. If we fail to independently control for all these effects it is likely that the effects of any one will be biased by its dependence on the others.

Our two-step method, explained earlier, attempts to separate out the relevant effects. The first step takes a panel of individuals and estimates the relationship between age and earnings to model the lifecycle pattern. Because this is a panel of people it is effectively a set of cohorts of individuals who are simply getting older. The second step fixes this lifecycle pattern and then estimates the time and cohort effects from the (pooled) cross section data.

Moreover, greater flexibility can be afforded by estimating *separate* equations for the earnings of graduates and non-graduates. This allows age – log earnings profiles to have both different slopes as well as different heights. LE impose the assumptions that: there are no cohort effects and there are no calendar time effects (although they do *impose* productivity growth in their simulation work). Moreover, they also assume that lifecycle age – log earnings profiles are parallel and independent of the other control variables.

While it is straightforward to allow for non-parallel age – log earnings profiles it is difficult to overcome the problem of identifying time and cohort effects separately from age effects. Fortunately the LFS data contain (from 1997) earnings information at wave 1 and 5, approximately 12 months apart. Moreover, BHPS contains information every year for up to 18 years. This information tells us how earnings have changed over one year in LFS, and over up to 17 years in the case of BHPS, of the lifetime for people of a particular age.

⁴¹ To the best of our knowledge there are no estimates in the literature that identify the causal effect of HEI type on earnings for the UK but our analysis suggests that selection is responsible for at least some of the differentials that we can observe in the raw data and in our simplest specification. It is difficult to see how one could resolve the extent to which the differentials are due to selection with available data.

Thus, by using the wage change data on people of all ages we can estimate how wages change from year to year across the lifecycle at a point in calendar time, using either LFS or BHPS.

The LFS is available for a large sample, but only over a one year interval. The BHPS is available for a small sample but over a much longer period of time. In both cases the data in the panel are used separately from the LFS cross section information so we choose between the two panel datasets on the basis of the precision of the estimates that they yield. We find that the LFS estimates of the lifecycle parameters are more precise than the BHPS panel estimates of the same parameters.

Table 8 shows the LFS estimates, based on observations that have two earnings observations. These estimates suggest that age-earnings profiles do indeed seem to be parallel for men (the reported coefficients that underlie the lifecycle are effectively identical for graduates as for non-graduates). However, for women, graduates tend to have steeper age earnings profiles than non-graduates.

Having estimated the first step we then constrain these parameters to their estimated values in the second step and estimate the remaining parameters *using the LFS cross section data*. So in Table 9 we impose these first step estimates on the earnings model and provide estimates of the remaining parameters. When estimating this second step, we found that cohort effects were statistically insignificant and so Table 9 presents estimates of the second step that omit them. The coefficients in Table 9 are similar to the earlier results. We find that there are no significant regional differences in graduate earnings differentials and these coefficients are not reported. We further find that the graduate vs non-graduate earnings differential is higher for non-white men, than white, by approximately 2% (i.e. -0.064 minus -0.084). This is a common finding in the literature which may reflect higher levels of labour market discrimination for non-graduates than for graduates. However, the graduate differential is approximately 4% less (-0.064 minus -0.028) for non-white women compared to white women.

Table 8 Estimates of the Lifecycle Age earnings profile - LFS

	Men		Women	
	D=0	D=1	D=0	D=1
η^D	0.13031***	0.12997***	0.07815***	0.09829***
γ^D	-0.00122**	-0.00113***	-0.00055***	-0.00071***

Note: D=0/1 refers to non-graduate/graduate. η^D corresponds to the age coefficient, and γ^D to the age squared coefficient in the baseline model.

* indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

Table 9 Estimates of the level of log earnings: LFS

	Men		Women	
	D=0	D=1	D=0	D=1
Time	0.01429***	0.01908***	0.01948***	0.02552***
Immigrant	0.01830	0.01116	-0.01947	0.00318
Non-white	-0.08339***	-0.06397***	-0.02787	-0.06408***

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

Importantly, these estimates imply that real earnings growth over time is *higher* for graduates than non-graduates - by 0.5% pa for men (0.019 minus 0.014) and 0.6% pa for women (0.025-0.019) - reflecting the skill biased technical change that has been occurring over recent decades. Of course, there is no guarantee that the future will be just like the past but these are our best estimates and we can reject the assumption that the time coefficients are the same for graduates and non-graduates, for both men and women. While these figures might appear small, over a working lifetime they compound to large differences in present values, especially for women. This is probably the most important difference between the estimates here and those in LE. Note that it was only possible to separately identify these effects through using the two-step methodology to estimate equations by graduate status.

We can extend this two-step analysis in a number of ways. For example, we can apply this two-step methodology to the LFS data broken down by subject in the first step, and by subject and cohort in the second step. These results are not reported here because of our concerns about differential ability bias by subject. We can also introduce cohort effects into the two step methodology, but we do not report those results here because we find no significant such effects – despite the expansion in HE across cohorts. Thirdly, when we control for drop-out status we find that the results remain the same - it seems that dropouts are effectively identical to those that had 2+ A-levels but never attended HE.⁴²

Finally, we can explore how the earnings differentials, on average for all subjects, might vary across the distribution of the unobservable determinants of earnings. Here we are thinking, specifically, of how the relationship between wages and education might vary across unobserved ability. This extension is known as *quantile regression* (QR) – unlike conventional regression estimation which models how the mean prediction responds to a change in a regressor, a quantile regression models how a change in a regressor would affect any chosen quantile in the distribution of predictions. This allows the researcher to model responses of the whole distribution (of earnings, in our case) to its determinants, not just the average of the distribution. Thus, QR attempts to address the question of how education affects individuals of different levels of ability. The interpretation of our most extensive QR results is complex so results that are equivalent to our simplest baseline specification (as in Table 3) are presented in the Appendix - but we show in Section 6 below the implications of two step results imply sizeable differences in NPVs across the ability distribution.

