

Britton, Jack; Dias, Mónica Costa; Goll, David

Working Paper

Can higher education policy boost intergenerational mobility? Evidence from an empirical matching model

IFS Working Papers, No. 23/06

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Britton, Jack; Dias, Mónica Costa; Goll, David (2023) : Can higher education policy boost intergenerational mobility? Evidence from an empirical matching model, IFS Working Papers, No. 23/06, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2023.0623>

This Version is available at:

<https://hdl.handle.net/10419/284232>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Jack Britton
Monica Costa Dias
David Goll

23/06

Working paper

Can higher education policy boost intergenerational mobility? Evidence from an empirical matching model

Can higher education policy boost intergenerational mobility? Evidence from an empirical matching model*

Jack Britton, *Institute for Fiscal Studies and University of York*,[†]
Monica Costa Dias, *Institute for Fiscal Studies & University of Bristol*,
David Goll, *Institute for Fiscal Studies & University College London*

Abstract

Higher education policy potentially has an important role in improving intergenerational income mobility. We use rich administrative data to estimate a novel two-sided matching model of sorting into field and institution within higher education. We use it to simulate a wide set of policies aimed at boosting mobility. We find that substantially improving population-level mobility statistics with higher education policy is very difficult. However, there is considerable scope to narrow gaps in outcomes between richer and poorer students who leave school with good attainment levels. From a set of policies aimed at influencing student demand, we conclude that by far the most effective policies would target what people study, rather than whether or where they study. However, the biggest improvements in terms of mobility come from an aggressive policy that targets the supply side by enforcing universities to give preferential admission to poorer students who graduate near the top of their secondary school class. This policy would substantially narrow earnings gaps between richer and poorer students with good prior attainment, without significant costs to the taxpayer.

Keywords: Higher education choices; tuition fees

JEL: I23, I24, I26

*The authors are grateful for the Department for Education for data access, financial support, and for comments. The policies considered in this paper were decided by the authors and are not necessarily reflective of current government policy thinking. We also thank the Economic and Social Research Council for funding through grant ES/S010718/1 and the Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (RES-544-28-50001) and NORFACE for funding through grant HuCIAW 462-16-090 for funding. All errors are our own.

[†]Corresponding author: jack.britton@ifs.org.uk. Institute for Fiscal Studies, 7 Ridgmount St., London, UK

1 Introduction

There is a large body of evidence suggesting that whether, where and what people study at university can make a dramatic difference to their subsequent earnings.¹ However, students from poorer backgrounds are less likely to go to university and, if they do go, they are less likely to enrol on high-earning courses. Reforms that address these socio-economic gaps in higher education pathways could therefore have important implications for intergenerational mobility. Indeed, Chetty et al. (2020) draw this conclusion, stating that “changing how students are allocated to colleges could substantially increase intergenerational mobility, even without changing colleges’ educational programs.” This is an appealing argument for policymakers, as it implies well-designed policy tweaks could have large payoffs.

However, it is not clear which policies could realistically achieve this boost in mobility, as in practice, the equilibrium match between students and universities is a complicated function of both the preferences of students for universities and the preferences of universities for students. For example, even if the government forced top universities to make preferential offers to poorer students, eligible students might still not want to accept those offers. Conversely, financial incentives such as targeted grants or fee exemptions might not have much effect if universities are reluctant to admit students with lower prior attainment.

In this paper, we develop and estimate an empirical two-sided matching model of sorting into field and university within higher education in the the United Kingdom (UK). The UK, alongside the United States, is one of the worst performing countries in the OECD in terms for intergenerational mobility.² The model takes into account both the preferences of students for universities and the preferences of universities for students. Through preference restrictions on one side of the market, the use of demand-shifting instruments such as distance and peer choices, and the use of significant reforms that occurred in 2012, we are able identify preference parameters on both sides of the market.

We use the model to evaluate several, hypothetical policies aimed at improving intergener-

¹See, for example Andrews et al. (2017); Anelli (2018); Britton et al. (2021); Chetty et al. (2020); Kirkeboen et al. (2016); Hastings et al. (2018); Dale and Krueger (2002) and Zimmerman (2019).

²See Corak (2013), where intergenerational mobility is measured by the extent to which the earnings of children are predicted by the earnings of their parents .

ational mobility. In general, we find that the scope for higher education policy to substantially affect population-level mobility statistics is limited. However, higher education policy *is* capable of dramatically reducing gaps in outcomes between those from richer and poorer backgrounds, conditional on leaving school with good qualifications. This finding is not inconsistent with Chetty et al. (2020), which also only looks at mobility amongst a specific subset (in their case, those who attend college).

