The Role of Uncertainty in Education Mismatch

Shihang Hou*

October 5, 2022

Abstract

In many developed economies, there is both a high share of students selecting into higher education (HE) and a large share of graduates who end up working in non-graduate occupations. Some policymakers, especially in the British context, have pointed to this fact as evidence that the level of education is inefficiently high. In this paper, I propose an explanation of overeducation based on the observation that the causal effect of HE on human capital is both heterogeneous and imperfectly observed by the student. I propose a model in which workers first decide whether to invest in higher education on the basis of limited information, and then match to jobs in a labour market. In this setting, workers could be over-educated if they do not benefit from HE after attending, and under-educated if they could have benefited but did not attend HE. Structurally estimating my model on UK data, I calculate that 18.2% of the population are over-educated and 14.7% are under-educated. In general, the equilibrium in the model is inefficient, because of a hold-up externality which causes under-investment and a congestion externality which causes over-investment. These externalities offset each other, and in general, it is ambiguous whether the level of education is inefficiently high or low.

JEL Codes: D81, I23, I26, J24
Keywords: Overeducation, efficient human capital investment, labour market matching

*University of Oxford, Department of Economics. Address: 10 Manor Road, Oxford OX1 3UQ, United Kingdom. E-mail: shihang.hou@economics.ox.ac.uk. I would like to thank my supervisor, Margaret Stevens, for her guidance throughout the writing of this paper. I would like to thank Abigail Adams-Prassl, Ian Crawford, Felicia Ionescu, Thomas Jorgensen, Romuald Meango, Volker Nocke, Jessica Pan, Barbara Petrongolo, Simon Quinn, Kjetil Storesletten, Alex Teytelboym, Greg Veramendi, and Basit Zafar for their helpful comments on various versions of the paper. I would also like to thank participants at the Warwick PhD Conference, the Surrey PhD Workshop in Microeconomics, and the Young Economist Symposium 2022 for their thoughts and comments. Finally, I would like to thank my fellow DPhil economics students at Oxford, particularly Luke Heath Milsom, Binta Zahra Diop, Hannah Zillessen, Nick Ridpath and Vatsal Khandelwal, for helpful discussions and moral support.

1 Introduction

In many developed countries, a large share of young people attend higher education and complete undergraduate degrees[1]. However, in many of these countries, a significant share of graduates end up working in occupations which typically do not require a college degree[2] (hereafter, I use the terms over-education and under-employment interchangeably to refer to this phenomenon). Despite the college-wage premium remaining large, the large shares of students selecting into higher education (hereafter HE) and subsequently working in non-graduate occupations has led to concerns that many students are inefficiently over-investing in higher education[3]. These concerns typically come with the suggestion that it may be optimal to lower the numbers of student selecting into HE to reduce the incidence of overeducation.

Most previous theories[4] about the existence of over-education have focused on the role of labour market frictions, which may make it more profitable for skilled workers to take jobs in non-skilled occupations than remain unemployed. This paper proposes an alternative explanation of under-employment based on the idea that the causal effect of HE on human capital is heterogeneous. If that were the case, then students with a low causal return only slightly improves their human capital and may nevertheless only match with a non-graduate job post-HE. Students with low returns may nevertheless select into higher education if they are not perfectly aware of their return, and make ex-ante decisions that can turn out to be ex-post mistakes. This explanation coheres with recent work including Carneiro et al. (2011), Carneiro & Lee (2011) and Heckman et al. (2018), which emphasise that returns to HE vary by ability and that there are compositional changes in the population of graduates when HE expands.

In fact, if it is possible to abstract from general equilibrium effects, then the marginal treatment effects methodology used in Carneiro et al. (2011) and discussed in Heckman & Vytlacil (2001) could be used to characterise the proportions of students who invest unwisely in higher education. However, abstracting from general equilibrium effects is typically not reasonable especially when

[1]Statistics in the US context have been presented in Jackson (2021). I present details in the UK context which is the focus for my empirical work in the upcoming section 2.

[2]This phenomenon is described as over-education or under-employment in the economics literature. Leuven & Oosterbeek (2011) provides a critical review of the literature on overeducation more generally. Clark et al. (2017) and Meroni & Vera-Toscano (2017) document that over-education seems to be a persistent phenomenon over individuals lifetimes.

[3]For example, in the UK context, the then-Universities Minister summarised this concern as follows:

> For decades we have been recruiting too many young people on to courses that do nothing to improve their life chances or help with their career goals....Since 2004, there has been too much focus on getting students through the door, and not enough focus on how many drop out, or how many go on to graduate jobs.

In 2022, in the Tory leadership contest, ex-Chancellor of the Exchequer Rishi Sunak similarly proposed to use the shares of graduates in graduate jobs as an indicator of which degrees were good degrees. The quote comes from a speech on 1 July 2020 at the NEON summit on widening access and mobility by then-Universities Minister, Michelle Donelan. The transcript was accessed online on 30 May 2022 at https://www.gov.uk/government/speeches/universities-minister-calls-for-true-social-mobility.

[4]An early treatment of overeducation was Albrecht & Vroman (2002), which develops a matching model with heterogeneous workers and jobs, and showed that underemployment could exist in an economy due to search frictions. High-skill workers may accept matches with low-skill jobs because the cost of search is too high. Jackson (2021) further argues that while in a similar kind of model, the level of underemployment does not imply anything about the optimal level of higher education. In his model, although subsidies to HE decrease the returns to HE and increase the underemployment rate, that may nevertheless be efficient because it increases overall production and reduces the rate of unemployment.

considering counterfactuals involving large changes to the share of students choosing HE. It is likely that when the supply of graduates changes, the magnitude of the wage return of HE also could change, due to changes in unit price of 'graduate' labour relative to that of 'non-graduate' labour[5].

I develop a model where the return to HE in terms of human capital is heterogeneous across workers, varying by an index which I call ability. This return is imperfectly observed at the point of choice; workers can only observe their grades in school, which I interpret as noisy signals of their underlying ability. As such, workers make an education decision based on an ex-ante expectation of their likely returns to education, but ex-post may experience regret if they are 'unlucky', i.e. they received a signal which convinced them that they were higher ability than they actually are. Post HE, workers match to jobs in a frictionless matching market based on their post-education human capital. This matching labour market generates general equilibrium price effects; when the share of workers in HE increases, workers have less bargaining power as there are more skilled workers in the economy and their wages conditional on their human capital fall as a result.

I emphasise two implications of my model. First, over-education is just one kind of what may be more generally seen as education mismatch, which is when a worker makes an education choice that is worse off for them ex-post. In equilibrium, there are both over-educated graduates and under-educated non-graduates ex-post, and I find in my empirical estimation that under-education typically leads to a greater welfare loss than over-education. The prevalence and cost of the ex-post education mismatch depend crucially on how well workers can predict their return to HE. Second, in general, the equilibrium of the model is not efficient even with ex-ante expected utility maximising agents because of two externalities. First, there is a hold-up externality where workers do not internalise the effect of their education on firms' profits because they cannot appropriate this surplus. Thus, they will under-invest in higher education. Second, workers do not consider the impact of their education choice on the bargaining power of other workers in the economy. This leads to over-investment in higher education relative to the social optimum. Because these externalities go in opposite directions, it is ambiguous whether the level of education is too high in the economy. In general, there is scope for welfare-improving policy.

I contribute empirical estimates of the extent and welfare costs of education mismatch by estimating a parametric version of my model on data from the UK using the Simulated Method of Moments (SMM). I find that students face substantial uncertainty about their eventual returns to university; the estimated correlation between unobserved ability and the signal is only 0.324. I estimate that 32.9% of workers would have been better off if they made a different education decision to the one they actually did. 18.2% of workers or 39.1% of graduates did not benefit from university despite going, and 14.7% of all workers or 27.6% of non-graduates would have benefited from university despite not going. The costs in utility terms are small on average but non-negligible. The average utility loss to going to university for overeducated workers is equivalent to earnings increases of £0.183 per hour on average, while the average utility loss to missing out on university for undereducated workers is equivalent to £0.264 per hour on average. These results suggest that while over-education is the focus of headlines in the education debate, under-education is a substantially more costly problem when it occurs, which is why workers tend to err on the side of over-investment. I find that in practice, it seems to be welfare-improving to reduce the level of education by 1.6pp by implementing a small revenue-neutral graduate tax.

[5]This idea was the basis of the seminal Katz & Murphy (1992) which led to later work on the race between education and technology in Katz & Goldin (2008). This argument is also made in the search and matching context by Acemoglu (1999) and Jackson (2021).

This paper proceeds as follows. I conclude the introduction by discussing the literature that my paper relates to. I begin by describing some stylised facts about UK higher education in section 2 that my model seeks to explain. I then describe my model in section 3. In section 4, I describe the empirical parameterisation of my model, identification and the data sources which I use. Section 5 describes the results of my estimation, the model fit, and analyses welfare within the model. Section 6 analyses worker welfare in this model given the estimated parameters. I consider whether policy makers can increase welfare using a compensated graduate tax or subsidy in section 7, relating the exercise to current policy proposals being considered by the UK government. Section 8 concludes.

1.1 Literature Review

In terms of subject matter, my paper relates to a few papers which have proposed mechanisms for why there might be over-education in the economy. Albrecht & Vroman (2002) first noted that labour market frictions could induce skilled workers to accept matches with low-skilled jobs. Jackson (2021) endogenises the education choice in the context of a labour market sharing some characteristics with that of Albrecht and Vroman, and characterises optimal policy in this context. Gottschalk & Hansen (2003) proposes a mechanism where some graduates have heterogeneous preferences for non-skilled occupations. I contribute to this literature by proposing a new channel based on the heterogeneity of possible returns. Relative to existing explanations, my mechanism can account not just for over-education, but also under-education, as in the data, there are typically some non-graduates who are observed to work in high-skilled occupations. My explanation also implies a new channel for policy intervention - improving the information that students have about their likely returns to higher education. Relative to these papers, I also can allow for more than two firm types (i.e. simple or complex, high-skilled or low-skilled) by modelling the distribution of job productivities as a mixture of occupation-specific job productivity densities. This is to my knowledge novel to the literature and allows me to make predictions about the degree composition of occupations.

My model is similar to that of Chade & Lindenlaub (2021), which also consists of an education decision which subsequently feeds into a matching model of the labour market. The education choice component resembles that studied in e.g. Cunha et al. (2005), while the labour market component was first studied in the labour market context by Sattinger (1993). Relative to their paper, I allow for the return to education to depend on a factor which correlates to cognitive ability. This allows the model to generate a dependence between signals of ability and the return to education, which have been found in Heckman et al. (2018). This generates a relation between these measures and the probability of attending university, which was accounted for in their paper entirely through a mechanism invoking psychic costs. There are also other papers that propose models with similar features to that in my paper. Guvenen et al. (2020) use an information revelation mechanism similar to the one in this paper to explain occupational mismatch. Shephard & Sidibe (2019) and Jackson (2021) both analyse frictional labour markets with endogenous binary education choices. Besides the incorporation of labour market frictions, my paper differs from theirs in that I incorporate uncertainty about returns to education. Macera & Tsujiyama (2020) use a model with a similar structure to analyse the impact of changes in the firm productivity distribution on wage inequality.

My paper relates to the many papers which analyse student decisions under uncertainty. The idea of uncertainty in wage returns and ex-post regret is not new. Altonji (1993) develops a theoretical framework that explicitly considers various sources of uncertainty faced by the student

in their decisions to invest in education. Cunha et al. (2005) discuss the identification of Roy models with both student heterogeneity known at the time of investment and uncertainty not forecast-able by the student. Arcidiacono et al. (2016) analyse student learning throughout college, where students gradually learn about their abilities through their performance in college classes. My contribution here is that my paper is the first to offer a framework relating ex-post education regret to observations of graduates in non-graduate occupations, thus drawing together the occupation-based over-education literature and the wage-based causal effects of higher education literature. This thus relates the concept of ex-post regret to observed education-occupation mismatches. My framework also allows me to quantify the extent of mismatch in the economy, and to analyse the impact of policy, and to structurally relate these conclusions to the amount of uncertainty workers face.

Finally, my paper relates to the literature on the efficiency of higher education. Charlot & Decreuse (2005) argue that in an economy with segmented labour markets and heterogeneous workers, low-ability workers over-invest in education because they do not internalise the effect of their education choice on the average productivity of graduates. Hopkins (2012) analyses a matching model where workers can signal their ability by investing in education and finds that equilibrium is not generally efficient. In his model, low-ability workers can under-invest and high-ability workers can over-invest relative to equilibrium. Finally, Jackson (2021) finds that on top of the hold-up and congestion externalities in my model, there may be a thick markets externality where more education investment results in the creation of more complex jobs. His calibration finds that in the US context, there is under-investment in education and an inefficiently low level of underemployment. Empirically, my paper suggests that in the UK context, there is slight over-investment in education.

2 Facts about the labour market for graduates in the UK

2.1 The higher education wage premium

In this section, I review some of the patterns and stylised facts that motivated the model described in the following section. I focus on the UK context[6], although I note that Jackson (2021) has documented similar patterns in the American context.

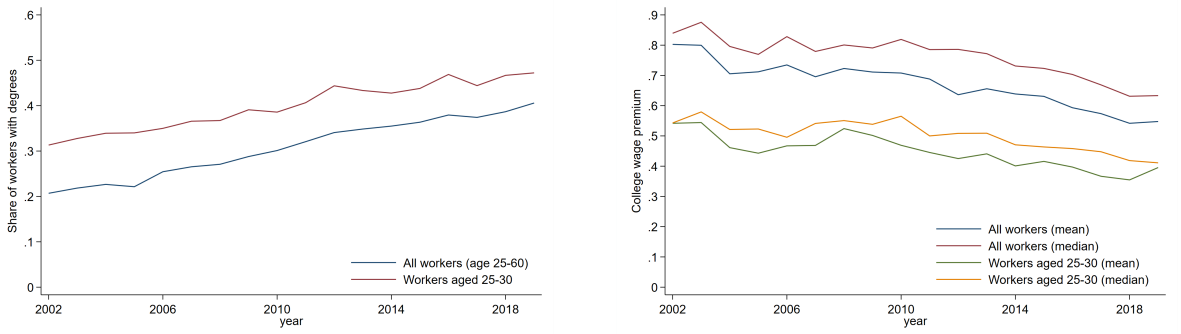
I focus on the replication of two well-known stylised facts about the UK graduate labour market. First, there has been a large and sustained increase in higher education in the UK over the last two decades[7]. Second, the previous papers in the literature have typically found that the wage premium has stayed roughly constant, despite the large expansion in graduate numbers in the labour market.

First, I plot the share of workers with a degree or equivalent from years 2002-2019 in table 1a. The first striking fact is that the share of workers with a higher education degree has increased dramatically over the last 17 years. In 2002, the share of 25 to 30 year olds with a graduate degree was 31.3%. Since then, there has been a roughly linear upwards trend, such that in 2019, 47.2%

[6]See Blundell et al. (2018) for a more detailed description of the history of UK higher education and the many policy changes that have led to the current situation.

[7]Other papers like Walker & Zhu (2008), Devereux & Fan (2011), and Blundell et al. (2018) also discuss the increase in the college participation rate and the college wage premium in the UK context.

Figure 1: Changes in the share of workers with degrees and the college wage premium from 2002-19



(a) Share of workers with degrees

(b) College wage premium

Data from the UK Labour Force Survey 2002-19. The mean college wage premium for each year as the coefficient on the indicator for whether the individual has a degree of a regression of log hourly earnings on dummies for 5-year age groups and the degree indicator. The median college wage premium is defined as the coefficient on the indicator for whether the individual has a degree of a quantile regression of log hourly earnings on dummies for 5-year age groups and the degree indicator.

of the same demographic group have a graduate degree, constituting an increase in share of about 15.9 percentage points. Since the population of 25 to 30 year olds has increased by about 45%, the number of graduates aged 25-30 years old in the labour market has increased by about 120% in the last 17 years.

Given such a large increase in the number of graduates, the supply and demand perspective of the graduate labour market would suggest that the wages of graduate workers should fall. This should in turn decrease the wage premium (defined as the ratio of the mean earnings of college graduates to non-college graduates). Some papers (e.g. Carneiro & Lee (2011)) further argue that when more students select into education, the marginal student is likely to be of lower quality and thus should lower the wage premium by changing the quality composition of the groups of graduates and non-graduates. To check this, figure 1b plots the mean and median college wage premium for the group of 25-30 year olds and for all workers.

Much of the older literature on the UK higher education premium has tended to conclude that the higher education expansion that started in the 90s did not substantially reduce the observed college wage premium. This seemed to be true from 2002 to 2010, when both the mean and median premium for workers aged 25-30 fluctuated around 50-55%. After this period however, there seems to be a clear downwards trend, such that the median premium is about 41% in 2019 while the mean premium has fallen to about 40%. Despite the fall, the observed wage premium still remains economically large, especially given the massive increase in the share of workers with degrees documented in figure 1a. These results seem to suggest that while the initial explosion in the number of graduates in the 90s did not seem to affect their labour market fortunes, subsequent cohorts of graduates in the 2000s seem to experience smaller wage premia relative to earlier cohorts. Nevertheless, the decline of the premium, by roughly 10-15 percentage points, is small relative to massive expansion in both the share and numbers of students now in education.

The calculated wage premium is also not an estimate of the causal wage return to education,

since it does not take into account selection into education and other forms of endogeneity. Since estimating a causal return to education credibly is a monumental task, I do not attempt it in this descriptive section. There have been three main papers estimating the return to higher education in a British context. All these papers estimate Mincer equations, accounting for heterogeneity using a matching approach; they take advantage of detailed observations of schooling and family variables to argue that any residual selection effects do not substantially affect the estimate. Blundell et al. (2000) use administrative data linked to the British Household Panel Survey and calculate that returns to a degree is 17% for men and 37% for women for a cohort born in 1958. Belfield et al. (2018a) similarly uses administrative data to find that the return at age 29 to be 8% for men and 28% for women for cohorts who attended university in the mid-late 2000s. Waltmann et al. (2020) simulate workers' life-cycle earnings for these recent cohorts and conclude that the return increases over the life-cycle, with the overall causal return being about 20% for both men and women, after taking into account loan repayments and the progressivity of the tax system. The fall in the headline return to HE from that calculated by Blundell et al. (2000) to that calculated in Belfield et al. (2018a) suggests that like the simple wage premium, the return to higher education has fallen over time.

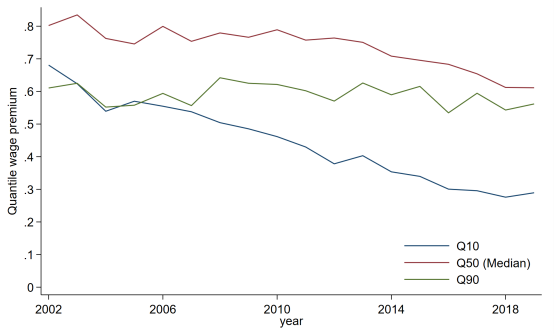
2.1.1 The percentiles and variances of the log wage distributions conditional on education

Many recent studies on the causal impact of higher education has highlighted the heterogeneity of the effect higher education has on different students[8]. Thus, it is possible that the mean or median wage premia may mask different and divergent trends for graduates at the tails of the wage distribution. To check this, I run a number of quantile regressions at the 10th percentile and 90th percentile to study how the wage premia between the 10th percentile graduates and non-graduates have changed. In figure 2a, I plot the coefficients on an indicator variable for whether the worker has a degree of quantile regressions on log hourly wage, including dummies for 5-year age groups as a control, for the 10th, 50th and 90th percentiles. As the blue line (for the 10th percentile quantile regression) shows, the difference of the conditional 10th percentile log wage between graduates and non-graduates has fallen substantially over the years.

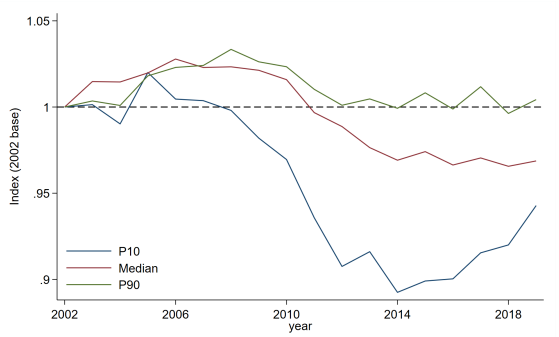
To check why this might be the case, I plot the indices of the log-wage of graduates at the 10th, 50th, and 90th percentiles, with 2002 as the base year (see 2b). A significant part of this story seems to be the substantial fall of the real wage for the 10th percentile of graduates, which at the trough in 2014, was more than 10% lower than in 2002. Thus, any theoretical model should ideally explain why the wages at the 10th percentile have fallen over the course of the substantial expansion of higher education.

[8] A related line of inquiry has been whether there are negative returns to higher education in the UK context. Belfield et al. (2018a) and Waltmann et al. (2020) find that for some subgroups analysed (students at particular institutions studying particular subjects), the return to education is negative. In the US context, Heckman et al. (2018) finds that returns for students with low test scores are near zero, but stops short of finding evidence of negative returns.

Figure 2: Changes in the distribution of earnings conditional on degree status



(a) Coefficients on quantile regressions at p10, p50 and p90



(b) Log wages for graduates at p10, p50 and p90

Data from the UK Labour Force Survey 2002-19. Controls for the quantile regressions in figure 2a include age group dummies.

2.2 Distribution of graduate employment into different occupations

2.2.1 The share of graduates in low skill occupations is increasing

A third stylised fact that is often found in the UK context is that there is a significant and increasing share of workers with graduate degrees working in so-called 'non-graduate' jobs[9]. Given the policy interest in the subject, especially since the government implicitly subsidises graduates who cannot fully repay their student loans, the UK Office of National Statistics has published a number of technical reports attempting to quantify overeducation using a variety of methods[10].

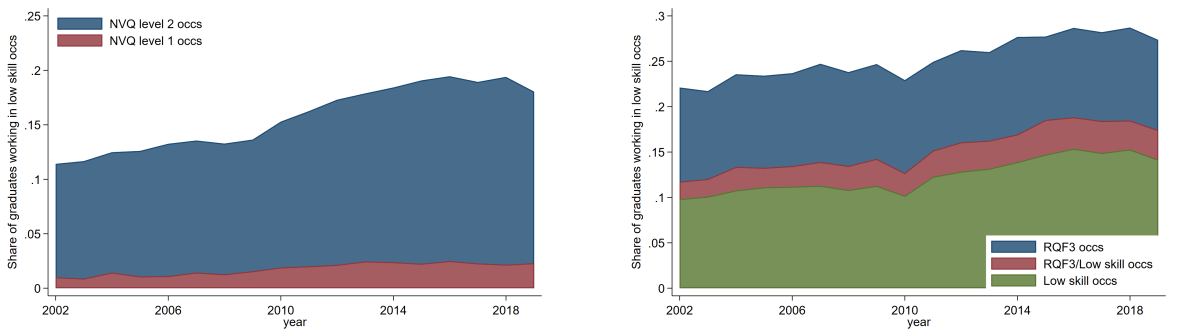
In this section, I first show that there are more graduates in occupations classified by the government statistical agencies as non-skilled occupations. The narrative that emerges from this exercise seems consistent; there are more graduates today working in occupations that would traditionally be considered 'non-graduate'. Figure 3a plots the share of graduates working in occupations considered to require a NVQ 1 or 2 level education (equivalent to having a good high school qualification), as classified by the ONS occupational classification. Over the studied time period, the share of workers working in low skill occupations increased from 11.4% in 2002 to a high of 19.4% in 2018. Figure 3b plots a similar graph, but instead uses a classification used by the UK Home

[9] As discussed in the literature review, the demarcation between graduate and non-graduate jobs is controversial, and a main concern of the large overeducation literature has been the clarification of the demarcation between the two categories. There are two main approaches to this issue. First, some researchers, including Green & Zhu (2010), and Green and Henseke (2016), argue that a classification should be based on some notion of an occupation requiring certain skills. Approaches which make use of existing skills classifications, such as the Office for National Statistics SOC classifications or the Home Office Appendix J classifications for occupation skill levels, would fall under this approach, as do approaches which turn to the task content of occupations to stipulate whether a job is 'graduate'. Second, other researchers have tried to classify occupations into 'graduate' and 'non-graduate' categories by examining outcomes, either the college-noncollege wage premium within an occupation (Gottschalk & Hansen (2003), O'Leary & Sloane (2016)) or the share of workers with degrees within that occupation (Elias and Purcell (2003)).