5.5 Employment probability

Finally, we explore the possibility that higher education affects the probability of employment as well as earnings conditional on employment. Using a baseline specification we find comparable results to LE. We allow for a lifecycle effect using our two step methodology, applied to the BHPS panel, and we also allow for several control variables in the second step applied to LFS cross section data. These results are reported in Tables 10 and 11. Here the dependent variable is whether the individual is employed or not, rather than earnings as before. Thus, these results capture the effects of the lifecycle and observable characteristics on the probability of employment. Table 10 shows the lifecycle

⁴² Three quarters of our HE dropouts had an age completed full-time education only one year greater than age at which they attained their highest qualification, which would be consistent with dropping out at the end of the first year of university study.

patterns while Table 11 shows the second stage. The lifecycle parameters capture the usual inverted U shape that we saw in Figure 6. The second step parameters show strong positive London and SE effects and strongly positive time effects for graduates vs. non-graduates and, typically, strong negative non-white effects. Rather than estimating a Probit model, these are *linear probability* models where the indicator variable for employment is regressed against the explanatory variables.^{43,44}

Table 10 indicates that both male and female graduates have steeper age-employment profiles than their non-graduates counterparts. With the age effects fixed, Table 11 suggests that the graduate non-graduate employment gap is 4% higher for non-white men, and 9% for non-white women. On the other hand, the native-immigrant employment differential is not precisely determined for either men or women. Finally, the coefficients on the time variable imply that the growth rate in the employment probability is significantly higher for graduates than non-graduates, regardless of gender. This is again consistent with the notion of skill biased technical change.

Table 10 Estimated Parameters of the employment probability: BHPS

	Men		Women	
	D=0	D=1	D=0	D=1
η^D	0.05183***	0.12830***	0.03764**	0.10521***
γ^D	-0.00071***	-0.00149***	-0.00062***	-0.00125***

Note: D=0/1 refers to non-graduate/graduate. η^D corresponds to the age coefficient, and γ^D to the age squared coefficient in the baseline model.

* indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

Table 11 Estimated Parameters of the employment probability: LFS

	Men		Women	
	D=0	D=1	D=0	D=1
Time	-0.00227***	0.00824***	-0.00673***	0.00787***
Immigrant	-0.02257	-0.01822***	-0.01794	-0.03890***
Non-white	-0.07687***	-0.03550***	-0.13424***	-0.04549***

Note: Regional coefficients not reported. * indicates statistically significant at 10% level, ** indicates statistically significant at 5% level, *** indicates statistically significant at 1% level.

⁴³ We also estimated the employment probability equations by degree subject but the reservations that we have about giving these estimates a causal interpretation in the earnings equations apply here too. Simple, one step estimates of the effects of each degree subject on employment are provided in the Appendix to facilitate comparison with LE.

⁴⁴ As with our earnings modelling we also controlled for drop-out status. The effect of being a drop out on male employment was very small and statistically insignificant. The effect on female employment was a statistically significant 2.5% reduction in employment probability.

6. Simulated Net Present Values

6.1 Method

Our estimation allows for six regions, two ethnicity types and two immigrant types. Thus the methodology takes predicted earnings and predicted employment for each of the 24 types of individuals and calculates the net present value (NPV)⁴⁵ of earnings, including any maintenance loan from the SLC while attending university, net of the opportunity costs incurred from not working while young (given by the predicted earnings for someone with 2+ A-levels and the same characteristics from the age of 18 to 21), and net of income tax, NI and VAT, and net of student loan repayments. Throughout we use the 2012 tax and SLC parameters and rules.

We assume that all degrees are 3 year duration. We have dropped medical students since their income is not well captured in our data. There are many pharmacy and architecture degrees that are of 3 year duration (and the LFS since 2002, which tells us about the age at which qualifications are obtained, broadly reflects this) – so we assume that such students join the labour market and get paid after 3 years. To the extent that such graduates may be taking degree programs that are longer than 3 years we will be underestimating their returns. Nurses face fees that are paid by the NHS, and they also receive a living allowance. We allow for all of these distinctions in our simulation work.

Gross earnings are predicted for each individual assuming that they are 18 in 2012 and work until the age of 68 (the current projected state pension age for someone 18).⁴⁶ For each individual we generate a large number of *draws* from the distribution of the error term, which has a zero mean, using the estimated variance of the error term (in the second stage earnings equation). We then add the draw to the predicted earnings and we do this for each draw. We assume that this draw is a fixed effect that remains switched on throughout the lifecycle.⁴⁷ This allows us to create a distribution in predicted earnings that reflects what we observe in the actual earnings distribution. The importance of this procedure stems from the nonlinear transformations that we make to the predicted gross earnings - for example, to allow for predicted earnings that exceed the higher rate tax threshold, and for earnings less than the £21,000 loan repayment threshold (this threshold will apply in 2015 so we deflate this to 2012 assuming 2% pa inflation, and we reflate with real earnings growth over time which we assume will be 1%).