We initially assess a set of policies that target the demand side of higher education. We find that offering free tuition and living cost support for poorer students, increasing living cost support, or conditioning living cost support on attending high status universities are not effective at improving mobility, even amongst those with good prior attainment. However, conditioning such grants on studying in higher-earning subject fields does make a difference. A significant proportion of poorer students shift into these subjects in response to these grants, and that this does improve mobility. This suggests that demand-side policies targeted at boosting mobility should focus more on *what* rather than *whether* or *where* people study.

We then turn to the supply-side, and evaluate a substantial reform to university admission policies. Specifically, the reform forces universities to offer priority admissions to students from poorer backgrounds whose test scores at age 16 are within the top 10% of people from their secondary (high) school in that year. This policy is similar to percent plan policies in Texas and Chile. We find large increases in degree quality for low SES students and substantial improvements in intergenerational mobility amongst those who leave school with good qualifications. This could in theory have significant fiscal costs if the overall match is less efficient - that is, if there is a high degree of complementarity between course quality and students' prior attainment. We allow for these complementarities, but find that, although they exist, the overall long-run fiscal costs of this policy are small. We conclude that policies targeting the supply-side offer the best opportunities for increasing income mobility.

Our model combines a non-transferable utility matching model of the HE market with a life-cycle model of consumption and earnings. In an initial period, potential students either match with a HE course or directly enter the labour market. The HE options available are heteroge-

neous; we incorporate 150 universities and three broad subject fields for students to match with.³ During the period we study, university prices (tuition fees) were fixed externally by the government, and student number controls were tightly regulated and binding. This market structure closely resembles the institutional setting of Agarwal (2015), and our matching model builds on his framework. Importantly, this means we do not need to identify university cost functions or make assumptions about the profit maximisation behaviour of universities.

After completing their degrees, students enter the labour market where they receive stochastic labour earnings that depend on their socio-economic background, school attainment and higher education match. The life-cycle component to the model allows us to incorporate the long-run effects of higher education on earnings. It also allows us to more accurately model the English student loan system, which allows students to pay for their university education through income-contingent deductions from their labour earnings throughout their working life.

The key assumptions of the model are as follows. First, we assume that within field, universities have a common preference ranking of students. This restriction is necessary because allowing for rich preference heterogeneity on both sides of the market creates significant challenges for identification. This assumption seems reasonable within the context of university admissions, where there is a large amount of vertical sorting of students based on test scores. Second, we assume that universities always prefer to fill their places as opposed to leaving them empty. Empirically, we provide evidence that annual enrolments very closely match student number controls, and that there is no correlation between proportion of places filled and course quality. Third, we assume that prospective students base their salary expectations on the salary outcomes of previous cohorts, and that student loans only affect their utility through their impact on their future net income. This is a strong assumption, but one that does not prevent us from being able to replicate the key patterns in the data, most notably the impact of the major reforms to tuition fees that occurred in 2012. Finally, we assume that the equilibrium match is stable. In this context, this assumption requires that any course that a student prefers to the one that they matched with would not be willing to accept them. We argue that the application process in the UK is sufficient

³The subject fields are Law Economics and Management (LEM), Science Technology, Engineering and Mathematics (STEM) and Arts, Humanities and Social Sciences (AHSS). A full list of the individual subjects included in each category are given in Appendix A.

to deliver a stable match, noting the fact that people can apply to multiple places, and can make use of the ‘clearing’ system if they are unhappy with their outcome.

We estimate the model using the Longitudinal Education Outcomes (LEO) dataset, which links together administrative school, university and tax records. The dataset includes everyone born between 1986 and 1996 who attended secondary school and took national tests in England at age 16.⁴ It has detailed information on examinations taken at ages 11, 16 and 18 including specific grades in specific subjects that we use to construct a two-dimensional skill index for each student that we refer to as ‘quantitative’ and ‘communication’ skills. We observe linked tax records for many cohorts of school leavers, including those who did not attend university. These cohorts span the substantial reforms to the higher education system that occurred in 2012, when tuition fees were increased from around £3,000 a year to around £9,000 a year.

The linked school-university records from LEO allow us to characterise the match between students and courses. We also use the school records to generate several instrumental variables that shift student preferences for courses without affecting university preferences. In particular, we can use variation in geographic proximity to different universities and quality of courses available in the local area to shift student choices. We show these instruments to be powerful predictors of both quality of course attended and field of study. These demand-shifters are crucial for separately identifying preferences on both sides of the market. Our panel data on earnings also allows us to account for unobserved permanent heterogeneity in students that can affect both university preferences and students’ future earnings.