[10] See Clegg (2017) and Saric (2019). Academic work in a UK context attempting to quantify overeducation include Elias & Purcell (2004), Green & Zhu (2010), and O'Leary & Sloane (2016).

Office to classify occupations into different skill levels[11]. Occupations were classified into low skill (requiring less than an A-level qualification), requiring an A-level qualification (RQF 3), requiring one year of a bachelor's degree (RQF 4), requiring an undergraduate degree (RQF 6), and requiring a postgraduate degree. At the start of the period, 9.8% of graduate workers worked in an occupation classified as low skill, but at the end, 14.2% of graduate workers do. 22.1% of workers were in an occupation requiring a high-school qualification or below at the start of the period, and at the end, 27.4% did. Thus, on the face of it, it seems that more graduate workers are now employed in so-called low-skill qualifications requiring a high-school degree or lower. Appendix A.1 shows that while the share of graduates in so-called low-skill occupations do vary between different subjects, non-negligible shares of graduates in most subjects, both in STEM and non-STEM, work in low-skilled occupations[12].

Figure 3: Share of graduates in low-skill jobs by two different classifications



(a) Share of graduates working in occupations requiring NVQ1/2 qualifications

(b) Share of graduates in low skill occupations by Home Office classification

Data from the UK Labour Force Survey 2002-19. The occupational classifications produced by the Office for National Statistics are assigned skill levels relating the qualification level (based on the old NVQ classification) required for the job. The appendix J classifications are produced by the Home Office for the then-purpose of classifying prospective migrants by their skill level according to the Regulated Qualifications Framework.

The interpretation of this fact is however not straightforward. As the number of graduates increase, the market readjusts[13], and jobs which were previously seen as non-graduate may be increasing taken up by graduates. This is in fact a feature of the model proposed in the following section, that as the supply of graduates changes, the composition of graduates within particular occupations similarly change.

[11]This classification has also been used in Aghion et al. (2019)

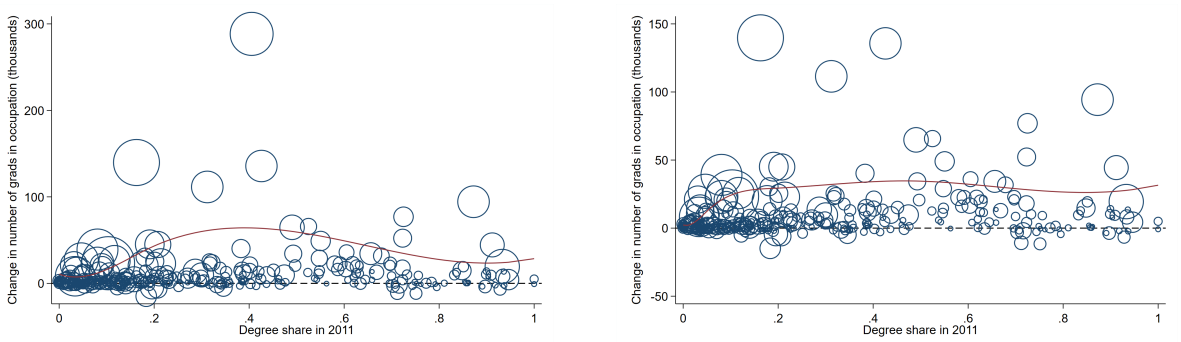
[12]This does not apply to degrees for vocational subjects, such as medicine, subjects allied to medicine, architecture, and to a lesser extent, education.

[13]Models in which the composition of jobs changes as a result of technological change or an increase in the number of graduates include Acemoglu (1999), Albrecht & Vroman (2002), and Gottschalk & Hansen (2003).

2.2.2 New graduates are increasingly working in occupations with previously low shares of degree holders

A natural question given the patterns in figure 1a showing a substantial increase in the number of graduates is what occupations these new graduates found employment in. To this end, I use information on the moments of hourly earnings within each occupation from the Annual Survey of Hours and Earnings (ASHE), matched to information about the share of workers with degrees within the occupation. Figure 4a plots a scatter diagram of the occupations, with the share of workers within the occupation with degrees in 2011 (the start of the available period[14]) on the x-axis, and the absolute change in the number of graduates (in thousands) on the y-axis[15]. A non-parametric line of best fit, weighted by the number of workers in that occupation in 2011, is plotted over the points. Figure 4b plots a similar scatterplot, but excludes the occupation with the largest change in figure 4a, nursing[16]. The size of the circles in both diagrams shows the size of the occupation in 2011.

Figure 4: Change in number of graduates within each occupation (2011-19)



(a) Change in number of graduates working in occ

(b) Figure 4a (excluding nurses)

Data on the share of workers with degrees from the UK Labour Force Survey 2011-19. Data on the number of workers in the occupation from the Annual Survey of Hours and Earnings. A non-parametric line of best fit is drawn in red in both diagrams. The size of the bubbles reflects the size of the occupations.

In each figure, I plot a non-parametric line of best fit for the relation between the initial share of workers in the occupation with degrees in 2011 and the change in the number of graduates in the occupation between 2011 and 2019. A general pattern that arises from these figures is that the occupations which received the greatest influx of workers were typically occupations with a fairly medium levels of workers with degrees at the start of the period[17].

[14]To maintain consistent occupational definitions, I used a period, 2011-19, for which SOC10 classifications were available.

[15]The absolute number of graduates was calculated by multiplying the number of workers in the occupation from the ASHE by the share of workers with degrees estimated from the LFS.

[16]From 2008, the minimum award for completing a nursing qualification became a degree, and so, new nursing graduates in subsequent cohorts have substantially increased the share of workers with degrees in the nursing profession. Since this change in classification arguably does not reflect any real economic change yet drives the roughly quadratic shape in figure 4a, it is removed as a robustness check in figure 4b.

[17]The occupation seeing the largest change is nursing, with 288 thousand new graduates entering the profession; this is likely to be a mechanical affair driven by a change in the classification of the nursing qualification. The occupations absorbing the next most graduates are respective "other administrative occupations" (+140 thousand),

Figure 5: Change in the share of workers with degrees within occupations from 2011-19



Data from the UK Labour Force Survey 2011-19. A non-parametric line of best fit is drawn in red. The size of the bubbles reflects the size of the occupations.

The patterns shown in figures 4a and 4b may potentially reflect only the relative sizes of occupations, not that their compositions have changed. Thus to complement those diagrams, I plot a scatterplot of the change in the share of workers with degrees within an occupation from 2011 to 2019 against the intial degree share in 2011 in figure 5. The figure shows that occupations with fairly low shares of workers with degrees experienced a substantial increase in the share of workers in the occupation with degrees, consistent with figure 4. Thus, any theory of the impact of higher education expansion on the graduate labour market has to account for the fact that most of these new graduates end up working in occupations with fairly low shares of workers with degrees.

2.3 Discussion

In this section, I have presented a number of facts about the graduate labour market in the UK.

1. There has been a substantial increase in the number of graduates in the labour market.
2. The average wage of graduates relative to non-graduates (the college wage premium) was largely constant for a significant period of time and has fallen in the last few years. It nevertheless is still significantly positive.
3. The relative college-high school wage ratios have fallen substantially particularly at the 10th percentile, but has not declined in the period for workers at the 90th percentile, suggesting substantial heterogeneity at different parts of the wage distribution.
4. The share of graduates working in 'non-graduate' jobs requiring high school skills has increased. There is a substantial share of graduates working in occupations classified as requir-

"sales accounts and business development managers" (+135 thousand), and "production managers and directors in manufacturing" (+111 thousand).

ing a high school degree (with numbers ranging from 10% to 25%).

5. Between 2011 and 19, new graduates found employment in occupations with previously low shares of graduates, supporting the hypothesis that the increase in graduate supply has changed to an extent which jobs are non-graduate.

The primary motivating puzzle that my model seeks to address is that the existence of a large college wage premium seems incongruous with about a quarter of the graduates working in non-graduate occupations. The natural conclusion, especially drawing on the conditional quantile regressions presented in fact 3, is that there is substantial heterogeneity in the return to higher education. This raises the question of why workers with low returns would invest in higher education to begin with. A natural conclusion here is that workers have some uncertainty about what their returns are, which leads to ex-post regret even with ex-ante optimal actions.

3 A Model of Education Investment and the Labour Market

In this section, I study a model in which workers first decide in an initial period whether or not to invest in education, and subsequently receive a wage depending on their heterogeneous characteristics and their education choices. The model of education choice in the first period is a Willis-Rosen type model of self-selection into education, but with incomplete information and uncertainty about returns, like in Cunha et al. (2005). The labour market is modelled with a frictionless matching model with transferable utility, like in Sattinger (1993)[18]. The central aim of the model is to rationalise the patterns in the wages and occupational outcomes of graduate and non-graduate workers with their choices to invest in higher education.

This section proceeds as follows. I describe the model and maintained assumptions in section 3.1, including the characterisation of equilibrium in this setting. I consider optimal private investment and optimal social investment in education in this model in section 3.2, and discuss the possibility of ex-post regret in this setting and how this might show up in the form of graduate workers matching to low-skill jobs.

3.1 Description of the model

3.1.1 Setting

The basic set-up is as follows. There are two masses of workers and jobs of equal length. As a primitive, workers are heterogeneous in their ability, denoted by a and distributed with density function $f_A(a)$. They also receive a noisy signal θ, interpreted as their school grades, which imperfectly reveal their underlying ability to the worker. I assume that the noisy signal is additive in ability and a white noise term ε, as follows.

[18]This is, in some ways, the simplest characterisation of a matching market with continuous types on both sides of the market. The main alternatives include two-sided matching with search frictions (Shimer & Smith (2000)), or reducing the number of types of either workers or firms (Acemoglu (1999), Shephard & Sidibe (2019)).

$$\begin{aligned}\theta &= a + \varepsilon \\ \varepsilon &\sim N(0, \sigma_\varepsilon^2)\end{aligned} \tag{1}$$

Jobs have varying levels of productivity, denoted by y. Suppose that there are K discrete occupations, indexed by k, where the distribution of productivity of jobs within that occupation is denoted by $f_y^k(y)$, and the share of jobs in that occupation in the economy is denoted by p_k, where $\sum_{k\in\{1,\cdots,K\}} p_k = 1$. The distribution of productivity in the entire economy is given by $f_y(y) = \sum_{k=1}^{K} p_k f_y^k(y)$[19].

There are two periods in the model. In the first period decide whether to invest in education and receive an instantaneous utility corresponding to their choice. Their education decision, denoted by $e \in \{0, 1\}$ and their initial ability determine their pre-labour market skill level, denoted by $s(a, e)$. It is the skill level s, instead of initial ability a, that determines the worker's return on the labour market. I assume that the skill function $s(a, e)$ is continuous and differentiable in a for both values of e. Furthermore, I assume that $\frac{\partial s(a,1)}{\partial a} > \frac{\partial s(a,0)}{\partial a}$ for all values of a, such that the difference of $s(a, 1) - s(a, 0)$ is increasing. Intuitively, this implies that the skill return to education is increasing in ability. Workers have heterogeneous preferences for education, which they know about prior to the education decision. The decision to invest in education is described in detail in section 3.1.3.

In the second period, workers with skill s and jobs with productivity y can match pairwise with each other to produce joint output $g(s, y)$. I assume that $g(s, y)$ is twice continuously differentiable, increasing and supermodular[20]. A model of this kind is known in the matching literature as a frictionless matching model with transferable utility and has been analysed thoroughly in other contexts. The resulting wage function for workers with skill s is the wage that individuals receive in the second period and is analysed in the following subsection 3.1.2.

In summary, the following assumptions are maintained throughout.

Assumption 1 *In the model, I maintain the following assumptions.*

1. *(Information structure of workers) Workers do not observe their initial ability a, and observe a noisy signal θ, related to a by the following expression. θ is the only information that the worker has about a and their return to education. ε has a mean zero, log-concave distribution.*

$$\begin{aligned}\theta &= a + \varepsilon \\ \varepsilon &\sim N(0, \sigma_\varepsilon^2)\end{aligned}$$

2. *(Assumptions on $s(a, e)$) The skill function $s(a, e)$ is continuous and differentiable in a for both values of e. Furthermore, I assume that $\frac{\partial s(a,1)}{\partial a} > \frac{\partial s(a,0)}{\partial a}$ for all values of a, such that the difference of $s(a, 1) - s(a, 0)$ is increasing in a.*

[19] In this model, occupations do not serve any direct role in the equilibrium match (i.e. there are no preferences for occupations), but they are important for two reasons. First, they help the estimation by providing information on the productivity of the job the worker is matched to, especially if information on the matched job is not available. Second, they lead to testable implications about the resulting share of workers with degrees and the college wage premium within occupations. I do not make any assumptions about the ordering of occupation-specific productivity distributions.

[20] In this case, the supermodularity condition is equivalent to $\frac{\partial^2 g}{\partial s \partial y} \geq 0$ (Topkis (1998)).

3. *(Assumptions on joint output $g(s,y)$) The joint output function $g(s,y)$ is increasing in both s and y, and twice continuously differentiable. The function is assumed to be supermodular, which is equivalent to the following condition since it is twice-continuously differentiable:* $\frac{\partial^2 g}{\partial s \partial y} \geq 0$.

3.1.2 The labour market and the wage function

In the labour market, workers and firms of each type have to choose a type of the corresponding worker to match with. They can make payments, in the form of wages to workers conditional on their skill $w(s)$, and in the form of profits to firms $\pi(y)$ to incentivise the agents to accept those matches. Their individual problems can be characterised as follows.

$$\text{Worker's problem for worker with skill } s^*: \quad \max_y g(s^*, y) - \pi(y) \tag{2}$$

$$\text{Firm's problem for firm with prod } y^*: \quad \max_s g(s, y^*) - w(s) \tag{3}$$

This is a well-studied question in the field of matching theory[21], and in this section, I only present the results and intuition relevant to my application. Post-education, suppose that the distribution of worker skill can be described by distribution $f_s(s)$[22]. The primitives of the problem are thus the distribution of skills $f_s(s)$, the overall distribution of productivity $f_y(y)$, and the joint output function $g(s,y)$.

It turns out that under the assumption that the joint output function $g(s,y)$ is super-modular and continuously differentiable with respect to both variables[23], it is possible to find a set of prices $w(s), \pi(y)$ for the entire support of s and y, which solves the workers' and firms' problems by solving the first-order conditions for the workers' and firms' individual problems. This assignment is also the assignment that maximises total output. The optimal assignment that prevails under the equilibrium matching exhibits positive assortative matching, such that the most productive firm matches with the most skilled worker. Denote as a matching in this model a function $\mu(\cdot)$ that specifies for each argument s a value of y such that a worker with skill y is matched a job with productivity y. Since I have assumed the joint output function to be supermodular, then the optimal assignment that results is given as follows, where $F_s(s)$ and $F_y(y)$ denote the cdfs of the skill distribution and the productivity distribution respectively.

$$\mu(s) = F_y^{-1}(F_s(s)) \tag{4}$$

In equilibrium, the first-order condition of the maximisation problem faced by firms needs to be satisfied by the worker with skill s that it matches to. That is, for all values of s, the following condition should hold, pinning down the shape of the resulting wage function.

[21] The exposition of the problem, its solution and attendant proofs have been discussed in many other papers and books, including Galichon (2016), and Chade et al. (2017).

[22] The determination of $f_s(s)$ is discussed in sections 3.1.3 and 3.1.5.

[23] This insight that complementarity, represented by the property of supermodularity, is central to matching problems is due to Becker (1973). This model is discussed in more detail in Chade et al. (2017) and Galichon (2016).

$$\frac{\partial g(s, \mu(s))}{\partial s} = w'(s) \tag{5}$$

Analogously, the first-order condition of the maximisation problem faced by the worker fixes the shape of the profit function as follows, where $\mu^{-1}(y)$ specifies the skill of the worker that a firm with productivity y matches to.

$$\frac{\partial g(\mu^{-1}(y), y)}{\partial y} = \pi'(y) \tag{6}$$

I note at this point that these results suggest that the wage and profit schedules are convex functions[24]. Intuitively, this derives from the complementarity between worker skill and firm productivity, represented by the supermodularity of the joint output function. Because matching with a highly productive worker has a larger impact on joint productivity for a more productive firm than a less productive firm, the former can afford to pay more than the latter to hire the productive worker, while maintaining a higher profit. Because the relative effect on joint output increases with skill, the relative incentive that firms have to pay similarly increases with skill, leading to the convexity described.

We can then solve for the wage and profit functions by integration.

$$w(s) = w_0 + \int_{-\infty}^{s} \frac{\partial g(x, \mu(x))}{\partial x} dx \tag{7}$$

$$\pi(y) = \pi_0 + \int_{-\infty}^{y} \frac{\partial g(\mu^{-1}(x), x)}{\partial x} dx \tag{8}$$

The wage function is thus fixed up to location, given by w_0 which is a constant of integration. This indeterminacy is discussed in Sattinger (1993); one intuition is that since the incentives for matching with any type depends on the relative incentives, any wage function that preserves the shape of the function leads to the same assignment in Walrasian equilibrium. The values of w_0 and π_0 must be such that $w_0 + \pi_0 = g(s_0, y_0)$, where s_0 and y_0 denote the minimum (or infimum) value of s and y in the domain of the distribution of both variables. The precise value of w_0 relative to π_0 can be interpreted as the bargaining outcome between the least skilled worker and the least productive firm.

3.1.3 The decision to invest in education

Workers have heterogeneous preferences for higher education, denoted by η_1, and for non education, denoted by η_0. I assume that the heterogeneous preference shocks for education are independent of the other shocks in the model, i.e. of the information shock ε, as well as the initial distribution of ability a. Let both η_1 and η_0 be mean 0 distributions, and let $\Delta\eta \equiv \eta_1 - \eta_0$, denoting the difference

[24] To see this, note that for any $s' > s''$, $\mu(s') \geq \mu(s'')$, with non-equality for most continuous distributions of y. Then, by the assumption of supermodularity, $\frac{\partial g(s,\mu(s))}{\partial s}(s') > \frac{\partial g(s,\mu(s))}{\partial s}(s'')$, and so $w'(s') > w'(s'')$.

between the heterogeneous preference terms, also be a random variable with mean zero. Let κ denote the aggregate preference for education relative to no education. κ can be interpreted to account for any pecuniary costs of education as well as any unmodelled non-pecuniary preferences; thus if κ is positive, workers have a preference for higher education on average and vice versa. If a worker chooses to invest in education, they receive an instantaneous utility denoted by $\kappa + \eta_1$, and if they do not, they receive instantaneous utility η_0. They also receive utility from their wages in the subsequent period, with time preference factor β.

The value of choosing education option e for a worker with signal θ is thus as follows. Let $w^B(\cdot)$ denote a wage function which maps a level of skill to a wage on the labour market. $w^B(\cdot)$ defines the wage environment that workers believe will prevail in the labour market at the time of workers' educational choice. For simplicity, I assume that all workers believe that the same wage environment.

$$V(\theta, e, \eta_e, w^B(\cdot)) = (\kappa \times e) + \eta_e + \beta E\{w^B(s(a,e))|\theta\} \tag{9}$$

The difference between the

$$\Delta V(\theta, \Delta\eta, w^B(\cdot)) = \kappa + \Delta\eta + \beta\left\{E\{w^B(s(a,1)) - w^B(s(a,0))|\theta\}\right\} \tag{10}$$

Workers face a discrete choice problem. Their optimal education choice e^* given their beliefs $w^B(\cdot)$ therefore is given as follows.

$$e^*(\theta, \Delta\eta, w^B(\cdot)) = \underset{e\in\{0,1\}}{\arg\max}\, e \times \Delta V(\theta, \Delta\eta, w^B(\cdot)) \tag{11}$$

3.1.4 The distribution of skill in the economy

The distribution of skill in the economy is an important determinant of wages in the economy as it enters the wage equation through the matching function, equation 4. In this model, the distribution of skill in the economy is dependent on who chooses to go into education. When more workers are educated, there is a greater supply of skilled workers in the economy, which affects the wages that workers receive. Let $p(\theta, \Delta\eta)$ denote a function that maps $\mathbb{R} \times \mathbb{R} \to \{0, 1\}$. Intuitively, this function represents an education profile, which maps workers by their observed grade to whether they attend higher education. Heterogeneous preferences for education do not affect skill, which makes it irrelevant for determining the skill distribution. I denote by $P(\theta)$ the integral of $p(\theta, \Delta\eta)$ over the distribution of $\Delta\eta$.

$$P(\theta) \equiv \int p(\theta, \Delta\eta) dF\Delta\eta \tag{12}$$

The probability of investing in education (conditional on the signal value) allows us to compute the distribution of skill in the economy using the Law of Total Probability.

$$\begin{aligned} f_S(s) = \int_{\varepsilon\in\mathbb{R}} & P(s^{-1}(s,1) + \varepsilon) f_A(s^{-1}(s,1)) \left|\frac{ds^{-1}(s,1)}{ds}\right| + \\ & (1 - P(s^{-1}(s,0) + \varepsilon)) f_A(s^{-1}(s,0)) \left|\frac{ds^{-1}(s,0)}{ds}\right| dF_\varepsilon(\varepsilon) \end{aligned} \tag{13}$$

$$F_S(s) = \int_{-\infty}^{s} f_S(x) dx \tag{14}$$

We can also derive the distribution of skill for workers of each educational level as follows.

$$f_S(s|e=1) = \frac{\int P(s^{-1}(s,1)+\varepsilon) f_A(s^{-1}(s,1)) \left|\frac{ds^{-1}(s,1)}{ds}\right| dF_\varepsilon(\varepsilon)}{\int\int P(a+\varepsilon) dF_\varepsilon(\varepsilon) dF_A(a)} \tag{15}$$

$$f_S(s|e=0) = \frac{\int (1-P(s^{-1}(s,0)+\varepsilon)) f_A(s^{-1}(s,0)) \left|\frac{ds^{-1}(s,0)}{ds}\right| dF_\varepsilon(\varepsilon)}{\int\int (1-P(a+\varepsilon)) dF_\varepsilon(\varepsilon) dF_A(a)} \tag{16}$$

3.1.5 Equilibrium under Rational Expectations

To complete the model and define equilibrium, I first introduce two notions. An education profile p is consistent with wage beliefs w^B if:

$$\forall\theta \in \mathbb{R}, \forall\Delta\eta \in \mathbb{R} : p(\theta, \Delta\eta) = e^*(\theta, \Delta\eta, w^B(\cdot)) \tag{17}$$

Second, a wage function w is generated by an education profile p if the education profile implies a skill distribution (by equation 13),and thus a matching function (by equation 4) which imply w (by equation 7). That is to say, the skill distribution $F_S(\cdot)$ is a functional with p as an argument, the matching function $\mu(\cdot)$ is a functional that takes F_S as an argument, and $w(\cdot)$ is a functional that takes $\mu(\cdot)$ as an argument. Thus, this series of functionals give a mapping between each education profile p to a wage function w. For convenience, I denote this mapping of functions by R in the following text.