⁴⁵ We assume a discount rate of 3.5% throughout. The Green Book, http://www.hm-treasury.gov.uk/d/green_book_complete.pdf, suggests that a lower rate might be used for investments beyond 30 years.

⁴⁶ Below, we provide some robustness checks to see if the assumption of retirement at age 68 makes a substantive difference relative to 65 for men and 60 for women (assumed by LE).

⁴⁷ We are implicitly assuming that this error term represents some unobserved factor that affects earnings on a permanent basis. The alternative is to interpret this error term as “luck” which quickly disappears. We could adopt this alternative assumption by taking a new draw every year over the lifecycle. While we cannot identify which interpretation is correct from our estimates we would not expect this to make a substantive difference to the calculations.

Had we simply used the point prediction from the estimates then we would compress the predicted earnings distribution considerably.⁴⁸ This would imply that the chances of predicting an extremely low or high level of earnings would have been small – we would have far too few higher rate taxpayers and too few young graduates below the £21,000 limit. This would generate considerable biases in the results - both for students and the government.

We then compute the NPV associated with this stream of incomes net of taxes and loans. In some of the tables below we demonstrate the effect of allowing for the possibility that parental income is below £25,000, where scholarships and bursaries become relevant, and above £42,000 where they are no longer relevant. Throughout, we assume that fees are £9000 in England (and lower in other parts of the UK) and maintenance loans are £2700 per annum, and we assume that students from low income households receive a National Scholarship Programme (NSP) award of £3000 in the first year and £1000 thereafter, and that universities pay bursaries of £1000 each year. We also consider the case of no university scholarships. If fees were lower the NPV calculations would be correspondingly higher by approximately this lower amount times 3 (apart from a small adjustment for discounting).

Inspection of the VAT paid at different points of the net income distribution (see Figure 10.1 in <http://www.ifs.org.uk/budgets/gb2009/09chap10.pdf>) suggests that approximately 10% of total expenditure is paid in VAT across a wide range of the income distribution. Thus, we simply subtract 10% from predicted expected incomes to capture VAT for both graduates and non-graduates.⁴⁹

The employment probability is predicted from the employment model. We multiply the 24 employment predictions at each age by the earnings predictions at each age.⁵⁰

6.2 NPV results

We report figures per student. Approximate aggregate figures for revenue could be obtained by assuming that there is a flow of 480,000 undergraduate students each year. Table 12 shows results for a naïve model which uses estimates that: do not discriminate by major, assumes that a degree takes three years to complete, and makes no allowance for spells of non-employment over the lifecycle. The NPV figures for non-graduates are provided to give an idea of scale. The graduate premium reported here is the excess NPV associated with a degree over no degree.⁵¹ The government revenue refers to the revenue

⁴⁸ Note that the measure of fit (R^2) for our earnings equations imply that we explain only around 30% of the variation in log earnings. Neglecting the variance in the unobservable earnings determinants would result in very compressed predicted earnings distributions.

⁴⁹ We assume that all of this expected net income is spent. LE assumes that the marginal propensity to consume is 0.6 so that only 60% is assumed to be spent. Of course, in a lifecycle model all net income is spent at some stage in the lifecycle (except for any bequests) so that, in the long run, we are correct.

⁵⁰ There is no clear correct procedure for imputing incomes for those who are predicted not to be in employment. For the sake of simplicity we assume that those not in work have no income. An alternative would be to impute some level of welfare payment. We provide a robustness check below for this.

⁵¹ Our NPV results are best regarded as upper bounds since they are based on estimates that may be contaminated by ability bias. The same problem pervades previous research, including LE. However, the checks on the robustness of our results above, and in the Appendix, suggest that the upper bound is likely to be “tight” – i.e. the bias is probably small.

that is received from a graduate relative to a non-graduate. Note that this is a lifetime NPV not a *per annum* figure.

Table 13 allows for the employment probability to be less than 1. This accounts for the lower NPV figures than in Table 12. Since this probability is lower for non-graduates than graduates the graduate premium is larger and the government enjoys a further revenue bonus.

We have also attempted to allow for HE dropouts, which we find are approximately 9%. In Table 14 we explore the sensitivity of the results in Table 13 to dropouts. We assume that 9% of students drop out from each year group and that they then have the same wage and employment lifecycles as those non-graduates who are t years younger, where $t = 1/2/3$. That is, we assume that the higher education of dropouts has no effect whatsoever and that dropouts do not accumulate the work experience that those who did not enter HE have gained. We assume that 50% of the non-graduates in our control group are dropouts. Thus, in Table 14, dropouts appear in the non-graduates. The NPV of non-graduate lifecycle earnings net of the loan repayments arising from the dropouts in this group is lower than in Table 13. The NPV associated with the graduate premium is correspondingly larger. Note that the government revenue associated with graduates cannot now be entirely attributed to the university system. This system did, after all, admit those dropouts who end up being worse off than if they had never been admitted.

Table 15 allows for degree class effects (and hence uses only the more recent data for both graduates and non-graduates). Table 16 breaks the results down by subject, despite our reservations about what one might conclude from this. It is clear that there are some large differences between men and women for some of the smaller groups – there are few female engineers and few male Historian/Philosophers for example so that the precision of our estimates are likely to be weak. Table 17 varies the loan and scholarship profile with predictable results.⁵²

⁵² Allowing for the tuition subsidies to HE, which LE report as £891 pa per student, leaves the student NPV unchanged and reduces the benefit to the government by approximately 1%. The precise level of subsidy remaining in the system is unimportant for our findings.