We estimate the model using a minimum distance estimator that matches data moments to simulated moments from the model. We estimate using four cohorts who left school between 2006 and 2009. This was a benign period for higher education policy, with no changes to tuition fees or student support. We also simulate the 2010 cohort without formally including them in the estimation in order to assess out-of-sample fit. We are able to replicate sorting patterns extremely well, both in and out-of-sample.

As a further test of the out-of-sample predictive power of our model, we show the model is

⁴These are General Certificate of Secondary Education (GCSE) exams and are taken by almost all students in England.

able to replicate the effects of the 2012 higher education reforms on sorting patterns, even though neither of the cohorts immediately before or after these reforms are included in the estimation data. Although we confirm the result from previous work that these reforms had almost no effect on overall participation (for example, Azmat and Simion, 2020), we present reduced form evidence that the overall zero effect masks large drops in demand from the highest ability students, with their spots being filled in by lower ability peers. Our model successfully replicates these basic patterns, which we believe adds weight to the conclusions from our counterfactual policy experiments.

The model estimates are also revealing about the drivers of higher education choices. The parameter estimates suggest that students, on average, have a strong preference for AHSS subjects over LEM subjects. Students have a distaste for attending institutions far from their home, and this distaste is much greater for low SES students. The university utility parameters suggest that universities care a lot about the match between student skill types and the subject they are studying. Universities strongly prefer quantitative skills over communications skills for STEM subjects and vice versa for AHSS subjects. Finally, the wage parameters in the model suggest important differences in returns to university relative to the estimates obtained from OLS regressions. In particular, the model estimates suggest the returns to university quality are only around half as large as the OLS estimates. A key difference between the two approaches is that the model incorporates unobserved heterogeneity, which boosts earnings by around 4.5% and is also an important driver of selection into higher quality institutions.

The rest of the paper is organised as follows. Section 2 discusses how our paper connects to the literature. Section 3 discusses the LEO data and institutional background before showing some data descriptives. Section 4 outlines our model and Section 5 discusses estimation and identification. Section 6 then shows the fit of the model and estimated impact of the 2012 reforms before Section 7 runs counterfactual policy experiments. Section 8 concludes.

2 Literature

Modelling of the higher education market. Our paper contributes to the growing literature that models higher education choices. Many papers model higher education choices without formally mod-

elling preferences of the supply side (Keane and Wolpin, 1997, 2001; Arcidiacono, 2004; Wiswall and Zafar, 2014; Delavande and Zafar, 2019). Counterexamples are Arcidiacono (2005) and Kapor (2020), who focus on affirmative action policies, and Epple et al. (2006) and Fu (2014), who focus on equilibrium tuition and financial aid policies. These latter papers are all complicated by the fact that in the United States, universities have much more control over their tuition fees, financial aid and student numbers. Because of this, Agarwal (2015), which models the ‘medical match’ between junior doctors (‘residents’) and training hospitals (‘programs’) in the United States, is in fact a much closer institutional setting to ours and is therefore the closest paper to ours methodologically.

We extend Agarwal (2015) in four important ways. First, we increase the dimensionality of both student characteristics and the choices available to students. Students in our model have multidimensional skills and match to a course in a specific subject field at a specific university. This reflects the fact that, in the UK, students make both subject and university choices prior to entry. The fit between student skills and course matters for student preferences, for university preferences, and for subsequent earnings outcomes. Second, we relax the assumption of homogeneous supply-side preferences, allowing courses to rank students differently across different fields. For example, we allow STEM courses to have stronger preferences for quantitative skills than other subjects. Third, we allow for permanent unobserved student heterogeneity that affects both university preferences for students and student earnings. Students can therefore select into courses on characteristics that are unobservable to the researcher, enabling us to correct for potential endogeneity in the earnings equation. Fourth, we incorporate a dynamic lifecycle component to the model, which is important in our context as it facilitates explicit modelling of the English income contingent loans system.