To complete the model, I need to specify how students generate beliefs about the wage functions they will face on the labour market. This is an area where much research still need to be done. Many authors now realise that it is the perceived return to higher education that matters for studying students' investment decisions (Jensen (2010), Wiswall & Zafar (2014)), but there is still relatively little work on how students form such beliefs about returns. In the absence of more sophisticated theories, I resort to a simple assumption of rational expectations in this paper.

In this context, rational expectations imply that workers are able to anticipate the wages that actually obtain when they are on the labour market and set their beliefs to the predicted wages. If workers have rational expectations, then workers' beliefs w^B, the education profile generated by their optimal education choices p, and the resulting wage structure w have to agree.

Definition 1 *(Equilibrium) Rational expectation equilibrium obtains when workers have beliefs w^{B*}, an education profile p^* obtains, and wages w^* are generated such that:*

1. *p^* is consistent with wage beliefs w^{B*},*
2. *w^* is generated by p^*,*
3. *and $w^{B*}(s) = w^*(s), \forall s$.*

Note that the definition of "generated by" specifies a mapping of education profiles to wage functions, denoted by R. Similarly, the notion of consistency with a set of beliefs maps a function representing wage beliefs w^B to an education profile p. Denote this mapping T. The final condition

trivially maps wages to wage beliefs; denote this by S. Then the conditions for equilibrium imply the following mappings:

$$R(p) = w$$
$$T(w^B) = p$$
$$S(w) = w^B$$

Putting these mappings together, the equilibrium condition implies a self-mapping of the wage belief function to itself:

$$w^B = S(R(T(w^B))) \tag{18}$$

Another way to define equilibrium is thus as follows:

Definition 2 *(An equilibrium wage belief) Let Q denote the set of functions that map $\mathbb{R}_+$ to $\mathbb{R}_+$. A function $w^B \in Q$ is an equilibrium wage belief if it is a fixed point in the self-mapping $w^B = S(R(T(w^B)))$.*

It is also possible to analogously define the self-mapping in terms of education profiles p or wages w.

To conclude, I would like to set out two caveats. First, it is not obvious a priori that rational expectations is a good description of reality, particularly in the case of higher education decisions. In the case of higher education, students typically make one-shot decisions and often do not observe the outcomes of their decisions until a few years into the future. If a cohort of students underestimate the share of their peers who also select into education, they would simply find in their future that they would not have selected into education given the skill distribution which did occur. But they would not be able to go back in time to iterate on their decision to eventually reach the equilibrium outcome. However, it could be realistic that past cohorts would learn from the experiences of future cohorts. It would be interesting to study the evolution of wages and education profiles if beliefs were not rational but had an adaptative structure. I leave this for future extensions, but I note that the rational expectations solution must be a stationary point in such a series of education profiles.

In this paper, I abstract from a number of issues relating to the decision to invest in education that have been discussed in the literature, including heterogeneity in beliefs about both the pecuniary and non-pecuniary benefits of education (Belfield et al. (2020)), systematic differences between sexes, ethnic groups and socio-economic classes in their beliefs about the return to education (Patnaik et al. (2020)). I leave a detailed consideration of these factors in the context of my model to future work, but note that a straightforward way to consider them is in considering the effect of these factors on the propensity to invest in education $P^B(\cdot)$. With a specification of the relation between those sources of heterogeneity and $P^B(\cdot)$, it is conceptually straightforward to evaluate welfare changes in the context of a calibrated or estimated model.

3.1.6 Relation to the occupational structure

While the solution of the individual's problem and the determination of equilibrium in the model does not involve analysis of occupations, the mixed distribution set-up of the distribution of job

productivity allows the model to analyse how graduates match to occupations in this model. Conditional on a worker having skill s, the probability of the worker having a job in occupation k is as follows.

$$Pr(occ = k|S = s) = \frac{p_k f_Y^k(\mu(s))}{\sum_{l\in\{1,\cdots,K\}} p_l f_Y^l(\mu(s))} \tag{19}$$

There are two main advantages to this set-up. First, this allows occupations to be used in the identification of the model in a more natural way. Many past papers have used information on occupations, typically indices constructed from the O*NET dictionary of occupation titles (DOT), to identify structural matching models between workers and jobs. This assumes job heterogeneity within occupations. By allowing for a distribution of job productivities within occupations, I allow both the mean and variance of job productivities to affect the wages and matching to graduate workers in the model.

Second, the model's structure allows the causal effect of education be interpeted as the combination of two effects: the causal effect of education on a worker's productivity in any job, and the causal effect of education on a worker's probability of placing at a better occupation. This coheres with previous research on the channels of the education wage premium, such as Lemieux (2014), which finds that matching to better occupations accounts for at least half of the estimated college wage premium.

It is worth noting that some implicit assumptions made in the analysis. First, an important implicit assumption is that workers have no heterogeneous preferences for occupations. This, of course, is unrealistic; deviations in the predictions of the model for the wages and share of graduate workers in an occupation could plausibly be due to unobserved compensating differentials. Second, I assume that there are no frictions in the labour market part of the model, which implies that there are no occupational mismatch in the labour market. Thus, the entire source of mismatch in this model comes from informational frictions in the educational decision stage. The inclusion of either of these concerns will substantially complicate the matching problem that workers face, and increase the computational costs of the model. I leave the extension of this model to encompass those concerns for future work.

3.2 Optimal education investment

Under the assumptions maintained, it turns out that the optimal education decision, both under the full information case where a is known, and the imperfect information case where a is unknown except through θ, is increasing in a or θ respectively. This also allows heterogeneity in preferences for education to be handled simply in terms of variation in the threshold for educational investment. I first discuss the characterisation of the optimal education decision under both the full-information and incomplete-information settings considering only private utility returns. I then point out that the social planner in general would choose more investment than private individuals because there is a return to education captured by firms which cannot be appropriated by workers. Finally, I describe ex-post regret in this model, for individuals who made the ex-post welfare-reducing choice despite making the ex-ante rational decision, and show that this is entirely characterised by the relevant investment thresholds.

3.2.1 Full-Information Benchmark

Consider the case when ability is fully observable. Define by $e^P(a, \Delta\eta)$ the optimal education choice under full information for a worker with ability a and net preferences $\Delta\eta$. I assume for the subsequent analysis that workers are in rational expectations equilibrium where their beliefs about wages w^B are equal to actual wages.

$$\begin{aligned} e^P(a, \Delta\eta) &= \underset{e \in \{0,1\}}{\arg\max}\, e \times \{\kappa + \Delta\eta + \beta\{w(s(a,1)) - w(s(a,0))\}\} \\ &= \begin{cases} 1 \text{ if } \kappa + \Delta\eta + \beta\{w(s(a,1)) - w(s(a,0))\} \geq 0 \\ 0 \text{ otherwise} \end{cases} \end{aligned} \tag{20}$$

Denote for convenience $\Delta w(a) \equiv w(s(a,1)) - w(s(a,0))$. The equation defining optimality also implicitly defines a cut-off value of ability $a^P(\Delta\eta)$ for each level of education preference. We can visualise the optimality condition as a line in a-$\Delta\eta$ space. By implicitly differentiating the condition by $\Delta\eta$, we can further show that this line must be downward-sloping, thus defining a boundary, beyond which workers invest in education and before which they do not. This is visualised in figure 6 in the red line.

Proposition 1 *(Optimal private education choice under full-information)*

1. *The optimal education decision under full information is characterised by a cut-off point $a^P(\Delta\eta)$ for each value of $\Delta\eta$, which is the unique point that satisfies the optimality condition $w(s(a,1)) - w(s(a,1)) = -\left\{\frac{\kappa+\Delta\eta}{\beta}\right\}$. This cut-off is decreasing in $\Delta\eta$.*
2. *The optimal education decision under full information $e^P(a, \Delta\eta)$ is increasing in a, and $\Delta\eta$.*

Proof: *See appendix B.1.*

3.2.2 Imperfect information

Now, I consider my main setting where ability is not observable by the worker. Workers face an optimisation decision where they choose their education based on a signal θ, which is related to initial ability a through the structure specified by equation 1. The standard deviation of signal noise, σ_ε, determines how informative the signal is about the worker's underlying ability.

Denote by $e^I(\theta, \Delta\eta)$ the optimal education decision conditional on the signal and education preferences.

$$e^I(\theta, \Delta\eta) = \underset{e \in \{0,1\}}{\arg\max}\, e \times \{\kappa + \Delta\eta + \beta E\{w(s(a,1)) - w(s(a,0))|\theta\}\} \tag{21}$$

We can show that the optimal decision $e^I(\theta)$ follows an optimal cut-off point structure under the maintained assumptions. The proof uses results about comparative statics under uncertainty provided in Athey (2002), particularly theorem 2 in her paper. She considers problems where

workers aim to maximise a stochastic function $U(x,\theta)$ with respect to $x \in B \subset \mathbb{R}$, where $U(x,\theta)$ can be expressed as follows.

$$U(x,\theta) \equiv \int u(x,s) f(s;\theta) d\mu(s)$$

She provides conditions on the primitive functions $u(x,s)$ and $f(s;\theta)$ such that the argument maximising $U(x,\theta)$ is non-decreasing in θ and B. Theorem 2 states that two conditions are jointly sufficient:

1. The density function f is log-supermodular, which is satisfied when it obeys the monotone likelihood ratio (MLR) order.
2. The payoff function $u(x,s)$ satisfies single-crossing in (x,s); this means that for all $x' > x$, then $u(x',s) - u(x,s)$ crosses 0 at most once and from below.

The first condition is satisfied when $f(s;\theta)$ obeys the monotone-likelihood ratio order. Consider a family of distribution functions $\{\lambda(\cdot,\theta)\}_{\theta\in\mathbb{R}}$; this family of functions obey the MLR order if $\frac{\lambda(a,\theta'')}{\lambda(a,\theta')}$ is increasing in a whenever $\theta'' > \theta'$. It turns out that under the additive error structure assumed in 1, the conditional distribution satisfies MLR order in θ regardless of the initial distribution of a. Thus, given additive noise, the conditional distribution of the underlying random variable conditional on signal θ obeys MLR order with respect to the size of the signal θ as long as the error ε has a log-concave distribution (see a proof of this, see appendix B.2). This is assumed in assumption 3.1.1.

For the second condition to be satisfied, we have to show that $\kappa + \Delta\eta + \beta\{\Delta w(a)\}$ crosses 0 at most once and from below. This is satisfied when the term is strictly increasing, which was proved in appendix B.1 as part of the proof of proposition 1. Thus, we can prove that the optimal investment is non-decreasing in θ, conditional on $\Delta\eta$, by directly applying Athey's results.

Proposition 2 *(Partial information case under signals with additive, normally distributed noise) Denote by $e^I(\theta,\Delta\eta)$ a function that maps θ and $\Delta\eta$ to the optimal education decision of a worker with signal θ and preferences $\Delta\eta$ who maximises their expected utility under imperfect information. Then, conditional on Δeta, $e^I(\theta,\Delta\eta)$ is weakly increasing in θ under the maintained assumptions.*

Proof: *Direct application of Athey (2002), theorem 2, p. 200.*

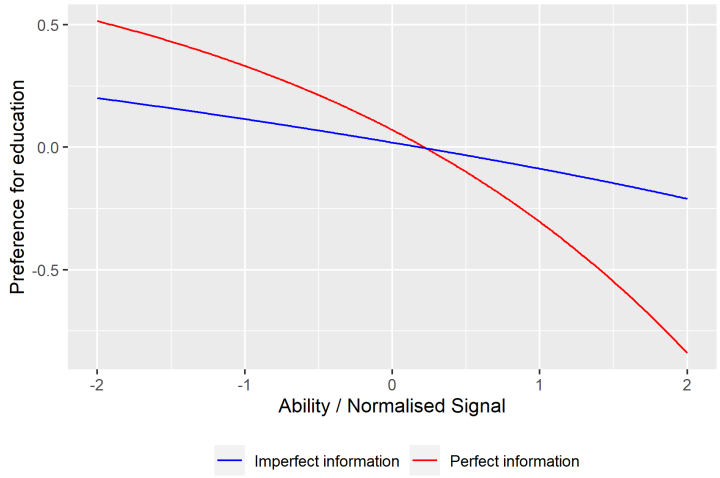
The corollary of proposition 3.2.2 is that since the condition for the worker choosing $e = 1$ over $e = 0$ is $V(\theta,1) - V(\theta,0) > 0$, then if $e^I(\theta)$ is non-decreasing in θ, then $V(\theta,1) - V(\theta,0)$ is also weakly increasing in θ. This implies that $E\{w(s(a,1)) - w(s(a,0))|\theta\}$ is also weakly increasing in θ.

Denote by $\theta^P(\Delta\eta)$ the cut-off value of θ such that given heterogeneous preferences $\Delta\eta$, a worker with signal $\theta^P(\Delta\eta)$ would be indifferent between investing and not investing in education.

$$E\left\{w(s(a,1)) - w(s(a,0))|\theta^P(\Delta\eta)\right\} = -\left\{\frac{\kappa + \Delta\eta}{\beta}\right\} \tag{22}$$

Since $E\{w(s(a,1)) - w(s(a,0))|\theta\}$ is weakly increasing in θ, $\theta^P(\Delta\eta)$ should be weakly decreasing in $\Delta\eta$, and $e^I(\theta)$ would be weakly increasing in $\Delta\eta$. Workers with a greater preference for higher

Figure 6: Boundary Lines Describing Optimal Education Investment in Perfect and Imperfect Info



This graph plots a parameterised example optimality condition under perfect information and imperfect information. The parameterisation comes from the empirical implementation, described in subsequent sections.

education would be more likely to invest in education across all values of θ. As is the case in the perfect information scenario, the cut-off function differs a boundary in two-dimensional θ-$\Delta\eta$ space which separates students who will invest in HE and those who will not.

Proposition 3 *(Cut-off strategy for optimal education choice under imperfect information) The optimal education decision under imperfect information is characterised by a cut-off signal $\theta^I(\Delta\eta)$, which is the point that satisfies equation 22. This cut-off is decreasing in $\Delta\eta$, such that workers with greater preference for higher education require a lower signal to induce them to choose to invest in education.*

Figure 6 plots a parameterised version the optimality condition with equality under perfect and imperfect information; as discussed, this condition implicitly defines a decreasing function in $\Delta\eta$ and a space which serves as a boundary for whether students will invest in HE. Students will invest in HE in imperfect information if they are to the right of the blue line, and will not invest in education if they are to the left of the line

3.2.3 Ex-ante rational but ex-post worse off workers

In the motivation of this paper, I discussed the problem of mismatch in the education system, in particular, the phenomenon of over-education, which is commonly described as when a graduate worker works in so-called non-graduate employment. This definition has motivated models with broadly two firm types; among others, Acemoglu (1999), Albrecht & Vroman (2002), and Jackson (2021) propose models where there are two kinds of firms, usually intuitively reflecting high-skill and low-skill firms. A major challenge with this kind of definition is that empirically, it is not clear

how to define which occupations or firms are high or low-skill[25]. Another challenge is that this approach does not capture the possibility that whether an occupation is considered high-skill may be endogenously determined by the availability of skill in the economy. The education composition of an occupation is likely to be different in an economy with few graduates and an economy with high HE completion.

My model offers an alternative, precise notion of education mismatch, based on the idea of ex-post regret. A worker experiences education mismatch when they choose a level of education ex-ante which leads than a worse outcome ex-post.

Definition 3 *(Education mismatch, over-education, under-education) A worker with ability a, signal θ, and education preferences $\Delta\eta$...*

1. *...experiences education mismatch if $e^P(a, \Delta\eta) \neq e^I(\theta, \Delta\eta)$.*
2. *...is over-educated if $e^P(a, \Delta\eta) = 0$ and $e^I(\theta, \Delta\eta) = 1$.*
3. *...is under-educated if $e^P(a, \Delta\eta) = 1$ and $e^I(\theta, \Delta\eta) = 0$.*

There are two kinds of workers who are ex-ante rational but ex-post worse off in my model. There are workers who do not select into education when they would have received a benefit ex-post; in this sense, they under-invest in education. There are then the corresponding type of worker who do invest in education when they do not receive a benefit ex-post; these type of workers over-invest in education. This notion of mismatch is consistent with the traditional phenomenon of over-education. Workers with high signals but low actual ability are more likely to match with lower productivity jobs, more of which are likely to be traditionally considered low-skill. Instead of being the definition of over-education, being in a non-graduate occupation is instead an indicator of being over-educated.

In rational expectations equilibrium, propositions 1 and 3 show that conditional on a worker's education preferences, their education choice under perfect and imperfect information depends on whether their ability or signal exceeds cutoffs, denoted by $a^P(\Delta\eta)$ and $\theta^I(\Delta\eta)$. Thus, we can define education mismatch in terms of these cut-offs as follows:

Definition 4 *A worker with signal θ and net education preferences $\Delta\eta$*

1. *is over-educated if they have $\theta > \theta^I(\Delta\eta)$ and $a < a^P(\Delta\eta)$.*
2. *is under-educated if they have $\theta < \theta^I(\Delta\eta)$ and $a > a^P(\Delta\eta)$.*

The shares of workers who end up over-educated or under-educated can be computed given the

[25] Leuven & Oosterbeek (2011) provides a critical discussion of the measurement issue. Many papers typically use an empirical measure which considers an occupation to be graduate if over 50% of workers in the field are graduates. This cut-off is largely arbitrary and in fact the share of graduates in occupations is not bimodal but largely continuous between 0 and 100%.

joint distribution of θ and a.

$$\begin{aligned} Pr(\text{overinvest}) &= Pr(a < a_A^P(\Delta\eta), \theta > \theta_B^I(\Delta\eta)) \\ &= \int_{-\infty}^{a_A^P(\Delta\eta)} (1 - F_\varepsilon(\theta_B^I(\Delta\eta)) - a)) f_A(a) da \end{aligned} \tag{23}$$

$$\begin{aligned} Pr(\text{underinvest}) &= Pr(a > a_A^P(\Delta\eta), \theta < \theta_B^I(\Delta\eta)) \\ &= \int_{a_A^P(\Delta\eta)}^{\infty} F_\varepsilon(\theta_B^I(\Delta\eta)) - a) f_A(a) da \end{aligned} \tag{24}$$

Outside of rational expectations equilibrium, mismatch can also occur if workers' beliefs about the wage schedule that they will face does not align with the actual wage schedule. I abstract from this issue in my empirical application, and leave the quantification of mismatch due to mistaken beliefs to future work.

3.3 Socially Optimal Education Investment

Suppose there were a social planner who could choose an education profile, that is, a mapping $p : \mathbb{R} \times \mathbb{R} \to \{0, 1\}$ that specifies a level of education investment for each level of grades and education preference. Denote by $\mathcal{P}$ the set of education profiles p. This planner aims to choose an education profile to maximise the sum of joint output and students' preferences in the economy. The social planner's objective function is as follows, denoted by W:

$$\begin{aligned} W[p] = \int_{\Delta\eta} \int_{\varepsilon} \int_{a} \{ & p(a + \varepsilon, \Delta\eta) \left[g(s(a, 1), \mu(s(a, 1))) + \kappa + \Delta\eta \right] + \\ & (1 - p(a + \varepsilon, \Delta\eta)) \left[g(s(a, 0), \mu(s(a, 0))) \right] \} \, dF(a) dF(\varepsilon) dF(\Delta\eta) \end{aligned} \tag{25}$$

Note that besides directly appearing in the equation, $\mu(\cdot)$ is also a functional in p, as p determines the supply of skill in the economy, which affects the match conditional on s. Denote by $\mu(s, p)$ the function with the function p as a secondary argument and denote by $D\mu(s; p_0)$ the derivative of $\mu(s)$ with respect to p, evaluated at p_0. By extension, this implies that $g(s, \mu(s))$ is a functional of p, and let $Dg(s; p_0)$ denote the derivative of joint output conditional on skill s, at education profile p_0[26].

It is in general difficult to solve the social planner's problem; since the object optimised over is a function, it is not straightforward to use conventional derivative based methods to solve the problem. To reduce the set of plausible candidate education profiles, I note that any candidate solution to the planner's problem must be monotonically increasing in $\Delta\eta$.

Proposition 4 *Let p^* denote the solution to the planner's problem:*

$$p^* = \underset{p(\theta, \Delta\eta) \in \mathcal{P}}{\arg\max} \, W \tag{26}$$

Then for all $\theta \in \mathcal{R}$, if $p^*(\theta, \Delta\eta) = 1$ and $p^*(\theta, \Delta\eta') = 0$, then $\Delta\eta > \Delta\eta'$.

[26] I define the functional derivatives precisely in the proof in appendix B.3.

***Proof**: Suppose this is not the case. Then, for some θ, $p^*(\theta, \Delta\eta) = 0$ and $p^*(\theta, \Delta\eta') = 1$, where $\Delta\eta > \Delta\eta'$. But it would be welfare-improving to reallocate an equal measure of workers with the lower $\Delta\eta'$ out of education and workers with the higher $\Delta\eta$ into education. Because this does not affect the supply of skill in the economy, this change will not affect the wage incentives.*

Let $\mathcal{Q}$ denote the subset of functions in $\mathcal{P}$ that satisfy the monotonicity condition. Functions in this set can be represented by specifying a cut-off value of $\Delta\eta$ for each value of θ, and for each element in this set, there is a corresponding element in the set of functions Ψ, of functions mapping the real line to the real line $\psi : \mathcal{R} \to \mathcal{R}$.

Denote the education profile that prevails under equilibrium by $\bar{p} \in \mathcal{P}$. I note at this point that the equilibrium education profile $\bar{p}$ is monotonic in $\Delta\eta$ because of proposition 3, and that it has a corresponding element in Ψ, which I denoted by $\bar{\psi}$.

It turns out that the education profile that prevails in equilibrium is not the socially optimal education profile.

Proposition 5 *Suppose the economy is in rational expectations equilibrium, where the education profile in equilibrium is $\bar{p}$. Then, $\bar{p}$ is not the socially optimal education profile, i.e. $\bar{p} \neq p^*$*

***Proof:** See appendix B.3.*

Inspecting the first-order condition of the derivative of W with respect to p at $\bar{p}$, it is possible to determine two main explanations for why equilibrium is not efficient. I term these two externalities, the hold-up externality and the congestion externality.

$$DW[\bar{p}] = \int\int\int \underbrace{\{\pi(s(a,1),\bar{p}) - \pi(s(a,0),\bar{p})\}}_{\text{Hold-up externality}} + \underbrace{\{\bar{p}(a+\varepsilon,\Delta\eta) \times Dg(s(a,1);\bar{p}) + (1-\bar{p}(a+\varepsilon,\Delta\eta)) \times Dg(s(a,0);\bar{p})\}}_{\text{Congestion externality}} dF_a dF_\varepsilon dF_{\Delta\eta} \tag{27}$$

The first-term representing the hold-up externality is straightforward to interpret. Workers do not manage to appropriate the full increase in joint output when they get educated, and thus, there is positive welfare to be gained from increasing education in the economy. This suggests that there is under-investment in education at rational expectations equilibrium, relative to social optimum.

The second-term denotes the indirect effect of changing the education profile on the skill distribution, and hence on the jobs that each worker matches with. The interpretation of the congestion externality is different from that in search and matching models. In those models, the congestion externality is the effect of workers' search or firms' vacancy posting on the value of search or vacancy posting. In my model, the congestion comes from the effect of education on the distribution of skill. A worker with any given skill level will match to more productive firms in a less skilled economy than a more skilled economy. Workers do not consider this effect on other workers when making their education decision and thus may not invest in education at the socially optimal level.

In general, the rational expectations equilibrium is thus not socially efficient. I show that this is indeed the case in my simulations of the model, and that the congestion and hold-up externalities counteract. It is generally ambiguous whether education is too high in my model.

4 Structural Implementation

I structurally estimate a parametric version of my model and use it to perform welfare analysis and counterfactual analysis. In this section, I will describe my empirical approach. Section 4.1 describes a parametric specification of the model described above. Then, I discuss the identification of my parameters and estimation in subsection 4.2. Finally, section 4.3 describes the data sources I use for my empirical application.