Table 12 NPV and Government Revenue per capita, for naïve model assuming full employment

	NPV		Government revenue	
	Non-graduates	Graduate premium	Non-graduates	Graduate premium
Men:	762	+148	515	+271
Women:	676	+181	418	+299

Note: Measurement unit: £1,000. Non-graduates here, and below, include the dropouts.

Table 13 NPV and Revenue for baseline model with varying employment over the lifecycle

	NPV		Government revenue	
	Non-graduates	Graduate premium	Non-graduates	Graduate premium
Men:	606	+168	406	+264
Women:	475	+252	287	+318

Note: Measurement unit: £1,000

Table 14 NPV and Revenue for baseline model with varying employment over the lifecycle and dropouts

	NPV		Government revenue	
	Non-graduates	Graduate premium	Non-graduates	Graduate premium
Men:	575	+199	294	+376
Women:	450	+277	206	+399

Note: Measurement unit: £1,000

Table 15 NPV and Revenue allowing for degree class effects

	NPV		Government revenue	
	Non-graduates	Graduate premium	Non-graduates	Graduate premium
Men:	611	+65 (2II) +141 (2I&I)	416	+128 (2II) +221 (2I&I)
Women:	496	+105 (2II) +190 (2I&I)	296	+108 (2II) +264 (2I&I)

Note: Measurement unit £1,000. 1 = first class degree, 2I = upper second; 2II = lower second

Table 16 NPV and Revenue by subject: Graduate premia

	NPV		Government revenue	
	Men	Women	Men	Women
Medical related	429	454	698	639
Nursing	170	-7	383	7
Bio/Vet/Agric	174	117	271	149
Physical Sci	237	123	366	151
Maths/Comp	100	243	146	306
Engineer/Tech	21	680	57	1108
Architecture	288	-193	460	-197
Social Study	-86	266	-88	354
Law	431	120	693	130
Economics	335	902	499	1478
Business & Man	256	149	391	177
Mass Com	3	95	30	112
Ling/Lang	161	123	258	161
Hist/Phil	557	113	954	149
Arts/Design	-111	111	-114	145
Education	103	396	174	548
Combined	332	326	497	442

Note: Measurement unit: £1,000

Table 17 NPV and Revenue with scholarships and bursaries

	NPV		Government revenue	
	Non-graduates	Graduate premium	Non-graduates	Graduate premium
<i>a) £3.3k fees with no bursaries/NSP, plus £2.7k maintenance</i>				
Men:	606	+180	406	+267
Women:	475	+267	287	+321
<i>b) £6k fees with no bursaries/NSP, plus £2.7k maintenance</i>				
Men:	606	+174	406	+266
Women:	475	+258	287	+320
<i>c) £9k fees with no bursaries/NSP, plus £2.7k maintenance loan (Baseline)</i>				
Men:	606	+168	406	+264
Women:	475	+252	287	+318
<i>d) £9k fees, no bursaries/NSP, no maintenance loan</i>				
Men:	606	+175	406	+266
Women:	475	+258	287	+320
<i>e) Low income parents - £9k fees with bursaries/NSP, and £3.25k maintenance grant (no maintenance loan)</i>				
Men:	606	+181	406	+246
Women:	475	+264	287	+300

Note: Measurement unit: £1,000. National Scholarship Programme award (NSP) is assumed to be £3k in the first year, £1k in subsequent years; University bursaries are assumed to be £1k in all years.

The estimates here, for both the average NPV to the student and the per student revenue effects for the government, are substantially larger than those reported in LE. Figures 17 & 18 present a comparison of WZ and LE. There are two main reasons for the discrepancy. First, our simulation methodology captures the variance in earnings more effectively than LE and so is designed to better compute higher rate tax liability – this will be particularly the case for graduate men who have higher earnings, on average, than graduate women, in part because of the latter’s lower participation rate. And, at the other extreme, we better capture the benefits to the student with earnings less than £21,000 for the same reason. This particularly affects women because female graduates have lower mean earnings than male graduates. The effect of better capturing the higher rate tax revenue from high earning individuals will decrease the gain in NPV from a degree; while the effect of better capturing the value of the loan subsidy to low earning individuals will increase the private benefit.

Second, our estimates themselves generate higher *gross* gains from a degree – mostly because we allow for differences in the trend in the graduate earnings differential associated with skill biased technical change. The NPV figures for the students, comparing our method with LE, are given in Figure 17.⁵³ The corresponding government revenue comparison is given in Figure 18. These figures correspond to the numbers in Table 13. The dominant effect is the greater predicted revenue associated with our estimated higher wage growth for graduates relative to non-graduates. However, there are differences between our work and LE that would generate either lower or higher private benefits and government revenue predictions. For example, our predicted earnings figures will have greater variance than in LE – this would result in more revenue (and lower net income) because of a greater chance of facing the higher rate of income tax, especially for men. Similarly, our greater variance in predicted earnings has implications for loan repayments that will affect net income figures. For example, we are more likely to predict incomes below £21,000 that would imply no repayments.