The returns to higher education. There is a large literature investigating the impact of attending different universities on earnings outcomes. Dale and Krueger (2002), Dale and Krueger (2014) and Mountjoy and Hickman (2020) suggest there are weak returns to course quality, but other papers (Black and Smith, 2006; Broecke, 2012; Hastings et al., 2013; Anelli, 2018; Dillon and Smith, 2020) suggest otherwise. Many of these papers depend on strong assumptions of selection on

observable factors. Kirkeboen et al. (2016) provide one method for circumventing these issues - our paper provides an alternative approach that enables us to identify causal estimates of the returns to quality by explicitly modelling selection on unobserved heterogeneity. Our estimates suggest that the returns to course quality are positive and economically significant, but that OLS overstates them by a factor of 1.5 to 2. We also contribute to the literature investigating returns to field of study Kirkeboen et al. (2016); Chevalier (2011); Anelli (2018) and our results are consistent with the idea that on average switching from a low to a high-returning field can matter a lot more for earnings than moving to a higher quality institution within the same field. Finally, we contribute to the small set of papers that have estimated match effects in higher education Dillon and Smith (2020); Mountjoy and Hickman (2020). Consistent with Dillon and Smith (2020), we find evidence of match effects, whereby higher ability students experience greater gains from attending higher quality courses.

The impact of grant aid support and tuition fees on higher education choices. There is a large literature looking at the impact of grants targeted at poorer students to attend university. Denning et al. (2019) investigate the impact of Pell Grants, finding positive effects on graduation and earnings. Marx and Turner (2018) suggest limited effects of Pell Grants on college enrolment or course quality, but Scott-Clayton and Zafar (2019) find evidence that other types of grants can affect these outcomes. Epple et al. (2006) also argue that financial aid could be targeted to improve access for poorer students to top colleges. We find that grants can affect participation of poorer students. This aligns with Dearden et al. (2014) who, also in the UK context, find that maintenance grants have a positive impact on participation among low SES students. We also find that targeted grants can affect field and course quality for low SES students.

There is less evidence on the impact of tuition fees for higher education, although the evidence from Germany and Canada has concluded that fees reduce participation (Hübner, 2012; Neill, 2009). Our estimates align with these findings. We also contribute evidence on the impact of the large tuition fee reforms that occurred in England in 2012. Azmat and Simion (2020) find that there was a negligible impact of these reforms on participation overall, but that there was a narrowing of the gap in participation between young people from richer and poorer backgrounds. They suggest

that the increase in tuition put people off, but that the (relatively modest) increase in maintenance support boosted participation of poorer students, outweighing the negative fee effects for them. Sá (2019) looks at the effect of these reforms on applications, finding large drops for all groups as a result of the reforms. Our paper contributes to understanding of the effects of these reforms by showing that there were responses amongst students with higher prior attainment, and that this can help to explain why socio-economic gaps narrowed (as high prior attainment students are much more likely to be from wealthier backgrounds).

3 Institutional background and data

3.1 Higher education policy in England

Students in England typically enter higher education within one or two years of their final year of secondary school (the school year in which they turn 18). Students study a specific subject (or combination of subjects) at one of around 150 universities in the UK, typically for three or four years. Throughout the period we are studying, the system was almost entirely public and was tightly regulated, with both student number controls and price caps. Student number controls were enforced for each university with large reductions in teaching grants if maximum allotted enrolments were exceeded,⁵ while tuition fee caps were legally imposed.

There have been several changes to the system, which we summarise briefly in Table 1.⁶ Our model is estimated using the 2006-2009 higher education cohorts. During this period, tuition fees were around £3,000 per year and students were able to borrow from the government to pay these fees via an income contingent loan. They were also eligible to borrow additional financial support of around £5,000 per year in ‘maintenance loans’ to help with living costs during study. These loans were also repayable via an income contingent loan.

Many graduates are unlikely to repay the full value of their loan. Students who borrowed over this period are required to repay 9% of their income above a threshold of around £15,000. No

⁵In Appendix F we show that annual enrolments during this period were close to the total student number control and that there is no relationship between university selectivity and how close universities were to their number cap. Exemptions to student number controls were introduced in 2012/13 and 2013/14, but were matched by reductions in student number controls that maintained the same overall tightness of the market. Student number controls were abolished in 2015/16.

⁶Since 2016 there have been further changes, including substantial reforms to loans repayments in 2022. These changes are beyond the scope of what we do here, but are interesting sources of variation for future analysis.

repayments are due if earnings are below £15,000. Interest rates are low and any outstanding debt will be written off after 25 years. The costs of any write-offs fall on the taxpayer.

In 2012, there was a major reform to the system. This reform included a large increase in the tuition fee cap to £9,000 as well as significant changes to the income contingent loan repayment terms. The repayment rate remained at 9% of income, but the repayment threshold increased to £21,000, and interest rates increased substantially.⁷ Finally, the repayment period was extended from 25 to 30 years.⁸ Combined, these changes dramatically increased the costs of higher education for borrowers with the highest lifetime earnings, while they actually reduced costs for the lowest-earning borrowers. We discuss this in more detail in Section 6.3.