4.1 Parameterization

To bring the model to data, I impose a number of parametric assumptions on the exogenous functions and distributions in my model, summarised below in table 1. Most of the parameterisations are chosen for simplicity and tractability; I assume that ability is normally distributed, with the mean and variance normalised to 0 and 1[27]. The signal noise is Gaussian, with variance σ_ε^2 which is interpreted as the relative degree of uncertainty that workers face about their true underlying ability. This parameterisation of the prior and the noise is commonly used in the economics of information, and a well-known result is that the resulting posterior distribution, that is the distribution of $a|\theta$, is also normal with a mean and variance that depend on the signal and the signal noise, i.e. $a|\theta \sim N(\frac{\theta}{1+\sigma_\varepsilon^2}, \frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2})$. The normality of the posterior distributions is convenient as it allows the use of quadrature methods to compute objects of interest, particularly the expected return to education conditional on the received signal.

To parameterise the joint output function, I follow the Cobb-Douglas parameterisation chosen by Chade & Lindenlaub (2021), where the joint output is given by $qs^{\gamma_1}y^{\gamma_2}$. This implies that the slope of wages as a function of skill is parameterised by $q(\gamma_1)s^{\gamma_1-1}\mu(s)^{\gamma_2}$. This parameterisation satisfies supermodularity if the parameters q, γ_1 and γ_2 are greater than 0, and if the domain of s and y is restricted to $\mathbb{R}_+$. Because only the shape of the wage function is determined, and not the location, I assume a minimum wage that will be estimated from the data, denoted by w_0.

To parameterise the skill function, which relates underlying ability and education to skill, I use a simple exponential form, where a worker with skill a has skill $\exp(a)$ if uneducated and $(1+\delta)\exp(a)$ if educated. This satisfies the assumptions made on the skill function in section 3, particularly that the marginal increase in skill to ability is greater for educated than non-educated workers. At this point, I note that the parameterisation that was chosen assumes that there are no effects on the occupational match and wages of higher education besides through its effects on skill. This model does not contain signalling (since firms observe worker skill, there is no asymmetric information) and sheep-skin effects.

I also parameterise the distribution of job productivities as a mixture distribution with K components, where K is the number of occupations. The full specification of this distribution involves specifying the weights, a $K \times 1$ vector denoted by $p = \big(p_1, \cdots, p_K\big)$, and the distribution for each component distribution. I assume that the job productivity within each distribution is log-normal with mean μ_y^k and variance σ_y^k. These moments are collected in the two $K \times 1$ vectors μ_y

[27]This assumption is commonly made, and gave rise to Pigou's paradox, which asks why wages are lognormal when ability is assumed to be symmetric. The seminal contribution of Roy (1951) can be interpreted as a response to this paradox. An interesting discussion of this early work is in Heckman & Sattinger (2015).

and σ_y respectively. While this parametric assumption may still be restrictive, I argue that it is an improvement on current approaches in the literature, which simply assumes that each occupation is associated with a particular productivity level[28].

Finally, I assume that preferences for higher education are consist of two components, an aggregate net preference for higher education, κ, and non-systematic preferences for higher education are parameterised by $\eta(1)$ and $\eta(0)$, which I assume to be mean 0 type I extreme value distributions with scale parameter $\frac{1}{\xi}$. The difference $\eta(1) - \eta(0)$ is logistically distributed with location parameter 0 and scale parameter $\frac{1}{\xi}$. For simplicity, I assume β in equation 9 is equal to 1. It is not separately identified from the scale of the net preference for education ξ.

Table 1: Parametric Specification

No	Object	Notation	Parametric Form	Parameters
	Exogenous functions			
1	Skill function	$s(a,e)$	$\exp(a)(1+\delta e)$	δ
2	Joint output function	$g(s,y)$	$qs^{\gamma 1}y^{\gamma 2}$	$q, \gamma 1, \gamma 2$
	Exogenous distributions			
3	Ability distribution	$f_A(\cdot)$	$N(0,1)$	-
4	Dist. for signal noise	$f_\varepsilon(\cdot)$	$N(0,\sigma_\varepsilon^2)$	σ_ε
5	Demeaned dist. for net heterogeneous educ pref	$f_{\Delta\eta}(\cdot)$	Logistic with loc 0 and scale $\frac{1}{\xi}$	ξ
6	Aggregate preference for HE	κ	-	κ
	Other parameters			
7	Minimum wage	-	-	w_0
8	Job productivity distribution	$f_Y(\cdot), f_Y^k(\cdot)$	$logN(\mu_y^k, (\sigma_y^k)^2)$	μ_y^k, σ_y^k

There are two aspects of the parameterisation which may be seen as extreme. First, the model starkly predicts that workers with the same ability should receive the same wage, and thus any variation in wages is attributed to the worker's observed signal deviating from their actual ability. A reasonable objection is that there may be sources of variation in observed wages that is not due to this mechanism, particularly data-driven sources of variance like observation error. Unfortunately, adding observation error to the model leads to problems with optimisation, which suggests that the variance of the observation error and the variance of the signal noise is not well identified separately. Thus, my estimates are likely to overstate the true uncertainty in the system if there is significant observation error in the data.

Second, my mechanism for assigning observed occupations to workers assumes that workers are indifferent to the occupation choice beyond the productivity of the job offered. This excludes concerns like preference for particular occupations, e.g. for prestige reasons. Implementation of these elements in my model is unfortunately not straightforward; incorporating these elements requires a more complex matching problem than has been described thus far. As I shall show in the discussion of model fit, this simple assumption seems to fit outcomes for non-graduate workers well but workers seem to have a preference for skilled preferences even when the productivity of the job they match to is the same. Thus, my model captures the dependence of the share of graduate workers in skilled occupations on the observed signal, but misses the level.

[28]This is the case in Chade & Lindenlaub (2021), as well as the occupational mismatch literature (e.g. Guvenen et al. (2020); Lise & Postel-Vinay (2020)).

4.2 Identification and Estimation

To understand whether the model is identified, I start by noting that it is possible to construct the distribution of skill in the economy under a particular set of parameters with an estimate of the share of workers who get degrees, conditional on θ, according to equation 13, conditional on δ and σ_ε. I assume that I can observe the components of the productivity distribution by observing the occupation shares, and occupation-specific means and variances, p_k, μ_y^k, σ_y^k for all $k \in \{1, \cdots, K\}$. Chade & Lindenlaub (2021) note that, following Ekeland et al. (2004), the parameters of the joint output function, q, γ_1, γ_2 and the minimum wage w_0 are identified when the distributions of ... In my case, this means that it is possible to identify parameters q, γ_1, γ_2 and the minimum wage w_0 conditional on parameters ε and δ.

Given the joint output parameters, the education technology parameter δ is governed by the wage return to education conditional on a; since a is not observed, the natural counterpart is the wage return conditional on θ. In a straightforward way, if δ is large, there should be a larger gap between the observed wages conditional on θ for $E[w|e = 1, \theta]$ and $E[w|e = 0, \theta]$. The signal noise standard deviation σ_ε is identified by the variance of wages conditional on θ. In the model, the fact that workers with the same signal can be observed to have different wages is explained by the signals being correlated imperfectly with true underlying ability. Thus, if workers with the same signal have significantly varying wages, this must be because the observed signal is fairly weakly correlated with true underlying ability. Conversely, if σ_ε was small, then we should observe that the signal θ is tightly correlated with wages, especially conditional on education. Thus, estimates of the joint output parameters and w_0 help pin down the education technology parameter δ and the signal noise parameter σ_ε. These parameters in turn determine the empirical skill distribution, which pins down the joint output parameters and w_0.

Finally, the parameters κ and ξ cannot be identified without taking a stand on workers' beliefs about their returns to education. To estimate these parameters, I assume that workers have rational expectations, that is, they have knowledge of the parameters of the model and can accurately assess their likely returns in expectation prior to education investment, and that the model is in equilibrium, i.e. that the observed distribution of workers choosing education is consistent with their incentives to pursue education[29]. In this case, since the parameters $\delta, q, \gamma_1, \gamma_2, \sigma_\varepsilon$ and w_0 are sufficient to construct workers' incentives conditional on their signal θ, I can identify ξ and κ using a binomial logit model, especially since I assume that the heterogeneous preference variables $\eta(1), \eta(0)$ are distributed according to extreme value type I.

This discussion suggests the following procedure for identifying the model parameters:

1. Non-parametrically[30] construct the share of workers who choose higher education in the data conditional on the observed signal θ at a set of points, and denote this $\hat{P}(\theta)$.
2. Estimate the parameters governing wages, the occupational match and the returns to education conditional on $\hat{P}(\theta)$ using simulated method of moments. This recovers the estimated partial parameter vector, $(\hat{\delta}, \hat{q}, \hat{\gamma}_1, \hat{\gamma}_2, \hat{w}_0, \hat{\sigma}_\varepsilon)$.

[29] It is possible that the model is not in equilibrium and thus that a different set of workers would choose education than is actually observed under rational expectations. This would make identification of the preference parameters κ and ξ impossible. However, this cannot be ruled out conclusively ex ante.

[30] I do this by essentially constructing a histogram, and taking the mean share of workers with degrees at 19 points across the signal distribution.

3. Using $(\hat{\delta}, \hat{q}, \hat{\gamma}_1, \hat{\gamma}_2, \hat{w}_0, \hat{\sigma}_\varepsilon)$ and $\hat{P}(\theta)$, I can calculate the return to education $\Delta(\theta) \equiv E[w(s(a,1)) - w(s(a,0))|\theta]$. Under the assumptions that workers have rational expectations and that the model is in equilibrium, I can identify κ, ξ using a binary logit model, with $\Delta(\theta)$ as the RHS variable.

4.2.1 Estimation of the joint output function, education technology function, and signal correlation

To jointly estimate $\delta, q, \gamma_1, \gamma_2, \sigma_\varepsilon$ and w_0, I use the simulated method of moments (Gourieroux et al. (1996); Adda & Cooper (2003)). The basis of this estimation method is to simulate datasets based on the parameters, and to find a set of parameters that produce moments that best approximates moments computed from the data. Let φ denote the vector of parameters to be estimated, and let $\hat{m}$ denote a vector of n_m targeted moments computed from the data. For a given set of parameter vectors, I simulate the model R times with N observations in each simulation, and compute a vector of equivalent moments from the simulated data. I then average the R vectors over the simulations to derive the average simulated counterparts of the moments in $\hat{m}$, which I denoted by $\tilde{m}(\varphi)$[31].

Define by $q(\varphi) \equiv \hat{m} - \tilde{m}(\varphi)$, which is a $n_m \times 1$ vector. Then the estimator is given as follows:

$$\hat{\varphi} = \arg\min_{\varphi} q(\varphi)' W q(\varphi) \tag{28}$$

, where W is a $n_m \times n_m$ positive semi-definite weight matrix. While any choice of W yields consistent estimates of the parameters φ, the optimal weighting matrix is the inverse of the covariance matrix of the moments used, which I compute using the bootstrap. Additional details about computation are available in appendix C.

Using this estimation strategy, it is crucial to choose the moments that can identify the parameters specified. Table 2 summarises the moments that I use for my estimation. There are 126 moments in all, in seven categories. The mean log wage within log-wage deciles (categories 1 and 2) help to identify the joint output function by summarising the shape of the wage distribution. The means, especially conditional on education and signal quintile, help to identify δ, which governs the return to education. The variances as well as the quartiles of log wages conditional on education and signal quintile help to identify σ_ε, which governs how good the signal is of underlying ability. Another important moment is the R^2 of regressing log earnings on a high degree polynomial conditional on education status. In a world without noise, the signal is as good as observing ability and thus the R^2 should be 1. As the degree of noise increases (i.e. σ_ε increases), the R^2 declines as the signal no longer perfectly tracks wages conditional on education status. Thus, this moment is important for identifying the degree of noise in the grade signal of ability.

4.2.2 Estimation of education preference parameters

Conditional on estimates of δ, σ_ε, the parameters of the joint output function, q, γ_1, γ_2 and w_0 from the estimation procedure described in the previous section, as well as the probability of higher education conditional on the signal, I can construct the expected return to education conditional on the signal, $\Delta(\theta) \equiv E[w(s(a,1)) - w(s(a,0))|\theta]$.

[31] I use 50 simulations, each with 2000 simulated observations.

Table 2: Moments

No.	Moments category	Number of moments
1	Mean log wage within income deciles	10
2	Mean log wage within income deciles cond. on education status	20
3	Log wage quartiles cond. on signal quintile	20
4	Log wage quartiles cond. on signal quintile and education	40
5	Mean and variance of log wages	2
6	Mean and variance of log wages cond. on education	4
7	Mean and variance of wages cond. on signal quintile	10
8	Mean of wages cond. on signal quintile and education	10
9	R^2 of regressing log earnings on a polynomial of grades conditional on degree	2

Suppose that workers have rational expectations, and the system is in equilibrium (i.e. the incentives to invest in education align with the observed choices to undergo education). Under this assumption, I then estimate the preference parameters $\psi = \{\kappa, \xi\}$, the location and scale of the extreme value type I distribution of $\eta_1 - \eta_0$, by the Generalised Method of Moments (GMM). The natural moments are the shares of workers with degrees conditional on the worker's signal. Denote by $\hat{\psi}$ the estimated values of ψ. Suppose that $\hat{P}_\theta$ denotes a vector containing the probability of having a degree conditional on $\tilde{\theta}$, where $\tilde{\theta}$ denotes a set of values of θ corresponding to the 5th quantile to the 95th quantile of the distribution of θ. Let $\tilde{P}_\theta(\psi)$ denote the probability of investing in education implied by $E[w(s(a,1)) - w(s(a,0))|\theta]$ and ψ at the points $\tilde{\theta}$. Then, the minimisation problem underlying the estimation is as follows. The standard errors are computed based on the Jacobian of the moments to the parameters (Greene (2003)).

$$\hat{\psi} = \arg\min_{\psi} \left(\tilde{P}_\theta(\psi) - \hat{P}_\theta\right)' W \left(\tilde{P}_\theta(\psi) - \hat{P}_\theta\right) \tag{29}$$

There are two reservations with my education choice set-up as it stands. First, this estimation approach does not account for the possibility that preferences for higher education may differ substantially across demographic groups. Suppose that there are G distinct demographic sub-groups, with the generic sub-group indexed by g. Then, a straightforward strategy is to use data on the observed higher education choice shares for each sub-group g, $\hat{P}_g(\theta)$, to estimate sub-group specific choice parameters. While I do not pursue this strategy in this paper for sample size reasons, it is conceptually straightforward to account for this concern in future research with a bigger dataset.

Second, this method does not allow for psychic costs of education that are correlated with workers' grades. This is, for example, in contrast to Chade & Lindenlaub (2021) which attribute all the correlation between ability scores and education choice to such costs. Without additional variation of earnings not due to differences in grades, such psychic costs are not identified in my model[32]. To the extent that this is a big concern, my estimation method would overstate the degree to which workers are responsive to pecuniary incentives in education choice, and understate the

[32] A promising method to address this concern may be to use variation across cohorts due to wage variations across the business cycle or changes in tuition fees due to policy. I do not have sufficient observations in this dataset to pursue this strategy and leave this for future research.

willingness to pay for higher education.

4.3 Data

To estimate the model, I need two sources of data. First, I need the set of moments that are used to identify the parameters in the model, described in table 2, which requires observations on labour market wages, whether the worker has a degree, a measure of their observed ability, and an observation of the occupation they ended up matching to (for analysis of untargeted moments). I use a sample of individuals from the Understanding Society longitudinal study[33]. This dataset allows administrative linkage to data on high school grades in a national exam, the GCSE, which are the grades that students rely on when applying to higher education. I pool observations from individuals born across five years from 1988-1993, who would have been 18 and considering higher education around the years 2006 to 2011[34], and use this as my main sample. For this sample of students, I have data on their signal, whether they attended higher education, a set of controls, and an unbalanced panel of earnings between the ages of 25 to 32.

For wages, I use hourly labour income, net of taxes and transfers. I face two issues which commonly affect empirical analyses of the relation between education choice and the labour market. First, the 'correct' wage measure to consider in calculating a worker's return to education is their total life-time earnings, which is unfortunately not observed in most datasets. Second, wages are typically the influence of many observable factors, a large number of which is not considered in my model. To address both issues, I run a Mincer regression of log earnings on age, grades, and a set of controls[35], on a panel of earnings for my sample. I adjust earnings using the coefficients from the regression to control for the effects of sex, ethnicity and wave fixed effects. I then take the predicted earnings at age 30 as the relevant earnings that I consider in the structural estimation[36]. Considering earnings at age 30 is common in the literature on higher education given the data typically available and practical constraints on collecting data on the whole life-cycle; see for example similar measures in papers including Blundell et al. (2000); Beffy et al. (2012); Delavande & Zafar (2019); Belfield et al. (2018a). It is helpful to consider these age 30 wages as a proxy for lifetime earnings.

As a measure of the signal that individuals consider, I use a variable recording students' performance at their Key Stage 4 exams[37] (Total GCSE/GNVQ new style point score). I normalise this

[33]Understanding Society (University Of Essex (2022)) is a longitudinal study which follows households over a long period of time, collecting data on a wide variety of topics for people within the household of all ages. It is a successor of the British Household Panel Survey (BHPS). The linkage to administrative education data is provided by Education & University Of Essex (2022).

[34]The government's policy towards higher education funding changed often in the late 1990s and 2000s, and thus I deliberately choose to pool cohorts which faced similar conditions for higher education funding. In this period (from the 2005/06 academic year to 2012/13), students typically borrowed about £6000 per year, under student loan plan 1. These debt details are taken from a government research report, accessed at https://researchbriefings.files.parliament.uk/documents/SN01079/SN01079.pdf.

[35]The controls I include are wave fixed effects, sex, and ethnicity.

[36]For each individual, I compute the mean residual of the regression summarised in table 3 over the periods that they are observed for in the panel. I add this mean residual to predicted earnings conditional on their grades predicted at age 30.

[37]In the UK, the education system is based on five key stages. Key Stage 4 refers to two years typically called Years 10 and 11, when students are between 14-16. This stage is capped by National Examinations, most typically the GCSE (General Certificate of Secondary Education), although other vocational qualifications are available. The

Table 3: Results of a Mincer regression on sample

	(1) Log net hourly labour income
Degree=1	0.0721***
	(3.43)
Normalised KS4 score	0.0752***
	(6.02)
Degree x Normalised KS4 score	0.0462
	(1.86)
Female	-0.0781***
	(-3.95)
Age	0.0363***
	(5.96)
Observations	3122
Individuals	1128

t statistics in parentheses

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

This table presents results of a Mincer-style regression of log earnings on degree status interacted with KS4 scores, age, gender, ethnicity and year fixed effects, estimated on an unbalanced panel dataset of workers born between 1988-93. The standard errors were clustered at the individual level.

score within the student's academic year; due to attrition, the distribution of this normalised score in my sample is not standard normal. There is a growing literature documenting that grades play a significant role in providing individuals with information on their ability (Tan (2022)) and their comparative advantage in subject choice (Avery et al. (2018); Li & Xia (2022)). Similarly, my set-up proposes that students learn about their relative aptitude for university and their heterogeneous return from higher education from their grades at GCSE. An important issue is that the researcher does not observe the individual's true information set and thus may be liable to overestimate the degree of uncertainty individuals face[38]. Workers may have further private information about their returns than is revealed through grades. I leave a satisfactory resolution of this important issue to future work and note that if private information is indeed significant, my results represent an overstatement of the degree of individual uncertainty.

I thus construct a dataset comprising 1113 observations containing three variables: whether workers have a degree at age 32, their adjusted log wage at age 30, and their normalised KS4 grade, which I interpret as the worker's signal of their underlying return. The summary statistics for the dataset from which the moments are computed are tabulated in table 4.

Table 4: Summary statistics

Variable	N	Mean	Sd
Log hourly labour earnings net of taxes and transfers	1113	2.43	0.26
New style KS4 point score, normalised within student's academic year	1113	0.22	0.91
Whether worker has a degree by age 32	1113	0.49	0.50
Female	1113	0.53	0.50
Non-white ethnicity	1113	0.27	0.44
In 1988 birth cohort	1113	0.17	0.38
In 1989 birth cohort	1113	0.20	0.40
In 1990 birth cohort	1113	0.18	0.39
In 1991 birth cohort	1113	0.16	0.36
In 1992 birth cohort	1113	0.16	0.36
In 1993 birth cohort	1113	0.13	0.34

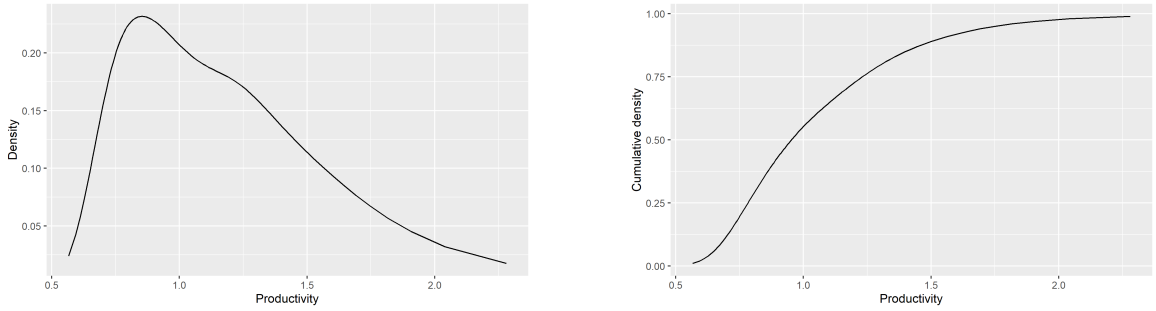
Second, I need data on the distribution of the productivity of jobs that workers match to, that is, the distribution of $f_Y(\cdot)$. As a proxy for productivity, I use the distribution of firm fixed effects within an occupation, calculated from a two-way fixed effect regression of log wages on individual worker and firm-occupation pair fixed effects. I assume that within each occupation (at the SOC00 3 digit level), the distribution of firm productivity is log-normal, with the mean and variance of the log productivity given by the mean and variance of occupation-firm pair fixed effects within the firm. More details about the computation of the fixed effects can be found in Hou & Milsom (2021). The effective distribution of firm productivity is plotted in figure 7. There is precedence for interpreting these fixed effects as informative about productivity (e.g. in Card et al. (2018)); Hou

point score used in this paper is a system used to encode students' performance in the range of exams into a numerical point score.

[38]In Cunha et al. (2005); Cunha & Heckman (2016), Cunha and Heckman discuss an empirical strategy for estimating workers' information sets using theoretical restrictions from the permanent income hypothesis and data on consumption. Another approach which is theoretically less onerous is to ask respondents in surveys about their expectations over outcomes, and to compare it against realised outcomes in the future. Unfortunately, these strategies are not available with the data I have.

& Milsom (2021) document that these fixed effects are correlated with firm revenue per worker.

Figure 7: Productivity distribution in empirical exercise



(a) Productivity density

(b) Productivity cumulative density

Panel 7a plots the density of productivity used in the empirical exercise, while panel 7b plots the cumulative density function of productivity used in the empirical exercise.

5 Empirical Results

Table 5 summarises the headline parameter estimates that I obtained from my estimation procedure. These results are conditional on a non-parametrically estimated probability of investment in education function conditional on θ, $\hat{P}(\theta)$.