6.3 Sensitivity Analysis

In Figures 19 and 20 we simulate the effects of changing the assumptions that we adopt in our work. The baseline figure above is for a retirement age of 68 and no payments to those not in employment. We consider the effect of changing the retirement age to 65/60 as in in LE (from 68). We also consider the consequences of allowing for welfare payments during periods of non-employment. Paying £70 pw to those not in employment has almost no effect on these results because the incidence of non-employment for graduates is relatively low, especially for men. The effect of an extended working life is large for women - because the extension is for 8 years rather than the 3 for men. In the spirit of our forward looking philosophy it seems most appropriate to assume retirement at 68 in the remaining analysis that is concerned with fees/loans.

⁵³ We explored the extent to which our results differ from LE. If, like them, we had imposed equal wage growth on both graduate and non-graduates we would reduce NPVs for males by 32% (from 168k to 115k), and for females by 26% (from 252k to 186k). Further imposing equal employment growth rate would reduce NPVs by another 26% and 51% for men and women respectively (to 85k and 95k). Thus the flexibility of our modelling is crucial to the findings. Nonetheless, if we did adopt these extreme restrictions we would still predict effects approximately as large as those in LE.

Figure 17 Per student NPV of the private benefit associated with obtaining a degree (£k per student)




Figure 18 NPV of government revenue associated with obtaining a degree (£k per student)




Figure 19 NPV of private benefits (£k per student): sensitivity

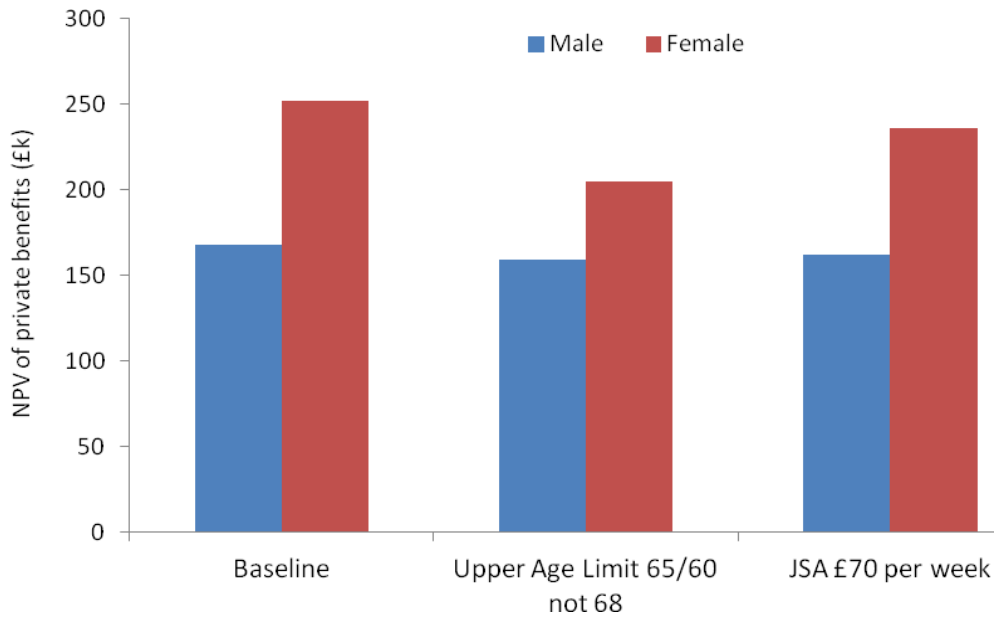


Figure 20 NPV of government revenue (£k per student): sensitivity


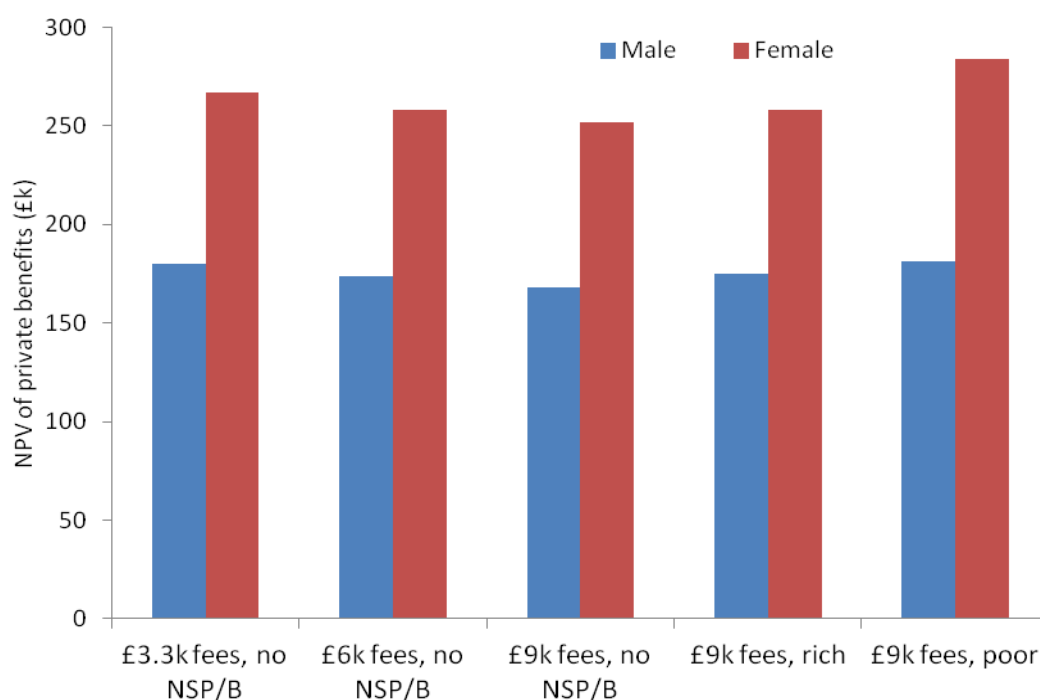


Figure 21 NPV of private benefits (£k per student): sensitivity

Note: NSP/B indicates receipt of a National Scholarship Programme and University Bursary.