Table 5: Parameter Estimates

No.	Parameter	Notation	Value	SE
Stage 1				
1	Signal noise	σ_ε	2.92	0.0000846
2	Skill return to education	δ	0.364	0.00269
3	Joint output function scale	q	7.83	0.172
4	Joint output function - exponent on s	γ_1	0.344	0.00713
5	Joint output function - exponent on y	γ_2	0.0318	0.0226
6	Minimum wage	w_0	4.46	0.0225
Stage 2				
7	Location parameter of het pref for educ relative to no educ	κ	-11.0	0.00447
8	Scale parameter of het pref for educ/no educ	ξ	11.6	0.00417

In subsection 5.1, I discuss the estimation of the variance of signal noise, the parameter governing the education technology function, the parameters of the joint output function, and the constant of the wage function ($\hat{\delta}, \hat{q}, \hat{\gamma}_1, \hat{\gamma}_2, \hat{w}_0, \hat{\sigma}_\varepsilon$). I loosely refer to this group of parameters as labour market parameters. In subsection 5.2, I discuss the empirical education investment function in θ, $\hat{P}(\theta)$, as well as the estimation of κ and ξ, which I refer to as the choice parameters.

5.1 Estimated labour market parameters and goodness-of-fit

The main feature of these estimates is that the estimate of σ_ε is large, implying that the correlation between actual ability and observed grades is 0.324. As I shall show in the subsequent welfare analysis, this implies that there will be substantial scope for actual returns to education to differ from students' expected returns to education.

The other parameters are not as naturally interpretable. Instead, to evaluate how well the model performs in fitting the data, I analyse how well the simulated moments match the actual moments of the data. Table 6 presents the means and variances of log wages, overall as well as conditional on whether the worker has a degree, for both the actual data and for the simulated outcomes from my model. The simulated means match the actual moments very well. The simulated variances are still within the confidence intervals of the estimated data variances, although my model implies that the log wage variance of graduates is higher than the log wage variance of non-graduates. This is not what I find in this dataset, but is in line with evidence from other papers like Lemieux (2006).

Finally, I plot the R^2 of a regression of log earnings on grades for the subsets of graduates and non-graduates. In my model, if grades are perfect signals of ability, the share of variation captured by grades should be 1, and if grades are non-informative about ability at all, the share of variation captured should be 0. The data suggests that the R^2 of this regression is 0.125 for graduates and 0.0659 for non-graduates, while my model predicts a value of around 0.09 for both groups. While these simulations lie within the 95% confidence intervals, it seems likely that the lower R^2 for non-graduates than graduates is a feature that is not captured by my model (e.g. through multi-dimensional skills).

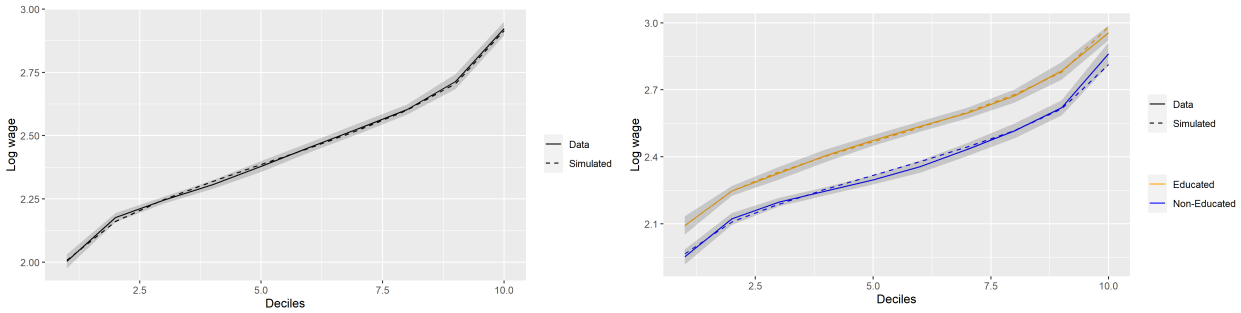
Table 6: Targeted log wage moments, overall and conditional on degree

Statistic	Data	Conf Interval	Simulated
Mean log wage	2.43	[2.42,2.45]	2.43
Mean log wage (e=1)	2.51	[2.49,2.53]	2.51
Mean log wage (e=0)	2.36	[2.34,2.38]	2.36
Variance log wage	0.0688	[0.0634,0.0743]	0.0675
Variance log wage (e=1)	0.062302	[0.0553,0.0693]	0.0652
Variance log wage (e=0)	0.064314	[0.0566,0.0720]	0.0588
R^2 of regressing log wages on grades (e=1)	0.125	[0.0737,0.176]	0.0917
R^2 of regressing log wages on grades (e=0)	0.0659	[0.0251,0.107]	0.0964

Panel 8 plots the actual and simulated log-wage quantiles, while panel 8b plots the quantiles conditional on education. My model is able to achieve a good match to the wage quantiles, both overall and conditional on education status.

Furthermore, my model is able to capture the interdepence between grades, education status and the individual's wage. Figure 9 plots the mean wage conditional on a worker's signal quintile and whether they are educated. Again, my model fits the data relatively well. In appendix D, I plot the fit of the simulated model to means of log earnings within four quartiles conditional on signal quintile and degree status, and show that the model is able to fit even those more detailed moments well.

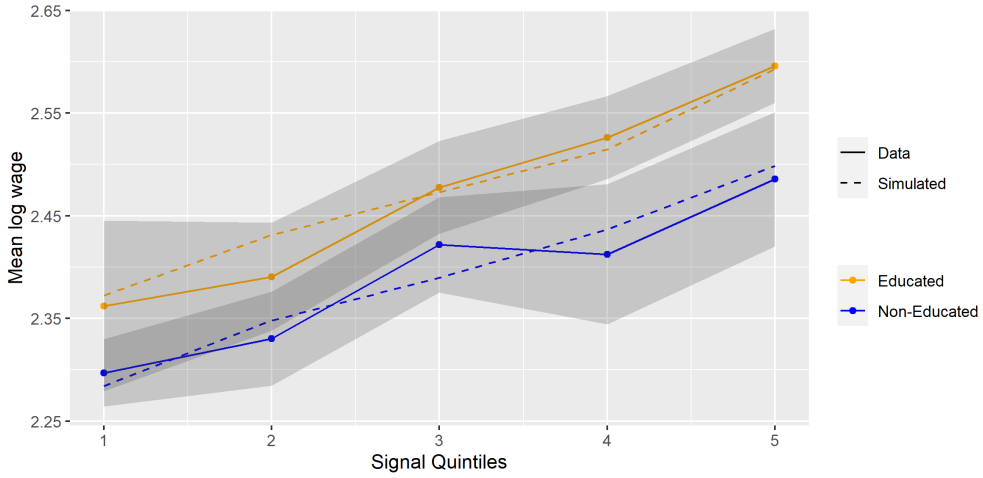
Figure 8: Actual and simulated wage quantiles



(a) Simulated and actual wage quantiles (b) Simulated and actual wage quantiles cond. on educ

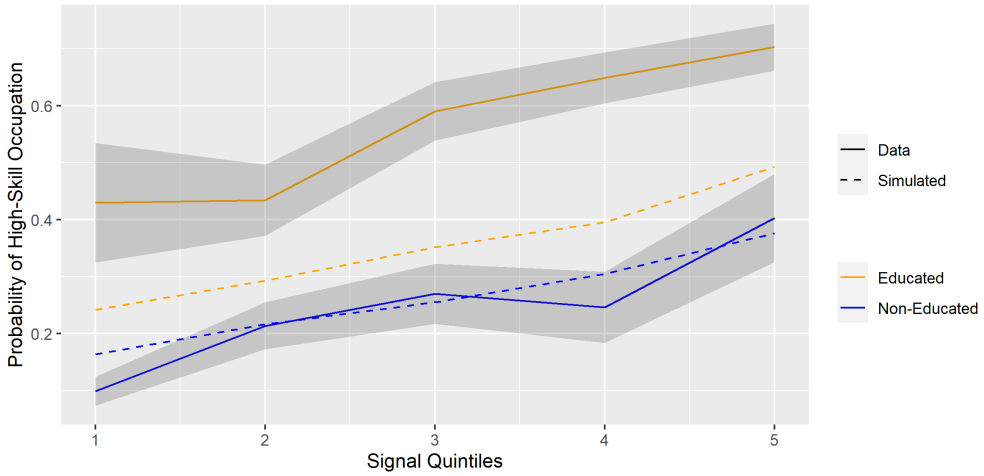
Panel 8a plots the mean log wage within each wage decile in the data in the solid line and the simulated equivalents in the dashed line. Panel 8b plots the actual and simulated wages for educated workers (in orange) and non-educated workers (in blue). The shaded ribbon plots the confidence intervals of the estimates from the data.

Figure 9: Actual and simulated mean wages conditional on education and signal quintile



This figure plots the mean log hourly earnings conditional on education (non-graduates in blue and graduates in orange) and five grade quintiles for both the actual data moments (in the solid line) and the simulated data moments (in the dashed line). The shaded ribbon represents the 95% confidence interval for the data moments.

Figure 10: Actual and simulated probabilities of matching to high-skill occupations



This figure plots the probability of being in a high-skill occupation conditional on education (non-graduates in blue and graduates in orange) and signal quintiles for both the actual data moments (in the solid line) and the simulated data moments (in the dashed line). The ribbons give the 95% confidence intervals for the data moments.

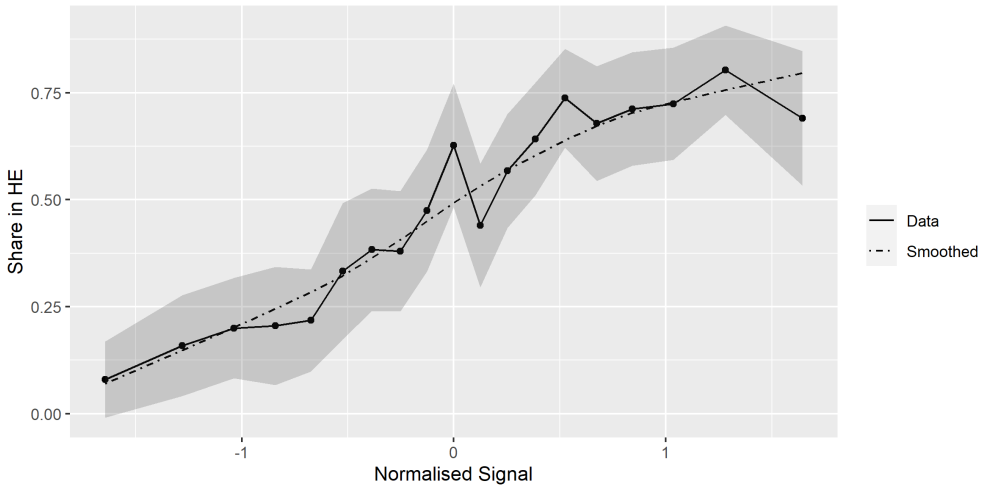
5.1.1 Untargeted moments

Finally, my model makes predictions about the probability of a worker being in particular occupations, conditional on their signal quintile and education. Given a worker's rank in the skill distribution post-education, it specifies a job productivity that the worker matches to. Conditional on this job productivity, a worker randomly matches to an occupation, with the probability of being in occupation k conditional on job productivity s given by equation 19. I then classify these SOC00 3 digit occupations into high skill/low skill categories based on a classification used by the Home Office for immigration guidance purposes.

To analyse how well my model fits the data on the matching of workers to occupations, I plot the share of workers in high-skill occupations, conditional on degree and five signal quintiles, in figure 10. While my model is imperfect in capturing the levels of graduates in high-skill occupations, it captures the qualitative facts that there are substantial shares of workers in both education groups in high-skill jobs and that the share of workers in high-skill jobs increases with the observed signal for both groups. This lends support to the modelling approach that there is an underlying skill index that determines both wages and the occupational match, of which the degree and grades are only indications. However, the significant mismatch in the level of graduates in high-skill jobs suggests that my characterisation does not fully capture sorting into occupations.

An interesting element to this failure to match the high-skill probability of graduates is that while my model underestimates the probability of working in a high-skill occupation for graduates and overestimates the probability for non-graduates, it captures the wages of both graduates and non-graduates quite well. This suggests that the failure in modelling the match to occupations does not affect the model's specification of wages. This could imply that while the model specifies the right productivity of the job that both graduates and non-graduates match to, it does not correctly

Figure 11: Estimated share of workers with degrees conditional on θ



This figure plots the share of workers with degrees conditional on θ at 19 points across the signal distribution. These shares were calculated by averaging the share in a 0.05 interval around each point, corresponding to the 5th to 95th quantiles of θ. The confidence intervals use standard errors calculated by the boostrap. The blue mixed line represents a smoothed approximation of the points, calculated by locally weighted smoothing, which is used to compute the distribution of skill in the estimation.

specify how graduates choose the occupation conditional on the matched job productivity. This could be indicative that graduates may have non-pecuniary preferences for occupations classified as high-skill (e.g. if they are white collar occupations), even when they are remunerated at the same rate as low-skilled occupations.

5.2 Estimated choice parameters

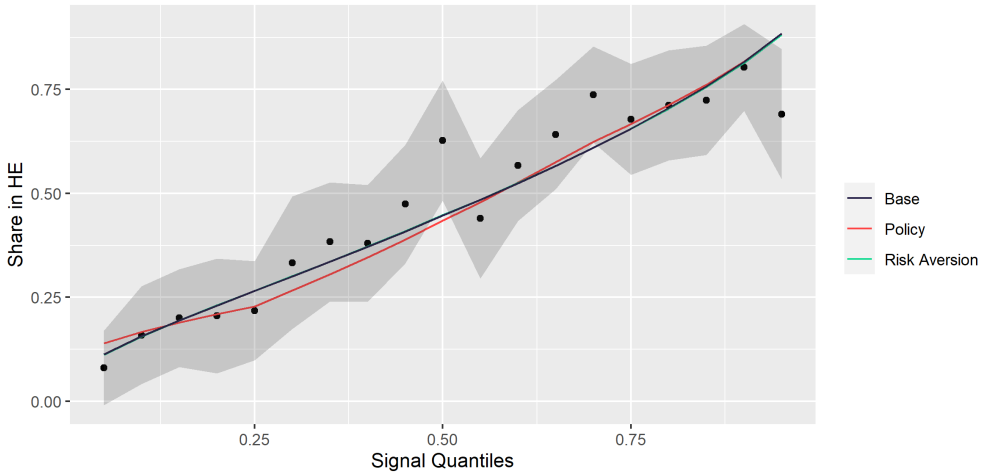
Figure 11 plots the shares of students who chose to go to higher education within a local region around 19 points across the signal distribution[39]. I compute standard errors by the bootstrap. To summarise this relation between probability of choosing higher education and θ as a function, I smooth the observed points using locally estimated scatterplot smoothing. As expected, workers are more likely to invest in education if they receive a higher grade, with choice probabilities ranging from 8.0% to 80.3% over the grade distribution.

I follow the procedure set out in section 4.2.2 to estimate the parameters κ_1, κ_2 and ξ. The parameters are summarised in table 5.The fit of the conditional choice probabilities to the observed data is shown in figure 12 below, in the black line. The simple base choice model achieves a good fit to most points of the data, besides the point at the median signal and the point at the 95 signal percentile.

I now turn to considering two extensions to this simple model of education choice. First, in this

[39]For example, the first point summarises the share in higher education from the 2.5th to 7.5th quantile of the signal.

Figure 12: Predicted probabilities of investing education conditional on θ



This figure plots the share of workers with degrees conditional on θ at 19 points across the signal distribution. The confidence intervals (in dashed black lines) use standard errors calculated by the boostrap. The mixed line represents a smoothed approximation of the points, calculated by locally weighted smoothing, which is used to compute the distribution of skill in the estimation. The blue solid line plots the predicted probability of investment in higher education conditional on θ considering the income-contingent loan policy, and the red solid line plots the predicted probability of higher education choice conditional on θ neglecting the income-contingent loan policy, using estimates in table 7.

paper, I have abstracted from risk aversion in my discussion. As a robustness check, I then estimate a version of the choice set-up when workers have CRRA utility as follows, where w denotes their wage:

$$u(w) = \begin{cases} \frac{w^{1-\zeta}-1}{1-\zeta} \text{ if } \zeta \in [0,1) \cup (1,\infty) \\ \log(w) \text{ if } \zeta = 1 \end{cases} \tag{30}$$

The risk aversion parameter is identified by the relative concavity of the log-odds function in θ relative to the expected returns function to θ. If the log-odds function is highly concave when the expected returns function is convex, then, this suggests that workers' choices are driven by a degree of risk aversion. The coefficient of risk aversion parameter is estimated as part of the GMM estimation described before. The results are presented in column 3 of table 7. I estimate a very low degree of risk aversion, consistent with the fairly non-concave shape of the conditional choice probabilities conditional on grades. Although the minimised value is lower than it was for the base model, the improvements are fairly minor. I plot the predicted probabilities of the model with CRRA utility in green in figure 12. The line coincides basically perfectly with the black line until the end, where it does better in fitting the points at the top end of the signal distribution.

Second, students typically pay for higher education in the UK with an income-contingent loan, where they are not required to repay the loan in full if their income is not sufficiently high. This offers a significant level of insurance to workers and is very generous for workers with lower grades, who have to pay a much lower expected tuition. This changes the incentives that workers face; in appendix E, I describe further how the implementation of the tuition policy changes workers' incentives under the estimated parameters. I estimate a set of parameters where workers face

Table 7: Choice Parameters under Various Specifications

Parameter	Base	Policy	Risk Aversion
ξ	11.5	24.2	17.4
	(0.00447)	(0.00948)	(0.440)
κ	-11.0	-16.2	-12.0
	(0.00417)	(0.00626)	(0.0717)
ζ			0.127
			(0.00765)
Minimised Value	600.	719.	582.

expected returns that would prevail under the tuition policy. These estimates are presented in column 2 of table 7 above. The implied choice probabilities are also plotted on figure 12 in a red line. The minimised value suggests that this model is a substantially worse fit to the moments, driven by a bad fit for the middle of the signal distribution. This may suggest that students do not fully internalise the generosity of the income-contingent loan, especially for workers with lower grades.

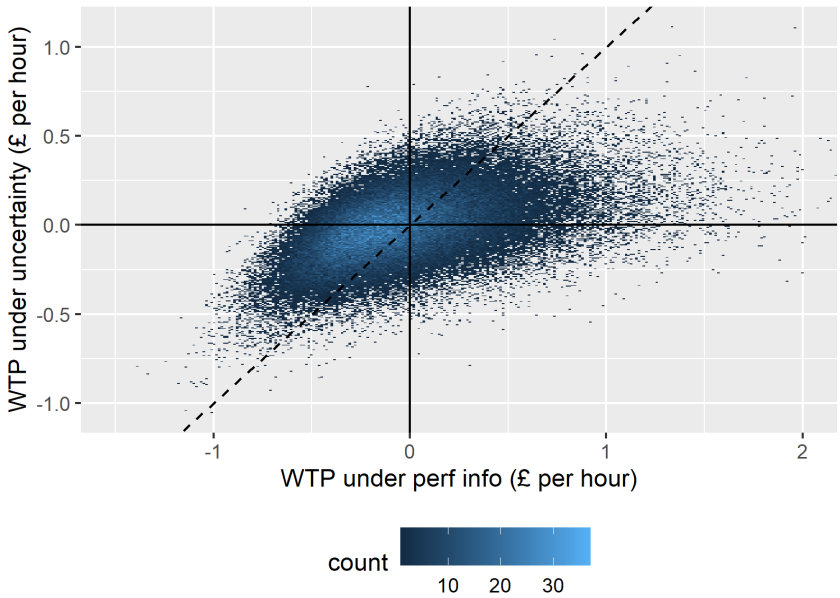
6 Welfare

6.1 Welfare in terms of willingness to pay for higher education

As section 3.2.3 describes, my model implies that imperfect information about ability can lead workers to make ex-ante optimal decisions which nevertheless lead to ex-post regrettable outcomes. Do these frictions lead to significant utility losses? The structural estimation of my model, described in section 4, allows me to answer this question by simulating the model and analysing the welfare properties of the simulated agents.

In this section, I use as a natural measure of welfare, the difference between the value a worker places on investing in education and not investing in education measured in terms of pounds per hour earnings (at age 30). Intuitively, this measure can be thought of as the willingness to pay (WTP) for an investment in education; that is the amount in pounds per hour at age 30 equivalent that an uneducated worker is willing to give up to be in higher education (on top of any other tuition or psychic costs already captured by the choice parameters). Conversely, a worker with higher education has to be compensated up to their WTP to be induced out of higher education (on top of being refunded any tuition or psychic costs). A worker's willingness to pay would depend on their net preferences for education $\Delta\eta$ and either their ability a under perfect information or their signal θ under imperfect information. A worker with a positive willingness to pay would select

Figure 13: Simulated Joint density of actual and expected WTP for higher ed



This figure plots the distribution of actual (conditional on unobserved ability) and expected (conditional on the signal) willingness to pay for higher education in a simulated dataset generated from the model and the estimated parameters. Each square represents a set of binned values for WTP under perfect info and imperfect info, and the colour of the square denotes the count of points within the bin (with N = 100,000). The dashed line denotes the 45 degree line.

into higher education, and vice versa[40].

With the estimated parameters of my model described in the previous section, I can simulate the distribution of WTP under both perfect information and imperfect information. Figure 13 plots the joint density of WTP under perfect and imperfect information with a heat map, with a lighter blue colour indicating a higher density and a darker blue colour indicating a lower density. The dashed line is the 45 degree line; being to the left of the 45 degree line means that a worker has a greater WTP with imperfect information than under perfect information. This means that the information imperfect leads them to overvalue attending higher education. Being to the right of the 45 degree line implies that the worker undervalues attending higher education as their WTP under imperfect information is small than their WTP under perfect information. In general, not being on the 45 degree line implies that there is a gap between a worker's willingness to pay under perfect and imperfect information.

The axes also have an important interpretation in the diagram. If a point is above the x-axis,

[40]The WTP under perfect information is the difference between the wages under education and no education, and the difference between heterogeneous preferences for education and no education. The WTP under imperfect information is the difference between the value functions, specified in equation 9.

$$\text{WTP}_{\text{perf}}(a, \Delta\eta) = \beta\left(w(s(a,1)) - w(s(a,0))\right) + \kappa + \Delta(\eta) \tag{31}$$

$$\text{WTP}_{\text{imperf}}(\theta, \Delta\eta) = V(\theta, 1) - V(\theta, 0) \tag{32}$$

this implies that the worker the point denotes would invest in higher education under imperfect information, and vice versa. If a point is to the right of y-axis, then the worker would invest in higher education under perfect information. The diagram is thus divided naturally into four quadrants. Being in the top right quadrant implies that the worker would have invested in HE under both perfect information and imperfect information, while being in the bottom left quadrant implies that the worker would not have invested in HE under both information scenarios. Being in either of these quadrants and not being on the 45 degree line implies that the worker has a mistaken willingness to pay due to the information friction, but not one sufficiently wrong that it leads their education choice to diverge under perfect and imperfect information. This explains why the loss distributions are censored to the left; small 'mistakes' in the valuation of WTP often do not induce workers to change their education choice, and thus do not cause any utility loss.

On the other hand, being in the top left quadrant implies that a worker has a positive WTP under imperfect information but a negative WTP under perfect information, meaning that they would select into higher education despite receiving a negative utility return. Being in the bottom right quadrant implies the opposite, that a worker has a negative WTP under imperfect information and a positive WTP under perfect information; these workers end up not investing in higher education when they would have received a positive utility return. Workers in the top-left quadrant are over-educated, and workers in the bottom-right quadrant are under-educated, under the terminology in section 3.2.3.