We present, in Figure 21, the simulation results, from Table 17, for the NPV of private benefits for a variety of types of individual. The first block is for the old fee regime for an individual who is assumed to not receive bursaries or scholarships; the second is for £6,000 fees; the third is for the fees of £9,000 currently being charged by many institutions; the fourth block is for £9,000 fees for a student from a high income background who does not take a further loan to cover maintenance; while the final block is for a student from low income background who receives a (£3250) maintenance grant as well as a scholarship/bursary package. We can see the clear effect of the impact of higher fees by comparing the first, second and third blocks. The fourth block is for a student from a rich household who does not require a maintenance loan and is not eligible for a maintenance grant, while the final block is for a student from a low income household who is eligible for a maintenance grant and so does not take a maintenance loan.

Finally, in Table 18 we show the NPV on students and government on average, and broken down by decile of the distribution of log earnings. The differences in the results across successive quantiles is relatively small through most of the distribution, but there are distinctly greater effects much higher up the distribution which suggests that HE provides a complement to innate ability. This gives an idea of how the distribution of the net gains from a degree affect different points in the distribution of earnings. The net gains are fairly stable up to the median, but rise somewhat as we get close to the top one or two deciles.

Table 18 NPV and Revenue across Deciles

MALES	Average	1st	2nd	3rd	4th	5th	6th	7th	8th	9th
Student	168	173	168	163	161	167	157	154	166	206
Government	264	251	280	237	239	240	284	247	276	327
FEMALE										
Student	252	247	240	241	241	245	255	252	285	265
Government	318	297	302	303	263	293	317	334	387	372

Note: Measurement unit: £1,000

6.4 Summary

Our benchmark estimates of the private benefits in terms of the NPV of the lifecycle difference in earnings are large: on average, £168k for men and £252k for women. These are economically significant differences relative to the figures in LE. We also estimate much larger returns to the government through higher tax revenues. There are good reasons why our methodology better captures these returns.

A good part of the differences are due to lifecycle gains associated with the higher real earnings growth for graduates relative to non-graduates. This is a reflection of the skill-biased technical change that has prevailed over recent decades. However, our robustness checking showed that even if we imposed the assumption that these differential real earnings (and employment) growth were equal for graduates and non-graduates we would still find effects that were of the same order of magnitude as LE.

A corollary of the high estimated real wage growth of graduates relative to non-graduates, and hence relative to average earnings, is that the rate at which graduates pay down their debt will be faster, and the interest rate they face while doing so will be higher. These implications are incorporated in our simulations for graduate net of loan repayment incomes and the higher tax payments are reflected in the simulated stream of revenue for the government. In addition, the estimates would have positive implications for the cost of the implicit subsidy that the government provides through the design of the loan system.

7. Conclusions and Further Work

This report has estimated the returns to HE and extended the analysis in previous research in a number of directions. In the course of this it has examined the robustness of the results reported in LE in a number of ways. The important extension that we implement in estimation is to allow much greater flexibility in the specification of the statistical model - which permits us to distinguish between lifecycle, cohort and time effects on hourly earnings. We find that this flexibility has important implications.

While we were concerned that the expansion of the supply of graduates in the 1990s may have reduced the graduate – non-graduate earnings differentials, we were not able to find statistically significant differences between post and pre expansion cohorts. It could be that any effect was masked by the coincident recession that may have affected non-graduate earnings more than graduate earnings.⁵⁴ Further progress in this dimension is likely to be difficult with available data.

Our estimation methodology also allows the lifecycle of wages to differ between graduates and non-graduates. The differences for men were very small and, while they were larger for women they were still not statistically significantly different.

The important difference we found was that there was higher secular growth in the real earnings for graduates than for non-graduates over *time* – consistent with the idea that technical progress over time has been biased towards skilled labour, as has been reported in many other studies. Even though the annual difference was small, projecting this forward across a whole lifecycle has important consequences for the NPV of the graduate earnings differential. Thus, we find larger NPVs of both gross and net private benefits, and the return to the government using our estimates.

An important concern that proved to be unfounded was the possibility of ability bias – arising from the correlation between educational attainment and unobserved ability. The fear was that our estimates would be biased upwards. We found that our estimates, using richer data in BHPS, were not affected by the inclusion of controls that might capture such bias. We were also able, using the BHPS data, to establish that the effects of HEI type on the degree premium were relatively modest when controls for ability were included – even though those controls were not likely to eliminate all of the effect of unobservable ability differences. While we estimate the effects of a degree, on average, we also break down this average by subject. Here we are less optimistic that our estimates (and similar ones in LE) can be relied on. We were not able to subject these estimates to the same checks because BHPS does not contain information on subject. It seems likely that ability differs by subject so that we are likely to overestimate the effect of some degrees more than others.⁵⁵

⁵⁴ The paucity of data on youth unemployment in the past prevents us from pursuing this further. Further research may do well to exploit regional variation in unemployment over a shorter period of time using available youth unemployment data. And there may be some merit in using benefit count unemployment data to delve further back in time.