In table 8, I summarise the shares of the simulated population in each of the four quadrants: workers who invest and would benefit, those who invest but would not benefit (over-education), those who do not invest but would have benefited (under-education), and those who do not invest and would not have benefited. The noise around individuals' ability implies significant levels of

Table 8: Shares of workers by education choice and whether they receive a net benefit

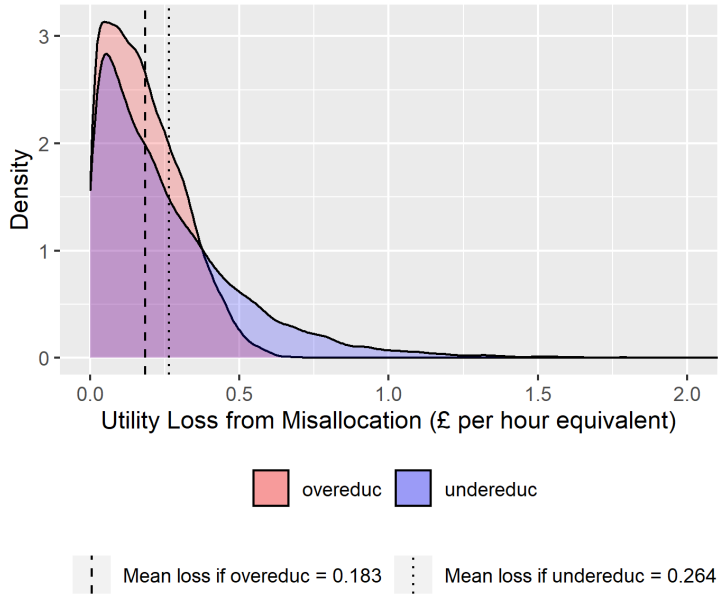
	Would invest	Would not invest	Total
Would benefit	0.286	0.147	0.433
Would not benefit	0.182	0.385	0.567
Total	0.468	0.532	1.00

This table summarises the simulated population by whether they would invest in higher education and whether they would experience a net positive utility from doing so ex-post.

misallocation; 39.1% of graduates are 'over-educated' in the sense that they go to university despite a net negative return and 27.6% of non-graduates are 'under-educated' in the sense that they could have benefited from university if they had gone. In total, 32.9% of the population would have been better off if they made a different education choice.

How costly is the misallocation of higher education? In figure 14, I plot the distribution of welfare losses in terms of willingness to pay, conditional on being over-educated and being under-educated. In general, the losses are small and right-skewed, in large part because they are censored to the left. Small deviations between willingness to pay under the perfect information and imperfect information scenarios may not be sufficient to induce the worker to change their education choice, and thus do not lead to misallocation. The mean utility loss in expected return terms comes out to be £0.183 per hour (£323.35 per year) conditional on being over-educated, and £0.264 per hour (£465.22 per year) conditional on being under-educated. These numbers are approximately

Figure 14: Distribution of welfare loss due to misallocation



This figure plots the distribution of the utility loss from misallocation, conditional on being over- and under-educated, from a simulated dataset with the estimated parameters from the structural estimation exercise. The results are presented in terms of willingness to pay, in pounds per hour terms.

2.5% and 3.18% of the average wage in the dataset. The mean loss from being over-educated is smaller than the mean loss from being under-educated, which suggests that while over-education is a more salient problem because one can observe graduates matched to low-skill jobs in the data, under-education is on average more costly to those who experience it. This is driven by long right tail of utility loss conditional on being under-educated; table 9 below summarises the quantiles of the utility loss in £ per hour money-metric terms. Intuitively, this is driven by the convexity of returns to education with regards to ability. An under-educated worker is a high ability worker who mistakenly thinks that they are low ability, and the counterfactual is that they may miss out on a large return higher education as a result. On the other hand, a over-educated worker is a low ability worker who mistakenly thinks that he is high ability; their counterfactual earnings without education investment is also relatively low, and the degree of mismatch is likely to be low as a result. This asymmetry of potential gains and losses lead workers to err in favour of investing in education; thus, the share of over-educated workers exceeds the share of under-educated workers.

Table 9: Quantiles of WTP loss due to misallocation (£ per hour equivalent)

	25%	50%	75%	90%
Overeducated	0.0797	0.163	0.269	0.365
Undereducated	0.0829	0.194	0.368	0.586

This table summarises four quantiles of utility loss for over- and under-educated workers from a simulated dataset with the estimated parameters from the structural estimation exercise. The results are presented in terms of willingness to pay, in pounds per hour terms.

There are no directly analogous results to mine in the literature, but this result falls within a plausible range implied by other studies on the UK higher education system. Waltmann et al. (2020) uses an administrative dataset including tax data and education records to estimate lifetime returns to education and find that about 20% of students would experience negative lifetime pecuniary returns to education. They also find that the government is expected to make a loss on the degrees of 40% of men and 50% of women due to the UK's income contingent loan scheme. Thus, my finding that 39.1% of graduates do not benefit from university in aggregate falls within the ballpark of their results.

6.2 Welfare in terms of value for money

Although willingness to pay is a natural concept of welfare to focus on, the policy debate, especially in the UK, has focused on the concept of 'value for money' or 'earnings potential'[41]. Intuitively, this seems to refer to the exclusion of non-pecuniary motivations for attending university, roughly implying that students should attend higher education if their decision passes a cost-benefit analysis considering only pecuniary returns to university net of pecuniary costs.

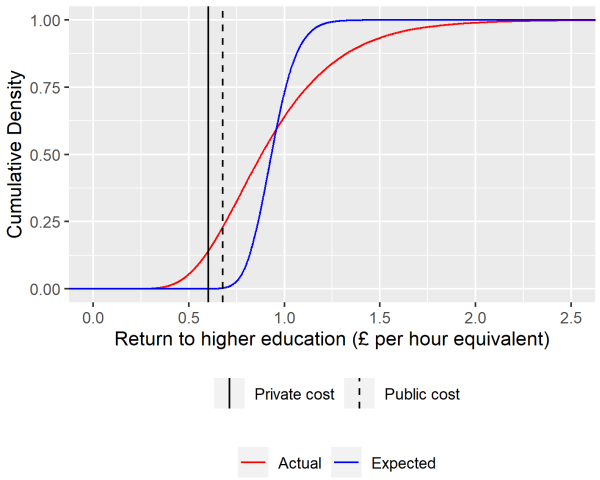
In this paper, I interpret value for money roughly as the pecuniary wage return to attending university exceeding the pecuniary costs of doing so, excluding the consumption value that workers may receive from higher education. In this sense, a worker may attend university in my model despite not receiving value for money for doing so because of (1) they have a positive non-pecuniary preference for attending university, or (2) because the information friction leads them to believe that they will receive a positive pecuniary return despite not actually doing so.

Figure 15 plots the distribution of actual and expected wage returns to higher education in a simulated dataset with 20000 workers as cumulative density functions. A striking feature of these distributions is that expected returns (conditional on the grade signal) are much less variable than average returns (conditional on ability). This is to be expected, especially given the substantial noise with which grades correlates to actual underlying ability.

To conduct a full cost-benefit analysis, I need to specify the pecuniary cost of higher education. I consider two alternatives. First, I take an estimate of the cost of providing higher education from Belfield et al. (2018b), £25,000, and amortise this over 30 years at a 2.5% interest rate. This calculation implies the per-year payment to be £1194.441 per year, or £0.679 per hour (assuming a 40 hour work week for 44 weeks). Second, I take the total student loan that students take out on average, similar amortised over 30 years at a 2.5% interest rate. The average loan is £22,230, and is taken from Bolton (2019). The implied per year per hour equivalent cost is 0.599. Two vertical lines are plotted on the figure, representing two cost benchmarks. The dashed vertical line presents the cost of providing a higher education course during the period, including the private cost borne by the student and the public cost borne by the Treasury in terms of teaching grants given to universities. The solid vertical line represents the cost borne directly by workers, including tuition costs and living costs covered under the maintenance loan.

[41] Earnings potential was discussed by the former Chancellor of the UK, and Conservative party leader hopeful Rishi Sunak, who pledged in the 2022 leadership contest, that he would crack down on university courses, assessing them "through their drop-out rates, numbers in graduate jobs and salary thresholds, with exceptions for nursing and other courses with high social value". See https://www.theguardian.com/politics/2022/aug/07/rishi-sunak-vows-to-end-low-earning-degrees-in-post-16-education-shake-up.

Figure 15: Cumulative density of simulated actual and expected returns to higher ed



This figure plots the distribution of actual (conditional on unobserved ability) and expected returns (conditional on the signal) in a simulated dataset generated from the model and the estimated parameters. The red line plots the cumulative density of actual returns and the blue line plots the cumulative density of expected returns. The two vertical lines represent two cost benchmarks.

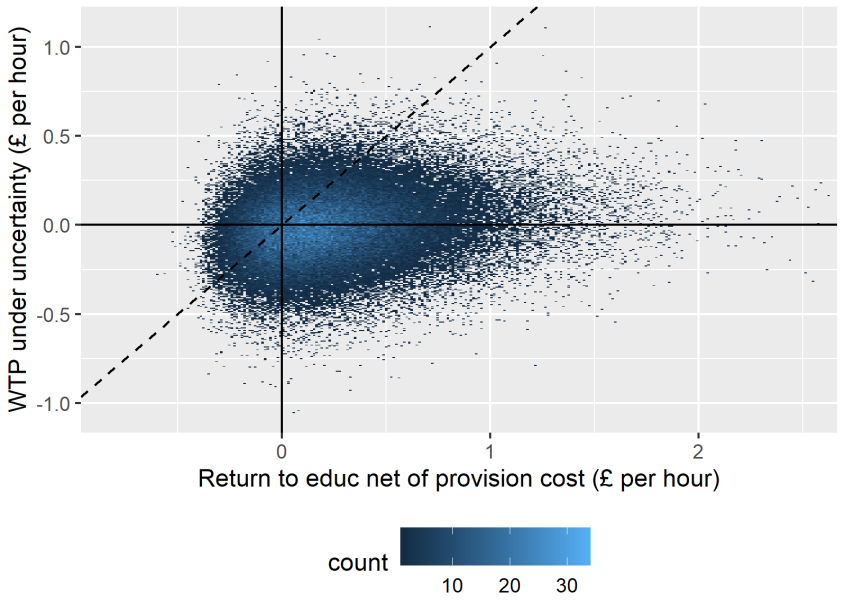
The figure implies that most workers in this cohort would receive a benefit net of either cost, in part because this cost, especially spread over a lifetime, is not particularly high. Only 14.3% of workers would not have received a return exceeding their private cost of education, and only 23.2% of workers would not have received a return exceeding the total cost of provision. Only 0.3% workers would not have had an expected income above the total cost of provision, implying that if workers considered only the pecuniary accounting of costs and benefits, more workers would have selected into higher education[42]. From this point of view, higher education was an extremely valuable investment for workers relative to its cost for the studied cohorts[43].

In practice, only 46.8% of workers do invest in education. Figure 16 plots the joint density of the return to education net of the cost of provision and the expected willingness to pay for higher education. The figure shows that going by the returns net of the cost of provision, most graduates receive a positive return to higher education, and only 18.3% of graduates would not have a positive return. On the other hand, most non-graduates are under-investing in education, and 72.5% of non-graduates could have benefited. In my model, this misallocation can be generated by two channels. First, the noisiness with which grades measures the actual labour market performance leads to an information imperfection which impedes choice. Second, heterogeneous net preferences, such as a high consumption value for university or particularly low costs due to family endowments, may lead some workers to find higher education value for money even with a lower labour market return. In table 10, I first summarise the share of workers who would invest in higher education, and then divide the population into workers who invest and receive a return greater than the mean preference benchmark, those who invest but receive a return lower than the benchmark, those who do not

[42]Note however that if all workers selected into higher education, the prospective return under the resulting skill distribution might not sustain all of them selecting into higher education after wages fall.

[43]Note that the per-year tuition cap was 3000 pounds for the studied cohorts. Since then, this cap has tripled to 9250 pounds per year in 2021.

Figure 16: Simulated joint density of actual return net of cost of provision and expected WTP for higher ed



This figure plots the distribution of actual (conditional on unobserved ability) return to higher education minus the cost of provision and expected (conditional on the signal) willingness to pay for higher education in a simulated dataset generated from the model and the estimated parameters. Each square represents a set of binned values for WTP under perfect info and imperfect info, and the colour of the square denotes the count of points within the bin (with N = 100,000). The dashed line denotes the 45 degree line.

invest but would have received a return greater than the benchmark, and those who do not invest and would not have received a return higher than the benchmark. Then, I perform a counterfactual exercise, where I isolate each of these two channels in turn to analyse which factor contributes more to the mismatch between returns and choice (keeping the skill environment constant).

Table 10: Breakdown of investment choices and whether return exceeds mean costs benchmark, under four counterfactual scenarios

	Het + Noise channels	Het channel only	Noise channel only	No channels
Share choosing degree	0.468	0.433	0.997	0.768
Would invest, VfM	0.382	0.429	0.767	0.768
Would invest, not VfM	0.0856	0.004	0.230	0
No invest, VfM	0.386	0.340	0.00108	0
No invest/Not VfM	0.146	0.228	0.00187	0.232

This table summarises the share of workers in higher education, and a breakdown by whether they would experience "sufficiently" high returns to higher education. I use as a benchmark for "sufficiently high" the level of returns that would cover the cost of providing higher education, and is computed to be £0.679 in this setting. This is denoted as VfM, short for "value for money". These shares are computed for four scenario, the base scenario with both uncertainty of returns and heterogeneous preferences, and three counterfactual scenarios alternative with heterogeneous preferences only, uncertain returns, and one with neither.

First, column 1 suggests that there is significant mismatch between wage returns in the system; of the 46.8% of workers who would choose university in the model, 38.2 percentage points, or 81.6% would generate a return greater than the cost of provision benchmark. Furthermore, 38.6 percentage points, or 72.5% of non-graduates, would have received a return higher than the mean preference benchmark but do not end up in higher education; this suggests that too few workers invest in education from a value for money point of view. Heterogeneous non-pecuniary preferences for higher education seems to drive under-investment in higher education, since even knowing the actual returns to higher education, 44.3% who would have received a net positive return would not have attended higher education. On the other hand, the uncertainty about true returns seems to drive over-education, leading workers who might not have benefited to invest in higher education. Thus, both channels are important in different ways in driving the mismatch between the net wage returns to education and workers' choices in the cohort studied.

7 Policy Relevance

Given the large degree of higher education expansion in the UK, many policymakers are worried that the UK has reached "peak graduates", and that the marginal graduate is not benefiting from higher education[44]. This view is summarised in a July 2020 speech by the Universities Minister,

[44] Although this paper is based on the UK context, underemployment of graduates is a concern throughout the world. In the US, Freeman (1976) and Caplan (2018) are two popular economics books which make the argument that the level of higher education is inefficiently high. The Federal Reserve Bank of New York also finds that about 30% of all college graduates are underemployed (statistics are available at https://www.newyorkfed.org/research/college-labor-market/index#/underemployment).

who argued that[45]:

> For decades we have been recruiting too many young people on to courses that do nothing to improve their life chances or help with their career goals....Since 2004, there has been too much focus on getting students through the door, and not enough focus on how many drop out, or how many go on to graduate jobs.

This concern has led policymakers to propose policy instruments which reduce the numbers of students going to university. For example, in a policy consultation in 2022, the UK government considered policies like student number caps and minimum grade thresholds for access to higher education finance as a response to the high cost to the Treasury of paying for the income-contingent loan system. The government has also effectively reduced the degree of the subsidy by increasing the loan repayment period from 30 years to 40 years[46].

However, if overeducation is not driven primarily by workers investing in education even when they expect low returns because of e.g. the consumption value of a university experience, but because of uncertainty about their likely returns as this paper has argued, then it is not clear that policies to change the aggregate preference for education will have positive welfare effects. In particular, policies that discourage university attendance would increase the share of workers who are under-educated and decrease the share of workers who are over-educated, and vice versa, with the aggregate welfare effects being ambiguous. Furthermore, reducing the share of workers in higher education can increase the wages of workers conditional on their education status by reducing the average skill of workers in the economy, changing the welfare returns through such general equilibrium effects.

To analyse the impact of policies to reduce the take-up of higher education, I simulate two counterfactuals, in which policy makers (1) institute a flat fee for graduates (e.g. by increasing tuition fees or instituting a flat graduate tax), and (2) institute a flat subsidy for graduates (e.g. by reducing tuition fees or instituting universal graduate subsidies). I make these changes revenue-neutral by instituting a flat lump-sum tax or subsidy on each worker equal to the average subsidy/tax paid to workers in the population; because these changes would apply regardless of their education status, the compensation does not affect the incentives to invest in education. Denote the flat tax by τ, where a negative value implies a subsidy. The policy then modifies the value function, equation 9, as follows:

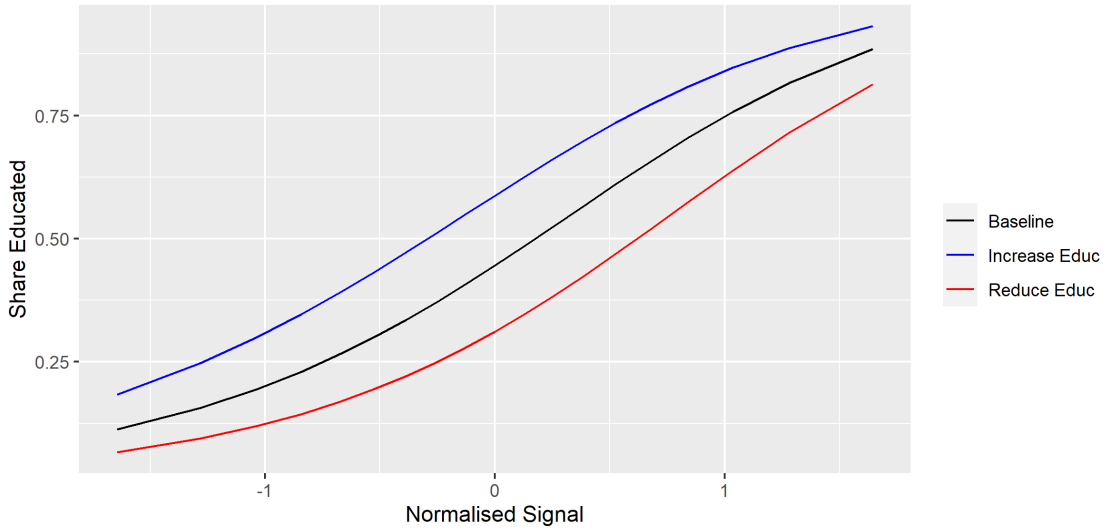
$$V^*(\theta, e) = \left((\kappa - \underbrace{\tau}_{\text{Graduate tax/subsidy}}) \times e \right) + \eta(e) + \underbrace{\left[\int P(\theta) d\theta \right] \tau}_{\text{Compensation}} + \beta E\{w(s(a,e))|\theta\} \tag{33}$$

The tax/subsidy is thus a revenue-neutral method of shifting the aggregate preference for higher education, and can be equivalently thought of as a tool to achieve an optimal level of higher education in the economy.

[45] From a speech on 1 July 2020 at the NEON summit on widening access and mobility by then-Universities Minister, Michelle Donelan. The transcript was accessed online on 30 May 2022 at https://www.gov.uk/government/speeches/universities-minister-calls-for-true-social-mobility.

[46] https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1057091/HE_reform_commanc paper-web_version.pdf

Figure 17: Distribution of welfare loss due to misallocation



This figure plots the equilibrium share of workers in education under two counterfactuals. The blue line plots the situation in which education preferences increase by the equivalent of 5p per hour, while the red line plots a situation in which education preferences decrease by the equivalent of 5p per hour.

I simulate the impact of a flat subsidy or tax equivalent to 5p per hour to graduate earnings[47]. These policies have the first-order effect of (1) reducing or (2) increasing the shares of workers who choose to invest in higher education for all values of the grade signal[48]. Figure 17 summarises the equilibrium effect on higher education choice of the counterfactual policies. As expected, a graduate subsidy increases the share of workers selecting into education, while a graduate tax decreases this share.

Table 11 below describes wages and higher education choice in each of the two counter-factual scenarios. The average wage increases by 0.7% with the graduate subsidy, and decreases with the graduate tax by 0.8%, driven entirely by more workers being in higher education. Wages conditional on education status followed the reverse pattern. In the scenario with the graduate subsidy, the average wage of graduates falls by 1.1%, and the average wage of non-graduates falls by 0.8%. In the graduate tax counterfactual, the average wage of graduates increases by 1.0%, and the average wage of non-graduates increases by 0.9%. An intuitive way to understand this is with the graduate subsidy, more workers get educated and the supply of skill increases in the economy. As a result, workers are less able to bargain for higher wages conditional on their skill because there are more alternatives. Similarly. in the graduate tax counterfactual, skilled workers are less common, and are able to command a greater wage for their skill level. In aggregate, the ratio between graduate and non-graduate wages decreases by 0.21 percentage points in the graduate subsidy scenario and

[47]This change appears small but is applied to hourly post-tax earnings at age 30 which is used as a proxy for lifetime earnings. The change is equivalent to about 0.4% of mean graduate earnings over the lifetime. It is equivalent to passing on the impact of a 10% increase or decrease in the cost of providing higher education provided by Belfield et al. (2018b), amortised over 45 years at a 2% interest rate.

[48]In the language of the model, this is equivalent to increasing or decreasing κ in equation 9. This is also functionally the same as policies to increase or reduce aggregate preferences for higher education.

increases by 0.11 percentage points in the graduate tax scenario. Conversely, firm profits increase when the supply of skill in the economy increases due to more education, and decrease when the number of graduates decreases. In a sense, firms in the model can free-ride on workers' higher education decisions, both because they get access to more skilled workers and because a larger number of graduates increases their bargaining power to lower worker wages. In the short run, the share of workers in skilled occupation remains largely constant[49], but as the supply of graduates increases, the share of both graduates and non-graduates in skilled occupations falls as a percentage.

Table 11: Descriptive Statistics in Counterfactual Scenarios

	Baseline	Graduate tax	Graduate subsidy
Share with degree	0.4662	0.3598	0.5773
Average hourly post-tax wage	11.38	11.29	11.46
Average hourly post-tax wage (graduates)	12.35	12.49	12.23
Average hourly post-tax wage (non-graduates)	10.59	10.67	10.49
University wage premium	0.1669	0.1715	0.1658
Average firm profits	-2.986	-2.996	-2.974
Share in skilled occs	30.97	30.89	30.86
Share grads in skilled occs	39.82	42.07	37.63
Share non-grads in skilled occs	23.24	24.60	21.59

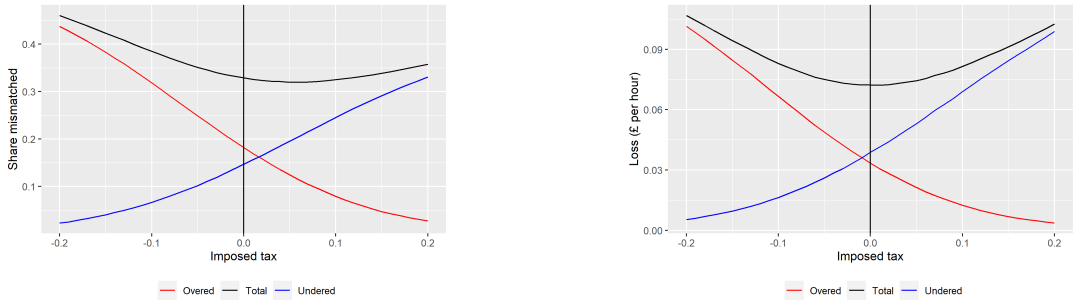
This table summarises some key descriptive statistics in each of three counterfactual simulations. The wage and wage premium are not inclusive of the simulated tax or subsidy. The firm profits are presented as negative as the constant in the joint output function is not identified and normalised to 0. This also explains why joint output is smaller than wages.