⁵⁵ It is difficult to see how this could be resolved with existing data. One possibility would be to merge data from HESA and NPD administrative records into LFS to provide a richer set of controls in LFS – but this is not presently possible.

In simulation, the important extension we pursue is to allow for the *variance* in predicted earnings. That is, we do not simulate using a single point prediction but, instead, simulate the whole estimated distribution in earnings. This has important implications for *net* private benefits since it allows us to capture the higher subsidy implied by the 30 year debt forgiveness limit and the £21,000 threshold in the new loan arrangements. This arises because we capture the subsidy received by the bottom tail of the earnings distribution. It also has important implications for government revenue since we capture the higher rate tax revenue associated with the top tail of the earnings distribution. **HE is an important and favourable investment for the government as well as for students.**

Both of these innovations combine to suggest that the simulated private benefits and the simulated revenue consequences of higher education are considerably underestimated in previous research. We have some evidence that suggests that these average effects are robust to unobserved differences across graduates and non-graduates.

However, the averages conceal considerable variation across students. Much attention has been given to the merit of studying specific subjects. Unfortunately, there are no data that allow us to examine the extent to which these results might be due to selection effects – the subjects with the higher returns may simply be reflect selection of better students into those subjects. Thus, no causal interpretation can be given to our results decomposed by subject.

Appendix – Further Results

Table A1 **Bounds on the returns to degrees (baseline specification)**

Relative correlation coefficient (Λ)	Men	Women
{0}: (OLS results)	0.205	0.268
(s.e.)	(0.004)	(0.004)
{0, 0.25}	[0.172, 0.205] (0.163, 0.213)	[0.250, 0.268] (0.242, 0.276)
{0, 0.5}	[0.138, 0.205] (0.128, 0.213)	[0.230, 0.268] (0.221, 0.276)
{0, 1}	[0.068, 0.205] (0.055, 0.213)	[0.188, 0.268] (0.174, 0.276)

Note: Λ = ratio of amount of selection on unobservables to the amount of selection on observables. Intervals in square brackets are the bounds themselves, while the intervals in the round brackets are 95% asymptotic confidence intervals.

{0, 0.25} means that 0 is the lower bound of the causal effect of a degree (i.e. where it is assumed that there is no selection on unobservables – as in OLS), and 0.25, for example, is the assumed maximum amount of selection on unobservables (correlation with the error term) relative to the selection on observable.

So 0.172 is the lower bound of the earnings premium if one feels that there is only a modest degree of selection on unobservables. The table suggests that it takes a large degree of selection on unobservables to drive the lower bound of the estimate of the causal effect of a degree to below 10%. The results suggest that ability bias is small enough to make us reasonably confident that the return to a degree is, indeed, high. For women, the bounds are particularly tight.

Table A2 shows that the predicted employment probabilities by gender and subject of degree. Consistent with the observed data, having a degree is associated with higher probability of being in employment, especially for women. There is some variation across degree subjects among graduates; however this is small relative to the difference between graduates and non-graduates, regardless of gender.

The QR results show, as expected, the estimated effect of a degree at the median is close to OLS. The estimated returns for women at the median are similar to that the bottom of the distribution but smaller at the top decile. For men the results at the top of the distribution are smaller than at the median - but still large enough to make a degree a strong investment, on average. The results at the bottom decile are even larger than at the median. To the extent that these QR results capture differences in ability across the distribution of graduates they are encouragingly close to the OLS results that we rely on in our simulation analysis.

Table A2 Employment probability estimates including subject of degree: LFS (baseline specification)

	Men	Women
<i>2+ A-Level</i>	<i>0.741</i>	<i>0.651</i>
Medical related	0.840	0.827
Nursing	0.855	0.817
Bio/Vet/Agric	0.832	0.808
Physical Sci	0.831	0.824
Maths/Comp	0.854	0.832
Engineer/Tech	0.835	0.837
Architecture	0.824	0.813
Social Study	0.840	0.812
Law	0.821	0.842
Economics	0.848	0.823
Business & Man	0.840	0.836
Mass Com	0.860	0.818
Ling/Lang	0.825	0.810
Hist/Phil	0.805	0.797
Arts/Design	0.841	0.824
Education	0.838	0.809
Combined	0.842	0.824

Table A3a Quantile Regression Estimates: LFS (baseline specification), Men

Variable	OLS	Bottom decile	Median	Top decile
Degree	0.20482***	0.23476***	0.20955***	0.16960***
Age	0.12402***	0.12356***	0.12493***	0.13780***
Age ²	-0.00134***	-0.00140***	-0.00133***	-0.00145***
Immigrant	0.00837	-0.03307	0.00762	0.01389
Nonwhite	-0.06424***	-0.08880***	-0.06434***	-0.03151*
London	0.23872***	0.20811***	0.22834***	0.33193***
SE	0.17217***	0.14723***	0.16821***	0.23056***
Wales	-0.05362***	-0.04561**	-0.06135***	-0.06374***
Scotland	0.04331***	0.03994***	0.03891***	0.03292***
NI	-0.07731***	-0.05852**	-0.06563***	-0.12932***
Time	0.00431***	0.00258***	0.00365***	0.00721***
Constant	-0.03008	-0.43556***	-0.05371**	0.11587***
N	61122	61122	61122	61122
R ²	0.29827			