7.1 Reducing the incidence and welfare cost of mismatch

Is it possible to reduce the incidence and welfare cost of mismatch through the proposed compensated tax scheme? I simulate the welfare effects of a tax for a range of values between -0.2 and 0.2, and consider minimising the number of workers affected by mismatch, or the total loss in WTP terms of mismatch. Figure 18a plots the share of mismatched workers under various compensated tax levels, while figure 18b plots the total welfare loss in WTP terms under various such levels. The optimal policy if we are aiming to minimise those affected is a tax of 0.06 pounds per hour, reducing the share mismatched from 32.9% to 31.9% under the no-policy baseline. However, this increases the average loss of the mismatch by 4% from 0.0723 pounds per hour to 0.0755. Intuitively, this follows from the earlier observation that under-education is more costly on average than over-education, and thus there are more over-educated than under-educated workers under ex-ante utility maximisation. Thus, the optimal tax in this situation promotes costly under-education (from 14.7% to 20.5%) to reduce the share of over-educated workers (from 18.2% to 11.4%). The optimal tax to minimise the degree of welfare loss in WTP is 0.01 pounds per hour, which is close to the baseline.

[49]Some papers like Shephard & Sidibe (2019); Blundell et al. (2022) emphasise that the supply of skilled jobs may change as the supply of graduates increases. This dynamic is not present in my model, and represents scope for future extensions to my framework.

Figure 18: Degree of mismatch under different compensated tax levels



(a) Share affected

(b) WTP Loss

This figure plots the share affected by mismatch and the welfare loss due to the information friction under various levels of compensated tax. The welfare loss is expressed in willingness to pay. The black line plots total loss, and the blue and red lines respectively plot the loss due to undereducation and overeducation.

7.2 Maximising overall welfare and workers' welfare

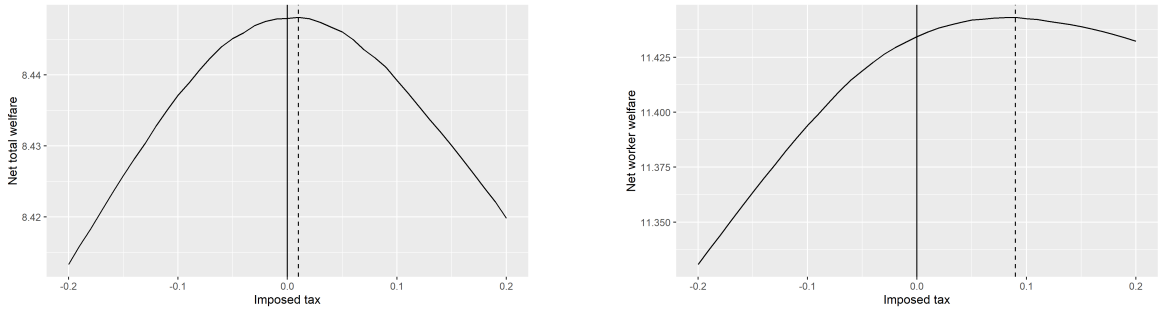
However, minimising the degree of mismatch is arguably not the relevant welfare metric; it may not matter if allocative inefficiencies are large if the economy features high profits and/or wages. Another natural metric is the sum of total utility in the economy, both from workers' wages and firm profits.

In general, subsidising higher education increases wages and profits by increasing the level of skill in the economy, while taxing higher education lowers wages and profits by doing the inverse, as previously discussed. However, the marginal increase in utility decreases as the number of workers in education increases. This is because, averaging over heterogeneous preferences for education, the marginal student when enrolment is low is likely to have higher returns than the marginal student when enrolment is high. Second, as the number of graduates increases, the supply of skill increases in the economy and allows firms to bid down the wages of all workers. Third, the marginal student is likely to have a lower net preference of higher education than workers who have already selected into higher education. Thus, while as the subsidy increases to the limit, the average profit and average wage in the economy increases, but at a decreasing rate. At some point, the marginal increase in utility is smaller than the average cost of higher education.

Figure 19a plots the total welfare in the economy for various compensated tax levels, summing up wages, profits and utility from attending higher education net of the cost of higher education (the share of workers with degrees multiplied by the unit cost of higher education). Figure 19b plots the total worker welfare conditional on the compensated tax, summing up wages, and utility from attending higher education but excluding profits. The resulting curves are all concave. The compensated tax that maximises total utility is a small tax equivalent to 1p per hour at age 30, while the compensated tax that maximises total worker welfare is a larger tax of 9p per hour at age 30.

If all workers are expected utility maximising in this model, why might the prevailing equilibrium nevertheless be suboptimal? There seems to be two sources of market failure in the model. First,

Figure 19: Net welfare under different compensated tax levels



(a) Net total welfare

(b) Net worker welfare

This figure plots the simulated counterfactual net welfare and net worker welfare under various levels of compensated tax. The welfare loss is expressed in willingness to pay. Figure 19a considers wages and utility from higher education, while figure 19b also includes average profits. The dashed vertical lines denote the utility-maximising value of the tax. The optimal tax is 0.01 in figure 19a and 0.09 in figure 19b.

because workers cannot appropriate all of the joint output in the bargaining process, they do not consider the increased profit that their education leads to for firms. They underinvest relative to the social optimum because they do not factor in the positive externality on firm profits of their educational investment. I call this the hold-up externality. For this reason, the optimal tax considering worker's welfare is higher than that considering aggregate welfare.

Second, the simulations suggest that workers' overinvest in higher education and that worker welfare would be higher if the extent of higher education were substantially reduced. This is likely to be a kind of congestion effect, where the higher level of education investment decreases the wages of similar workers around them; see that average graduate wages actually falls with greater educational investment. Such an externality arises because workers consider only their own welfare, but not the welfare of other workers in their economy and neglect the effect of their education decision on the wages and bargaining position of others. This externality leads to over-investment in higher education, and is offset by the hold-up externality. In my empirical application, the congestion externality dominates and implies that optimal policy is to slightly reduce the level of education from 46.2% to 44.6%. This is accomplished by imposing a tax equivalent to about 1 p per hour at age 30, or adding 368.37 pounds to the initial cost of education (amortised over 30 years at a 2.5% interest rate). If the policy-maker did not consider firm profits at all, the optimal policy would be to reduce higher education substantially, from 46.2% to 28.6%. This is accomplished by imposing a tax equivalent to 9 p per hour at age 30, or adding 3315.36 pounds to the initial cost of education (amortised over 30 years at a 2.5% interest rate).

A limitation of this analysis is that I have only considered policies in the set of a fully compensated taxes or subsidies that are entirely flat over the grade distribution. It may be possible to achieve greater efficiency by making the tax or subsidy conditional on student grades; for example, it is likely that the hold-up externality is greater on average for workers with high grades as they match to higher productivity jobs which generate greater profits. I leave this issue for future work.

8 Conclusion

There is often public concern about the level of investment in higher education, and that too many students go to university when they would not benefit from going. These concerns are frequently based on observations that many graduates end up working in occupations which typically do not require higher education to perform. However, causal studies of the wage returns to higher education typically find a substantial positive return, apparently contradicting these popular concerns. In this paper, I propose that both observations are consistent with a model in which returns to higher education are heterogeneous and uncertain at the point of university choice, and in which workers match to jobs in a Sattinger-style matching labour market after investing in education. Workers who are observed to work in non-graduate occupations despite getting a degree are interpreted in this framework as experiencing ex-post regret; they correctly invested in higher education with an ex ante expectation of a positive return but ended up with a negative return ex-post.

I estimate a parametric version of my proposed model on a pooled dataset of five cohorts born between 1988 and 1993, using data from the Understanding Society panel survey. I find that the degree of noise inherent in the signal is large; workers only have a signal with a 0.324 correlation to their underlying ability. I find that this leads 32.9% of workers in the economy to end up with ex-post regret. 18.2% of workers end up over-educated, that is, they invest in education but end up with a net negative utility return, and 14.7% are under-educated, missing out on a positive ex-post utility return. On average, the utility loss is in money-metric terms is £323.35 a year for over-educated workers, and £465.22 a year for under-educated workers. Being under-educated is on average more costly than being over-educated, which explains why more workers prefer to over-invest in education rather than vice versa.

I consider the welfare impacts of government policies where they impose compensated graduate taxes or subsidies to manipulate the returns to higher education. These policies relate to proposed policies by recent UK governments in response to a perceived glut of university graduates. I find that the optimal policy differs depending on the policy objectives of the government. I find that if the policy is to minimise the incidence or cost of mismatch in the economy, the government should instate a compensated tax to reduce the level of higher education attendance in the economy. For a welfare-maximising social planner, the equilibrium is in general efficient and trades off two kinds of externalities that work in opposite directions. I find that empirically, my estimates imply that the optimal policy would be to reduce the incidence of higher education by 1.6 percentage points by imposing a small graduate tax equivalent to 1 p per hour at age 30 (roughly equivalent to increasing the cost of education by 368 pounds).

There are many unanswered questions arising from my proposed model that I have not been able to address in this paper, which I leave for subsequent research. First, my empirical analysis has been based on observational outcomes data and an assumption of rational expectations. Recent research on the motivations behind higher education however has shifted towards using surveys about students' actual beliefs to decisions to select into higher education. It would be interesting to use those methods to study whether students believe an increase in the number of graduates would increase or decrease the college wage premium. This has consequences for whether expansions of higher education attendance would level out due to the college wage premium falling, or lead to escalating selection into higher education due to beliefs about being left behind in the matching market. Another promising avenue of research is to further consider how uncertainty about multiple

dimensions of skill could affect the analysis in this model. Furthermore, I have also glossed over the role of the tuition policy in the UK. It may be interesting to analyse with a more general life-cycle model the effect subsidies for higher education, especially for low-earning graduates, may have for incentives to invest in university degrees or to choose less productive courses in university. This is of particular interest to many countries given the high cost of subsidising universities to the Treasury. Finally, my framework leads to the possibility that graduate taxes or subsidies conditional on school grades may be more efficient than flat graduate taxes or subsidies. I hope to eventually extend the welfare analysis to consider a greater policy set than has been considered here.

References

Acemoglu, D. (1999). Changes in unemployment and wage inequality: An alternative theory and some evidence. *American economic review*, *89*(5), 1259–1278.

Adda, J., & Cooper, R. W. (2003). *Dynamic economics: quantitative methods and applications*. MIT press.

Aghion, P., Bergeaud, A., Blundell, R. W., & Griffith, R. (2019). The innovation premium to soft skills in low-skilled occupations. *Available at SSRN 3489777*.

Albrecht, J., & Vroman, S. (2002). A matching model with endogenous skill requirements. *International Economic Review*, *43*(1), 283–305. Publisher: Wiley Online Library.

Altonji, J. (1993). The Demand for and Return to Education When Education Outcomes Are Uncertain. *Journal of Labor Economics*, *11*(1), 48–83.
URL https://EconPapers.repec.org/RePEc:ucp:jlabec:v:11:y:1993:i:1:p:48-83

Arcidiacono, P., Aucejo, E., Maurel, A., & Ransom, T. (2016). College Attrition and the Dynamics of Information Revelation. Tech. rep., National Bureau of Economic Research, Inc. Issue: 22325.
URL https://EconPapers.repec.org/RePEc:nbr:nberwo:22325

Athey, S. (2002). Monotone Comparative Statics under Uncertainty. *The Quarterly Journal of Economics*, *117*(1), 187–223.
URL https://academic.oup.com/qje/article-lookup/doi/10.1162/003355302753399481

Avery, C., Gurantz, O., Hurwitz, M., & Smith, J. (2018). Shifting college majors in response to advanced placement exam scores. *Journal of Human Resources*, *53*(4), 918–956. Publisher: University of Wisconsin Press.

Bagnoli, M., & Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, *26*(2), 445–469.
URL http://link.springer.com/10.1007/s00199-004-0514-4

Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political economy*, *81*(4), 813–846. Publisher: The University of Chicago Press.

Beffy, M., Fougere, D., & Maurel, A. (2012). Choosing the field of study in postsecondary education: Do expected earnings matter? *Review of Economics and Statistics*, *94*(1), 334–347. Publisher: The MIT Press.

Belfield, C., Boneva, T., Rauh, C., & Shaw, J. (2020). What drives enrolment gaps in further education? the role of beliefs in sequential schooling decisions. *Economica*, *87*(346), 490–529. Publisher: Wiley Online Library.

Belfield, C., Britton, J., Buscha, F., Dearden, L., Dickson, M., van der Erve, L., Sibieta, L., Vignoles, A., Walker, I., & Zhu, Y. (2018a). The impact of undergraduate degrees on early-career earnings. Tech. rep., Department for Education.
URL https://www.ifs.org.uk/publications/13731

Belfield, C., Crawford, C., & Sibieta, L. (2018b). Long-run comparisons of spending per pupil across different stages of education. Publisher: Institute for Fiscal Studies.

Blundell, R., Dearden, L., Goodman, A., & Reed, H. (2000). The returns to higher education in Britain: evidence from a British cohort. *The Economic Journal*, *110*(461), F82–F99. Publisher: Oxford University Press Oxford, UK.

Blundell, R., Green, D., & Jin, W. (2018). The UK Education Expansion and Technological Change.

Blundell, R., Green, D. A., & Jin, W. (2022). The UK as a Technological Follower: Higher Education Expansion and the College Wage Premium. *The Review of Economic Studies*, *89*(1), 142–180. Publisher: Oxford University Press.

Bolton, P. (2019). Student loan statistics. Publisher: UK Parliament.

Britton, J., van der Erve, L., & Higgins, T. (2019). Income contingent student loan design: Lessons from around the world. *Economics of Education Review*, *71*, 65–82.
URL https://linkinghub.elsevier.com/retrieve/pii/S0272775717301668

Caplan, B. (2018). The case against education. In *The Case against Education*. Princeton University Press.

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, *36*(S1), S13–S70. Publisher: University of Chicago Press Chicago, IL.

Carneiro, P., Heckman, J. J., & Vytlacil, E. J. (2011). Estimating marginal returns to education. *American Economic Review*, *101*(6), 2754–81.

Carneiro, P., & Lee, S. (2011). Trends in Quality-Adjusted Skill Premia in the United States, 1960—2000. *The American Economic Review*, *101*(6), 2309–2349.
URL www.jstor.org/stable/23045644

Chade, H., Eeckhout, J., & Smith, L. (2017). Sorting through Search and Matching Models in Economics. *Journal of Economic Literature*, *55*(2), 493–544.
URL https://pubs.aeaweb.org/doi/10.1257/jel.20150777

Chade, H., & Lindenlaub, I. (2021). Risky Matching.

Charlot, O., & Decreuse, B. (2005). Self-selection in education with matching frictions. *Labour Economics*, *12*(2), 251–267. Publisher: Elsevier.

Clark, B., Joubert, C., & Maurel, A. (2017). The career prospects of overeducated Americans. *IZA Journal of Labor Economics*, *6*(1), 3.
URL http://izajole.springeropen.com/articles/10.1186/s40172-017-0053-4

Cunha, F., & Heckman, J. (2016). Decomposing Trends in Inequality in Earnings into Forecastable and Uncertain Components. *Journal of Labor Economics*, *34*(S2), S31–S65.
URL https://www.journals.uchicago.edu/doi/10.1086/684121

Cunha, F., Heckman, J., & Navarro, S. (2005). Separating uncertainty from heterogeneity in life cycle earnings. *Oxford Economic Papers*, *57*(2), 191–261.
URL https://EconPapers.repec.org/RePEc:oup:oxecpp:v:57:y:2005:i:2:p:191-261

Delavande, A., & Zafar, B. (2019). University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints. *Journal of Political Economy*, *127*(5), 2343–2393. Publisher: The University of Chicago Press Chicago, IL.

Devereux, P. J., & Fan, W. (2011). Earnings returns to the British education expansion. *Economics of Education Review*, *30*(6), 1153–1166.
URL https://linkinghub.elsevier.com/retrieve/pii/S0272775711000471

Education, D. F., & University Of Essex, I. F. S. (2022). United Kingdom Household Longitudinal Study; UKHLS; NPDUnderstanding Society: Waves 1- , 2008-: Secure AccessUnderstanding Society: Linked Education Administrative Datasets (National Pupil Database), England, 1995-2018: Secure Access. Version Number: 3rd Edition Type: dataset.
URL https://beta.ukdataservice.ac.uk/datacatalogue/doi/?id=7642#3

Ekeland, I., Heckman, J. J., & Nesheim, L. (2004). Identification and estimation of hedonic models. *Journal of political economy*, *112*(S1), S60–S109. Publisher: The University of Chicago Press.

Elias, P., & Purcell, K. (2004). Is Mass Higher Education Working? Evidence from the Labour Market Experiences of Recent Graduates. *National Institute Economic Review*, *190*, 60–74. Publisher: Cambridge University Press.

Freeman, R. (1976). The overeducated american. Publisher: Academic press.

Galichon, A. (2016). *Optimal transport methods in economics*. Princeton University Press.

Gottschalk, P., & Hansen, M. (2003). Is the proportion of college workers in noncollege jobs increasing? *Journal of Labor Economics*, *21*(2), 449–471. Publisher: The University of Chicago Press.

Gourieroux, C., Gourieroux, M., Monfort, A., & Monfort, D. A. (1996). *Simulation-based econometric methods*. Oxford university press.

Green, F., & Zhu, Y. (2010). Overqualification, job dissatisfaction, and increasing dispersion in the returns to graduate education. *Oxford Economic Papers*, *62*(4), 740–763.
URL https://academic.oup.com/oep/article-lookup/doi/10.1093/oep/gpq002

Greene, W. H. (2003). *Econometric analysis*. Pearson Education India.

Guvenen, F., Kuruscu, B., Tanaka, S., & Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, *12*(1), 210–44.

Heckman, J. J., Humphries, J. E., & Veramendi, G. (2018). Returns to Education: The Causal Effects of Education on Earnings, Health, and Smoking. *Journal of Political Economy*, *126*(S1), S197–S246.
URL https://www.journals.uchicago.edu/doi/10.1086/698760

Heckman, J. J., & Sattinger, M. (2015). Introduction to the distribution of earnings and of individual output, by AD Roy. *The Economic Journal*, *125*(583), 378–402. Publisher: Oxford University Press Oxford, UK.

Heckman, J. J., & Vytlacil, E. (2001). Policy-Relevant Treatment Effects. *The American Economic Review*, *91*(2), 107–111. Publisher: American Economic Association.
URL http://www.jstor.org/stable/2677742

Hopkins, E. (2012). Job market signaling of relative position, or Becker married to Spence. *Journal of the European Economic Association*, *10*(2), 290–322. Publisher: Oxford University Press.

Hou, S., & Milsom, L. (2021). The butcher, the brewer, or the baker: The role of occupations in explaining wage inequality. Publisher: University of Oxford.

Jackson, P. (2021). Equilibrium Underemployment.

Jensen, R. (2010). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics*, *125*(2), 515–548. Publisher: MIT Press.

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

Katz, L. F., & Murphy, K. M. (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *The Quarterly Journal of Economics*, *107*(1), 35–78.
URL https://academic.oup.com/qje/article-lookup/doi/10.2307/2118323

Lemieux, T. (2006). Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *American Economic Review*, *96*(3), 461–498.
URL https://pubs.aeaweb.org/doi/10.1257/aer.96.3.461

Lemieux, T. (2014). Occupations, fields of study and returns to education. *Canadian Journal of Economics/Revue canadienne d'économique*, *47*(4), 1047–1077. Publisher: Wiley Online Library.

Leuven, E., & Oosterbeek, H. (2011). Overeducation and Mismatch in the Labor Market. In *Handbook of the Economics of Education*, vol. 4, (pp. 283–326). Elsevier.
URL https://linkinghub.elsevier.com/retrieve/pii/B9780444534446000031

Li, H., & Xia, X. (2022). Grades as Signals of Comparative Advantage: How Letter Grades Affect Major Choices. *SSRN Electronic Journal*.
URL https://www.ssrn.com/abstract=4028958

Lise, J., & Postel-Vinay, F. (2020). Multidimensional Skills, Sorting, and Human Capital Accumulation. *American Economic Review*, *110*(8), 2328–2376.
URL https://pubs.aeaweb.org/doi/10.1257/aer.20162002

Macera, M., & Tsujiyama, H. (2020). Inequality, Productivity Dispersion, and Supply of Skills. Tech. rep., working paper, Goethe University Frankfurt.

Meroni, E. C., & Vera-Toscano, E. (2017). The persistence of overeducation among recent graduates. *Labour Economics*, *48*, 120–143.
URL https://linkinghub.elsevier.com/retrieve/pii/S0927537116303992

O'Leary, N., & Sloane, P. (2016). Too many graduates? An application of the Gottschalk–Hansen model to young British graduates between 2001–2010. *Oxford Economic Papers*, *68*(4), 945–967.
URL https://academic.oup.com/oep/article-lookup/doi/10.1093/oep/gpw027

Patnaik, A., Venator, J., Wiswall, M., & Zafar, B. (2020). The role of heterogeneous risk preferences, discount rates, and earnings expectations in college major choice. *Journal of Econometrics*. Publisher: Elsevier.

Quah, J., & Strulovici, B. (2009). Comparative Statics, Informativeness, and the Interval Dominance Order. *Econometrica*, *77*(6), 1949–1992.
URL http://doi.wiley.com/10.3982/ECTA7583

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, *3*(2), 135–146. Publisher: JSTOR.

Sattinger, M. (1993). Assignment models of the distribution of earnings. *Journal of economic literature*, *31*(2), 831–880. Publisher: JSTOR.

Shephard, A., & Sidibe, M. (2019). Schooling Investment, Mismatch, and Wage Inequality. *SSRN Electronic Journal*.
URL https://www.ssrn.com/abstract=3422854

Shimer, R., & Smith, L. (2000). Assortative matching and search. *Econometrica*, *68*(2), 343–369. Publisher: Wiley Online Library.

Tan, B. J. (2022). The Consequences of Letter Grades on Labor Market Outcomes and Student Behavior. *Journal of Labor Economics*, (p. 719994).
URL https://www.journals.uchicago.edu/doi/10.1086/719994

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.

University Of Essex, I. F. S. (2022). United Kingdom Household Longitudinal Study; UKHLSUnderstanding Society: Waves 1- , 2008-: Secure AccessUnderstanding Society: Waves 1-11, 2009-2020 and Harmonised BHPS: Waves 1-18, 1991-2009: Secure Access. Version Number: 13th Edition Type: dataset.
URL https://beta.ukdataservice.ac.uk/datacatalogue/doi/?id=6676#12

Walker, I., & Zhu, Y. (2008). The College Wage Premium and the Expansion of Higher Education in the UK*. *Scandinavian Journal of Economics*, *110*(4), 695–709.
URL http://doi.wiley.com/10.1111/j.1467-9442.2008.00557.x

Waltmann, B., van der Erve, L., Dearden, L., & Britton, J. (2020). The impact of undergraduate degrees on lifetime earnings. Tech. rep., The IFS.
URL https://www.ifs.org.uk/publications/14729

Wiswall, M., & Zafar, B. (2014). Determinants of College Major Choice: Identification using an Information Experiment. *The Review of Economic Studies*, *82*(2), 791–824.
URL https://doi.org/10.1093/restud/rdu044

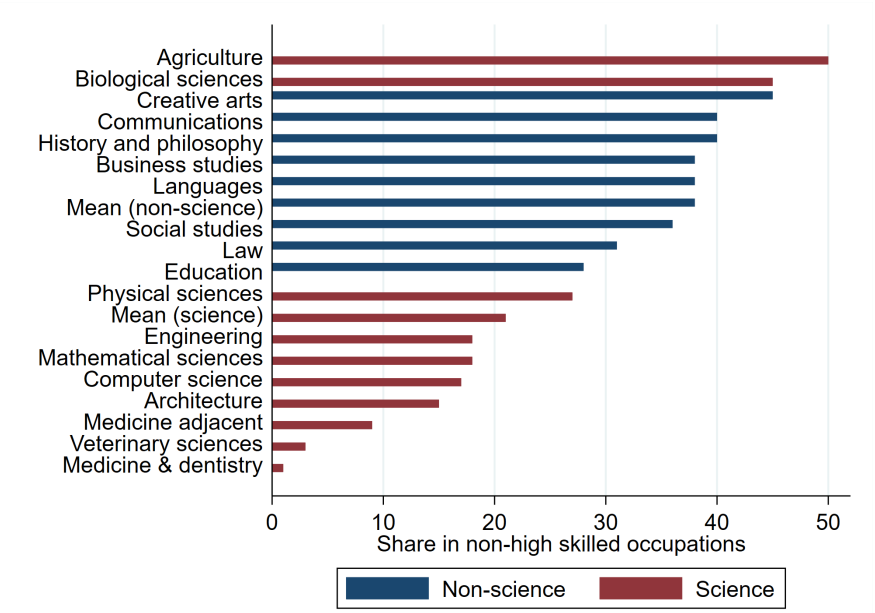
A Additional Stylised Facts

A.1 Employment in low-skilled occupations by subject

Figure 20 plots the share of graduates in non-high-skill occupations by their degree subject. In general, it is true that this rate is higher for workers in non-STEM subjects than in STEM subjects, with biological sciences being an important main exception. However, with the exception of a few vocational subjects like medicine and dentistry, veterinary science and medicine-adjacent subjects like pharmacy, most graduates even in STEM subjects experience at least 10% of workers in non skilled occupations, with the within-science mean being 20%. This suggests that the issues discussed in this paper also apply to STEM fields, even if it is not to a similar extent. Explicitly modelling a

multinomial education choice with choice of subject substantially complicates the matching problem considered and is not feasible with the current data used (the Understanding Society which forms the source for the core of the data work does not have information on university subjects). I hope to incorporate it in future work.