Table A3b Quantile Regression Estimates: LFS (baseline specification), Women

Variable	OLS	Bottom decile	Median	Top decile
Degree	0.26846***	0.27695***	0.27761***	0.21706***
Age	0.08611***	0.06242***	0.09648***	0.11161***
Age ²	-0.00094***	-0.00072***	-0.00106***	-0.00121***
Immigrant	-0.01564	-0.00164	-0.01693	-0.02683
Nonwhite	-0.04774***	-0.05911***	-0.05500***	-0.05295***
London	0.23689***	0.31104***	0.22086***	0.24229***
SE	0.07717***	0.05153***	0.07616***	0.10827***
Wales	0.00367	0.01758	0.00523	-0.01830
Scotland	0.02621***	0.04362***	0.03127***	-0.00104
NI	-0.01299	0.01823	-0.00948	-0.07578***
Time	0.00382***	0.00518***	0.00310***	0.00499***
Constant	0.58976***	0.55389***	0.41623***	0.54508***
N	58799	58799	58799	58799
R ²	0.22722			

References

- Altonji, J., T.E. Elder, and C.R. Taber (2005), "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools", **Journal of Political Economy**, 113, 151-184.
- Ashenfelter, O., C. Harmon, and H. Oosterbeek (1999), "Meta-analysis of the returns to education", **Labour Economics**, 6, 453-470.
- Bonjour, D., Cherkas, L.F., Haskel, J.E., Hawkes, D.D. and T.D. Spector (2003), "[Returns to Education: Evidence from U.K. Twins](#)", **American Economic Review**, 93, 1799-1812).
- Bound, J. and S. Tanner (2012), "Dropouts and Diplomas: The Divergence in College Outcomes", in E. Hanushek and S.J. Machin (eds), **Handbook of the Economics of Education**, Volume 4, North-Holland.
- Blundell, R.W., L. Dearden and B. Sianesi, (2005), "Evaluating the impact of education on earnings: Models, methods and results from the NCDS", **Journal of the Royal Statistical Society Series A**, 168, 473-512.
- Callender, C., D. Wilkinson, A. Gibson and C. Perkins (2011), "The impact of higher education for part-time students", UK Commission for Employment and Skills.
- Card, D. (2000), "The causal effect of education on earnings", in O. Ashenfelter and D. Card (eds.) **Handbook of Labour Economics** Vol 3a, North Holland.
- Chevalier, A., (2009), 'Does Higher Education Quality Matter in the UK?' Royal Holloway, University of London, **mimeo**.
- Conlon, G., and A. Chevalier, (2003), 'Does it Pay to Attend a Prestigious University?' **Centre for the Economics of Education (CEE) Discussion Paper P0033**.
- Conlon, G. and P. Patrignani (2011), "The Returns to Higher Education Qualifications", **BIS Research Paper 45**.
- Dolton, P.J., G.D. Inchley and G.H. Makepeace (1990), "The early careers of 1980 graduates: earnings, earnings differentials and postgraduate study", **UK Department of Employment** Research Paper No. 78.
- Eide, E., D. Brewer and R.G. Ehrenberg (1998). "Does It Pay to Attend an Elite Private College? Evidence on the Effects of Undergraduate College Quality on Graduate School Attendance", **Economics of Education Review**, 17, 371-376.
- Harmon, C. and I. Walker (1995), "Estimates of the economic return to schooling for the UK", **American Economic Review**, 85, 1278-1286.
- Harmon, C., I. Walker, and N. Westergaard-Nielsen (2001), **Education and Earnings in Europe**, Edward Elgar Press, Colchester.
- Hoekstra, M. (2009), The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach, **Review of Economics and Statistics**, 91, 717-724
- Hussain, I., S. McNally, and S. Telhaj, (2009), "University Quality and Graduate Wages in the UK", **IZA WP 4043**.
- Kahn, L. (2010), "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy", **Labour Economics**, 17, 303-316 .

- Krauth (2011), "Bounding a linear causal effect using relative correlation restrictions", Simon Fraser University, **mimeo**.
- Lindley, J. and S. Machin (2011), "Rising Wage Inequality and Postgraduate Education", **CEP Discussion Paper** 1075.
- O'Leary, N.C. and P.J. Sloane (2004), "The Return to a University Education in Great Britain", **IZA Working Paper** 1199.
- OECD (2010), **Education at a Glance**, OECD.
- Oreopoulos, P., T. von Wachter, and A. Heisz (2006), "Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates", **NBER Working Paper** 12159.
- Sloane, P.J. and N.C. O'Leary (2005), "The Return to a University Education in Great Britain", **National Institute Economic Review**, 193, 75-89
- Walker, I and Y. Zhu (2008), "The College Wage Premium and the Expansion of Higher Education in the UK", **Scandinavian Journal of Economics**, 110, 695-709.
- Walker, I and Y. Zhu (2011), "Differences by Degree: Evidence of the Net Financial Return to Undergraduate Study in England and Wales", **Economics of Education Review**, 30, 1177-118.

© Crown copyright 2013

You may re-use this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence. Visit www.nationalarchives.gov.uk/doc/open-government-licence, write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk.

This publication available from www.gov.uk/bis

Any enquiries regarding this publication should be sent to:

Department for Business, Innovation and Skills
1 Victoria Street
London SW1H 0ET
Tel: 020 7215 5000

If you require this publication in an alternative format, email enquiries@bis.gsi.gov.uk, or call 020 7215 5000.

BIS/13/899