Figure 20: Workers by share in non high-skill occupation by degree subject (2018-19)



Data from HESA outcomes survey for 2018-19 for graduates one-year from graduation. The red bars represent STEM subjects while the blue pairs represent non-STEM subjects. A high-skilled occupation is defined as an occupation which requires at least a degree, according to the ONS occupation classification, and consists of occupations with occupation codes beginning 1 to 3.

B Proofs

B.1 Proof of proposition 1

The optimality condition is as follows:

$$\kappa + \Delta\eta + \beta\Delta w(a) \geq 0$$

First, note that the LHS of the equation is strictly increasing in both $\Delta\eta$ and a. To show that $\Delta w(a)$ is increasing in a, note that:

$$\begin{aligned}\frac{d\Delta w(a)}{da} &= \frac{dw(s(a,1))}{da} - \frac{dw(s(a,0))}{da} \\ &= w'(s(a,1))\frac{\partial s(a,1)}{\partial a} - w'(s(a,0))\frac{\partial s(a,0)}{\partial a}\end{aligned}$$

By assumption, $\frac{\partial s(a,1)}{\partial a} > \frac{\partial s(a,0)}{\partial a}$. Furthermore, we know that the wage function $w(\cdot)$ is convex, so $w'(s(a,1)) > w'(s(a,0))$. Thus, $\frac{d\Delta w(a)}{da}$ has to be greater than 0; the wage return is increasing in

ability. Since the LHS is increasing in a and $\Delta\eta$, it follows that for each value of $\Delta\eta$, there must be a unique value of a which satisfies the optimality condition with equality.

Since the LHS is increasing in a and $\Delta\eta$, it follows that the LHS will only cross the zero threshold once, implying that e^P is increasing in a and $\Delta\eta$.

Let $a^P(\Delta\eta)$ denote the value of a which satisfies the optimality condition with equality, for any given value of $\Delta\eta$. Implicitly differentiate by $\Delta\eta$.

$$\kappa + \Delta\eta + \beta\Delta w(a^P(\Delta\eta)) = 0 \tag{34}$$

$$1 + \beta\frac{d\Delta w(a^P)}{da^P}\frac{da^P}{d\Delta\eta} = 0 \tag{35}$$

$$\frac{da^P(\Delta\eta)}{d\Delta\eta} = -\frac{1}{\beta\frac{d\Delta w(a^P)}{da^P}} \tag{36}$$

The RHS is negative since $\frac{d\Delta w(a^P)}{da^P} > 0$. Thus, we can show that the optimal cut-off $a^P(\Delta\eta)$ is decreasing in $\Delta\eta$.

The solution to each individual's education problem is completely characterised by $a^P(\Delta\eta)$, the cut-off level of ability for each preference state. We know that for any value of a smaller than it, the optimal decision would be to not invest in education, and similarly for any value of a greater than it, the optimal decision would be to invest in education. Also, $\frac{da^P(\Delta\eta)}{d(\Delta\eta)} < 0$, such that when a worker has a greater preference for university, it takes a lower ability threshold to induce them into education. Thus, $e^P(a, \Delta\eta)$ is increasing in a, and $\Delta\eta$.

B.2 Under additive noise, log-concavity of the density function of the error term is sufficient for MLR order

Under an additive structure as specified in equation 1, the joint distribution of θ and a can be derived by convolution as the product of the density function of a and of ε. The conditional densities can then be computed as follows:

$$\begin{aligned}
f_{\Theta,A}(\theta, a) &= f_A(a) \cdot f_\varepsilon(\theta - a) \\
f_\Theta(\theta) &= \int_{-\infty}^{\infty} f_A(a) f_\varepsilon(\theta - a) da \\
f_{A|\Theta}(a|\theta = \theta_1) &= \frac{f_{\Theta,A}(\theta_1, a)}{f_\Theta(\theta_1)} \\
f_{\Theta|A}(\theta|A = a) &= \frac{f_{\Theta,A}(\theta, a)}{f_A(a)} \\
&= f_\varepsilon(s - a)
\end{aligned}$$

Consider the ratio $\frac{f(a|\theta=\theta'')}{f(a|\theta=\theta')}$ where θ'' and θ' are arbitrary values such that $\theta'' > \theta'$. We begin by substituting the expression for the conditional density into the expression.

$$\begin{aligned}\frac{f(a|\theta=\theta'')}{f(a|\theta=\theta')} &= \frac{f_{\theta,a}(\theta'',a)}{f_{\theta,a}(\theta',a)} \cdot \frac{f_\Theta(\theta')}{f_\Theta(\theta'')} \\ &= \frac{f_A(a)f_\varepsilon(\theta''-a)}{f_A(a)f_\varepsilon(\theta'-a)} \cdot \frac{f_\Theta(\theta')}{f_\Theta(\theta'')} \\ &= \frac{f_\varepsilon(\theta''-a)}{f_\varepsilon(\theta'-a)} \cdot \frac{f_\Theta(\theta')}{f_\Theta(\theta'')}\end{aligned}$$

Since $\frac{f_\Theta(\theta')}{f_\Theta(\theta'')}$ is not a function in a, to establish MLR order in θ, we have to show that $\frac{f_\varepsilon(\theta''-a)}{f_\varepsilon(\theta'-a)}$ is increasing in a. Differentiating the expression and signing it, we get the following inequality for $\theta'' > \theta'$:

$$\frac{f'_\varepsilon(\theta'-a)f_\varepsilon(\theta''-a) - f'_\varepsilon(\theta''-a)f_\varepsilon(\theta'-a)}{[f_\varepsilon(\theta'-a)]^2} > 0$$

Since the denominator must be positive, we can examine the following identity:

$$f'_\varepsilon(\theta''-a)f_\varepsilon(\theta'-a) < f'_\varepsilon(\theta'-a)f_\varepsilon(\theta''-a)$$

Given that distribution functions have a positive range, we can rearrange terms as follows:

$$\frac{f'_\varepsilon(\theta'-a)}{f_\varepsilon(\theta'-a)} > \frac{f'_\varepsilon(\theta''-a)}{f_\varepsilon(\theta''-a)}$$

Thus, the condition for the conditional distributions obeying the monotone likelihood ratio order in the signal θ reduces to the following condition on the distribution of the error f_ε.

$$\theta' < \theta'' \rightarrow \frac{f'_\varepsilon(\theta'-a)}{f_\varepsilon(\theta'-a)} > \frac{f'_\varepsilon(\theta''-a)}{f_\varepsilon(\theta''-a)}$$

This condition is equivalent to the condition that $\frac{f'(x)}{f(x)}$ is monotone decreasing in x, and this is sufficient for the distribution f_ε to be log-concave (Bagnoli & Bergstrom (2005)).

The condition for the MLR order also holds when the errors are normally distributed; this can be computed directly as follows.

$$\begin{aligned}\frac{f_\varepsilon(\theta''-a)}{f_\varepsilon(\theta'-a)} &= \frac{\frac{1}{\sigma_\varepsilon\sqrt{2\pi}}\exp(-0.5\cdot(\frac{\theta''-a}{\sigma_\varepsilon})^2)}{\frac{1}{\sigma_\varepsilon\sqrt{2\pi}}\exp(-0.5\cdot(\frac{\theta'-a}{\sigma_\varepsilon})^2)} \\ &= \exp\left(-0.5\cdot\left(\left(\frac{\theta''-a}{\sigma_\varepsilon}\right)^2 - \left(\frac{\theta'-a}{\sigma_\varepsilon}\right)^2\right)\right) \\ \frac{\partial\frac{f_\varepsilon(\theta''-a)}{f_\varepsilon(\theta'-a)}}{\partial a} &= \left(\frac{\theta''-\theta'}{\sigma_\varepsilon^2}\right)\cdot\exp\left(-0.5\cdot\left(\left(\frac{\theta''-a}{\sigma_\varepsilon}\right)^2 - \left(\frac{\theta'-a}{\sigma_\varepsilon}\right)^2\right)\right) > 0\end{aligned}$$

B.3 Proof of proposition 5

Consider the subset of education profiles $\mathcal{Q}$, which correspond to a set of functions Ψ which maps $\mathbb{R}$ to $\mathbb{R}$. We can then rewrite the welfare functional in terms of ψ as follows:

$$W[\psi] = \int_{\Delta\eta}\int_{\varepsilon}\int_{a} \{1(\Delta\eta > \psi(a+\varepsilon))\left[g(s(a,1),\mu(s(a,1))) + \kappa + \Delta\eta\right] + \\ (1 - 1(\Delta\eta > \psi(a+\varepsilon)))\left[g(s(a,0),\mu(s(a,0)))\right]\} dF(a)dF(\varepsilon)dF(\Delta\eta)$$

Note that ψ is a vector space, since it simply maps the real line to the real line. Define a norm for the vector space Ψ as follows: $||\psi|| = \int |\psi(\theta)|d\theta$.

The functional W is differentiable at point $\bar{\psi}$ if there exists a linear operator $L(\psi)$ such that $\forall\epsilon, \exists\delta > 0$ such that $\forall\psi \in \Psi \setminus \{\bar{\psi}\}$ satisfying $||\psi - \bar{\psi}|| < \delta$, $\frac{W[\psi]-W[\bar{\psi}]-L(\psi-\bar{\psi})}{||\psi-\bar{\psi}||} < \epsilon$. Then $L(\psi)$ is the derivative of W at point ψ, denoted $DW[\psi]$.

A necessary condition for the optimality of any ψ as a maximiser of W is that $DW[\psi] = 0$. Writing out the derivative of W with respect to ψ, we get the following expression, noting that we can redefine functionals taking p as an argument to take ψ as an argument by replacing p with $p(\theta, \Delta\eta) = 1[\Delta\eta > \psi(\theta)]$. Noting that $g(s, \mu(s)) = w(s) + \pi(\mu(s))$, we can expand the first term.

$$\begin{aligned} DW[\bar{\psi}] &= \int\int\int \left\{g(s(a,1),\mu(s(a,1),\bar{\psi})) - g(s(a,0),\mu(s(a,0),\bar{\psi})) + \kappa + \Delta\eta\right\} + \\ &\quad \left\{1(\Delta\eta > \bar{\psi}(a+\varepsilon)) \times Dg(s(a,1);\bar{\psi}) + (1 - 1(\Delta\eta > \bar{\psi}(a+\varepsilon))) \times Dg(s(a,0);\bar{\psi})\right\} dF_a dF_\varepsilon dF_{\Delta\eta} \\ &= \int\int\int \left\{w(s(a,1),\bar{\psi})) - w(s(a,0),\bar{\psi})) + \kappa + \Delta\eta\right\} + \left\{\pi(s(a,1),\bar{\psi}) - \pi(s(a,0),\bar{\psi})\right\} + \\ &\quad \left\{1(\Delta\eta > \bar{\psi}(a+\varepsilon)) \times Dg(s(a,1);\bar{\psi}) + (1 - 1(\Delta\eta > \bar{\psi}(a+\varepsilon))) \times Dg(s(a,0);\bar{\psi})\right\} dF_a dF_\varepsilon dF_{\Delta\eta} \end{aligned}$$

Note that at the equilibrium profile $\bar{\psi}$, all workers are expected utility maximising at their education choice. This implies that it is not possible to improve the first term in the derivative by changing the education profile.

$$\int\int\int \left\{w(s(a,1),\bar{\psi})) - w(s(a,0),\bar{\psi})) + \kappa + \Delta\eta\right\} dF_a dF_\varepsilon dF_{\Delta\eta} = 0$$

Thus at the equilibrium education profile, the derivative consists of the remaining two terms:

$$\begin{aligned} DW[\bar{\psi}] &= \int\int\int \left\{\pi(s(a,1),\bar{\psi}) - \pi(s(a,0),\bar{\psi})\right\} + \\ &\quad \left\{1(\Delta\eta > \bar{\psi}(a+\varepsilon)) \times Dg(s(a,1);\bar{\psi}) + (1 - 1(\Delta\eta > \bar{\psi}(a+\varepsilon))) \times Dg(s(a,0);\bar{\psi})\right\} dF_a dF_\varepsilon dF_{\Delta\eta} \end{aligned}$$

In general, the last two terms are not in general non-zero. The surplus profits are positive, and thus for $DW[\bar{\psi}] = 0$, the last term must coincidentally exactly offset the profits term. When this is not the case, the necessary condition for optimality does not hold, implying that the equilibrium education profile is not the optimal education profile from the social planner's point of view.

C Computation and Estimation Details

The estimation procedure requires the simulation of earnings and grades conditional on the distribution of education choices in the economy. To derive the wage function, we require the evaluation of the cumulative density of each value of skill s, which requires numerical integration of a complex object $f_S(s)$, given by equation 13. The probability density function given by equation 13 itself requires integration over the signal noise term ε. I compute the integration in equation 13 using Gauss-Hermite quadrature, since the distribution of ε is assumed to be normal with variance σ_ε^2. I then compute the integration in equation 14 using standard numerical integration.

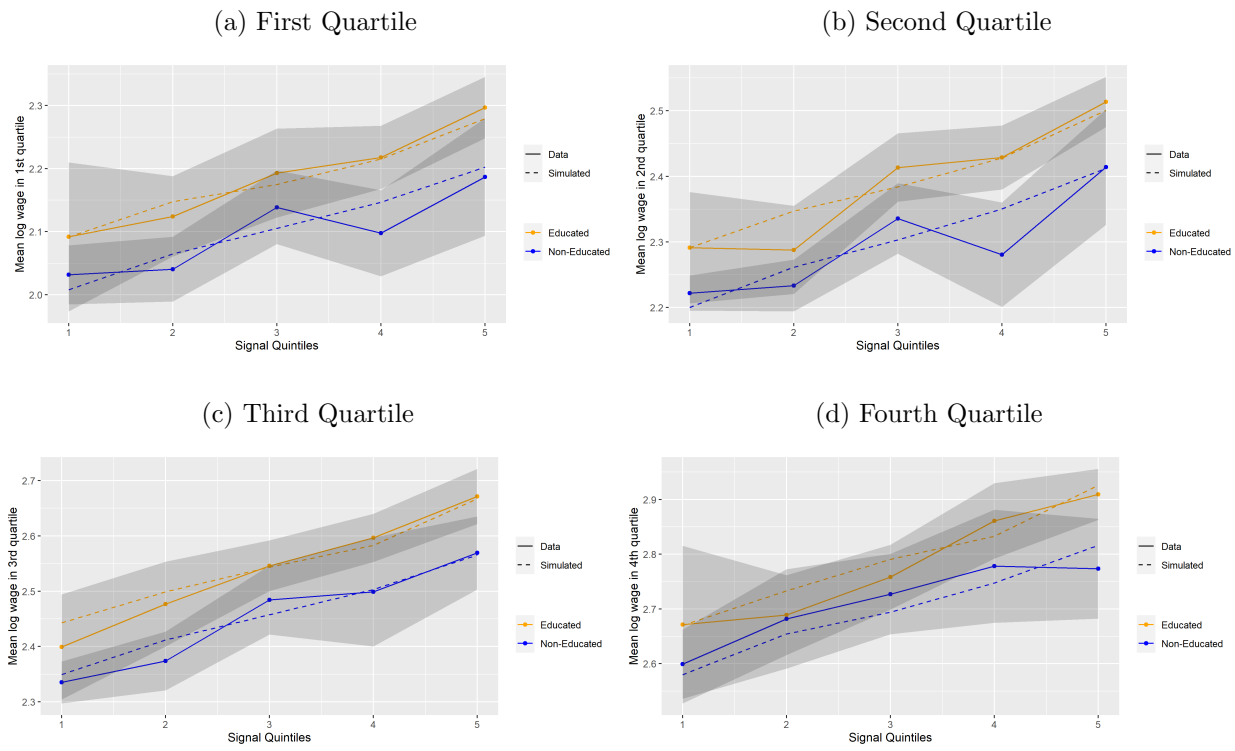
To compute the wage function, I first compute the derivate of the wage function, which has a simple functional form under my parameterisation: $w'(s) = q\gamma_1 s^{\gamma_1 - 1}\mu(s)^{\gamma_2}$. I approximate this function using a monotone cubic spline using Hyman filtering. To derive the wage condition on s, I integrate this spline approximation according to equation 7. Since each set of trial parameter values requires the computation of only a single wage function, I then approximate this wage function using a monotone cubic spline in s to save on computation when simulating the dataset.

Finally, to compute the value function in equation 9, it is necessary to take expectations over the wage function conditional on the signal θ. Under the parametric assumption of normality of both the distribution of a and the distribution of ε, the posterior distribution of $a|\theta$ is also normal. I use Gauss-Hermite quadrature to reduce the computational burden of computing these large expectation terms.

D Additional Targeted Moment Fits

The main figure in this appendix, figure 21 plots the mean log wage within four quartiles of log hourly wages, conditional on degree status and signal quintiles. These quartiles were used as moments in the structural estimation and are used to pin down both the degree of uncertainty about the signal, and the shape of the wage function. The simulated data seems to fit all the moments well, although the some moments are estimated with significant variance.

Figure 21: Simulated versus actual means of log wage within wage quartiles conditional on degree and grade quintile



Notes: This figure plots the mean log-wage within each of four log wage quartiles, conditional on degree status and five signal quintiles.

E Details about implementation of the income-contingent loan policy

In the UK, students typically pay for higher education in the UK with an income-contingent loan, first introduced under the Labour government in 1998. Under the terms of the loan, students borrow the sum required to pay for tuition fees. They then are obliged to make repayments, where the minimum repayment is a certain percentage of their income above some income threshold. They do not have to make any repayments if their income is below the specified threshold. Repayments stop after the initial loan plus interest is fully repaid, or a certain time period has elapsed[50].

The actual loan policy is implemented over a lifetime, which creates scope for dynamic complexities around the optimal repayment of the loan given a certain wage path. This is not captured in my framework as I focus on effectively a single period; in my implementation, I simplify it for a single period as follows. The parameters of the policy are the repayment threshold ι_1, the repayment rate above the threshold ι_2, and the initial loan sum ι_3. Then, the post-policy wage, given

[50]More details can be found on government websites such as https://www.gov.uk/government/statistics/student-loans-in-england-2021-to-2022/income-contingent-student-loan-repayment-plans-interest-rates-and-calculations-england. The UK income contingent loan policy has also been analysed in Britton et al. (2019).

an initial wage of w, is given by the following expression.

$$\text{post policy wage}(w) = \begin{cases} w \text{ if } w < \iota_1 \\ \iota_1 + (1-\iota_2)(w-\iota_1) \text{ if } w \in [\iota_1, \iota_1 + \frac{\iota_3}{\iota_2}] \\ w - \iota_3 \text{ if } w > \iota_1 + \frac{\iota_3}{\iota_2} \end{cases} \tag{37}$$

The expression says that a worker pays nothing under the scheme if their wage is below the income threshold ι_1. Past the threshold, they pay a share ι_2 of their income in excess of the threshold. Once this payment exceeds the initial loan sum ι_3, then the total payment is simply ι_3. I abstract away from the possibility of strategic repayment of the debt to minimise total repayment for high income workers, in this application. I calibrate the parameters as follows. I set the repayment threshold, ι_1, to be 15711 pounds per year, which is the post-tax average of repayment thresholds from 2016-19[51]. I convert this to pounds per hour terms by assuming that workers work for 44 weeks per year and 40 hours per week. The repayment rate is 9% as has been the policy rate for the loan plan taken by workers in my sample. Finally, I set the initial loan sum to be 22,229 pounds, amortised over 30 years at a 2.5% interest rate and converted to pounds per hour terms. The initial sum assumes on average 3.5 years in university, adding together the sum of the mean maintenance and tuition payments for 2006-2009[52].

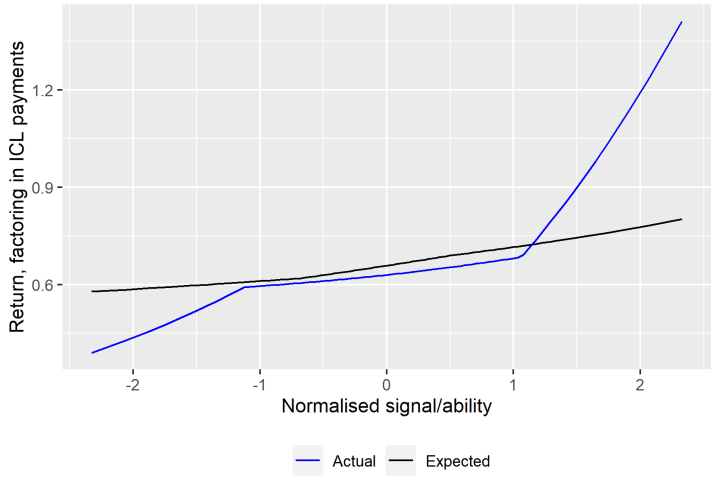
Figure 22 plots the return to higher education after factoring in the payment of the income contingent loan under perfect information conditional on ability (in the blue line), and in expectation conditional on grades (in the black line). The blue line shows that the details of the policy creates two kinks, which create non-linearity. The black line is substantially flatter but also kinked at around -0.8. The increase in the expected return is lower up to that point because the relative upset is depressed by the upside of being higher ability than expected is depressed by the flat portion of the blue line between approximately -1.2 and 1.2. Past that point however, the larger upside of the part of the blue line to the right of approximately 1.2 becomes more relevant, and increases the rate at which the expected return increases.

The policy effectively offers a significant level of insurance to workers, and workers who end up with sufficiently low earnings pay almost nothing for higher education. Importantly, the degree of insurance is greater for workers with low grades than those with high grades, as they are more likely than the latter to have low earnings later in life. Furthermore, workers are not liable to pay more than the original sum that they borrowed. Thus, the amount that a worker would pay stops rising after a certain level of income. In this sense, the greatest relative contribution comes from workers in an intermediate range of income.

[51] See https://www.gov.uk/guidance/previous-annual-repayment-thresholds

[52] The loan sizes are taken from student loan statistics produced for the British parliament by Paul Bolton for the Commons Library. See Bolton (2019). The report was accessed online at https://researchbriefings.files.parliament.uk/documents/SN01079/SN01079.pdf.

Figure 22: Post-income contingent loan return to higher education, conditional on grades



This figure plots the relation between ability and the return to higher education net of payments due because of the ICL policy in the blue line, under the parameters obtained in the estimation procedure (described in section 5). The black line plots the relation between ability and the expected return to higher education, post the implementation of the ICL policy.