Table 1: Approximate student loan rules by university cohort

	1998-2005	2006-2011	2012-2016
Tuition fees	£1,200	£3,000	£9,000
Fees borrowable	No	Yes	Yes
Maintenance support	Yes	Yes	Yes
Repayment rate	9%	9%	9%
Repayment threshold	£15,000	£15,000	£21,000
Interest rate	RPI	RPI	RPI + 3%
Repayment term	25 years	25 years	30 years

Note: The rules given here are simplified for illustrative purposes.

3.2 LEO data

We use the Longitudinal Education Outcomes (LEO) dataset, which is a new administrative dataset that links tax records to university and school records. We observe fully linked records for everyone who went to secondary school in England who was born between September 1985 and August 1996. There are approximately 600,000 people in each school-year cohort. We also have partially linked data (tax-university records, but no school records) going back to people born in the mid-1970s. This is summarised in Table 2.

⁷Specifically, they increased from the minimum of the bank rank rate and RPI to RPI plus up to 3%, depending on income and study status. The RPI is generally considered to overstate true inflation by around 1 percentage point and averaged around 3% in the years following the reform.

⁸There were also some relatively minor changes to student support. Maintenance grants increased by slightly more than the standard incremental increase, while a new ‘National Scholarship Programme’ (NSP) was introduced. However, the NSP was relatively small, and it was often given to students only after they had started at university (Chowdry et al., 2012).

Table 2: LEO data summary

Birth Cohort	Final School Year	School Records	University Records	(Usable) Tax Records	Tax Data Age Range
1975/76	93/94	x	✓	✓*	29 – 40
1976/77	94/95	x	✓	✓*	28 – 39
1977/78	95/96	x	✓	✓*	27 – 38
1978/79	96/97	x	✓	✓*	26 – 37
1979/80	97/98	x	✓	✓*	25 – 36
1980/81	98/99	x	✓	✓*	24 – 35
1981/82	99/00	x	✓	✓*	23 – 34
1982/83	00/01	x	✓	✓*	22 – 33
1983/84	01/02	x	✓	✓*	21 – 32
1984/85	02/03	x	✓	✓*	20 – 31
1985/86	03/04	✓	✓	✓	19 – 30
1986/87	04/05	✓	✓	✓	18 – 29
1987/88	05/06	✓	✓	✓	17 – 28
1988/89	06/07	✓	✓	✓	16 – 27
1989/90	07/08	✓	✓	✓	16 – 26
1990/91	08/09	✓	✓	✓	16 – 25
1991/92	09/10	✓	✓	x	16 – 24
1992/93	10/11	✓	✓	x	16 – 23
1993/94	11/12	✓	✓	x	16 – 22
1994/95	12/13	✓	✓	x	16 – 21
1995/96	13/14	✓	✓	x	16 – 20

Note: * indicates that we only observe this for people who attended university. Final school year is also the year individuals turn 18.

School records are drawn from the National Pupil Database (NPD). They include detailed exam scores from national examinations taken at age 11, 16 and 18. The NPD data also includes secondary school attended (including whether or not it is a private school), sex and parental socioeconomic status (SES), which we discuss in more detail below.

University records are provided by the Higher Education Statistics Agency (HESA). They cover all students attending a UK university and include information on the institution attended and subject studied, which we aggregate up to three broad subject areas.⁹ Throughout this paper, we define a students higher education outcomes (that is whether, where and what they are studying) based on their HESA record two years after finishing school.¹⁰ We only classify individuals as

⁹As introduced above, these are STEM, LEM and AHSS. Where students in the data study subjects that span more than one broad field, we assign them to the field that makes up the larger share of their degree.

¹⁰Dropout rates in the UK are low at only around 10% (see Table 3 from Belfield et al., 2018). We find that OLS estimates of our key earnings equation (see equation 1 below) are not sensitive to different decisions over the treatment of dropouts.

attending university if they are studying full-time for an undergraduate degree at this point. For simplicity, we categorise people into different cohorts based on their final year of school, and we assume that people compete for spots with other students in their cohort.¹¹ For our main estimation sample, approximately 35% of each cohort are enrolled in a full-time undergraduate degree two years after finishing school.

Tax records are provided by Her Majesty's Revenue and Customs (HMRC). We observe annual taxable income from employment, self-employment and from partnerships. The tax data we have includes the entire population between 2005/06 and 2016/17. We observe only annual earnings and the individual identifier generated by the DfE. The NPD and HESA records are hard linked based on individual student identifiers. These datasets are linked to HMRC records based on a fuzzy match from address and surname. Around 95% of the NPD are linked to the tax records. We drop individuals who are not linked.

3.3 Individual skills

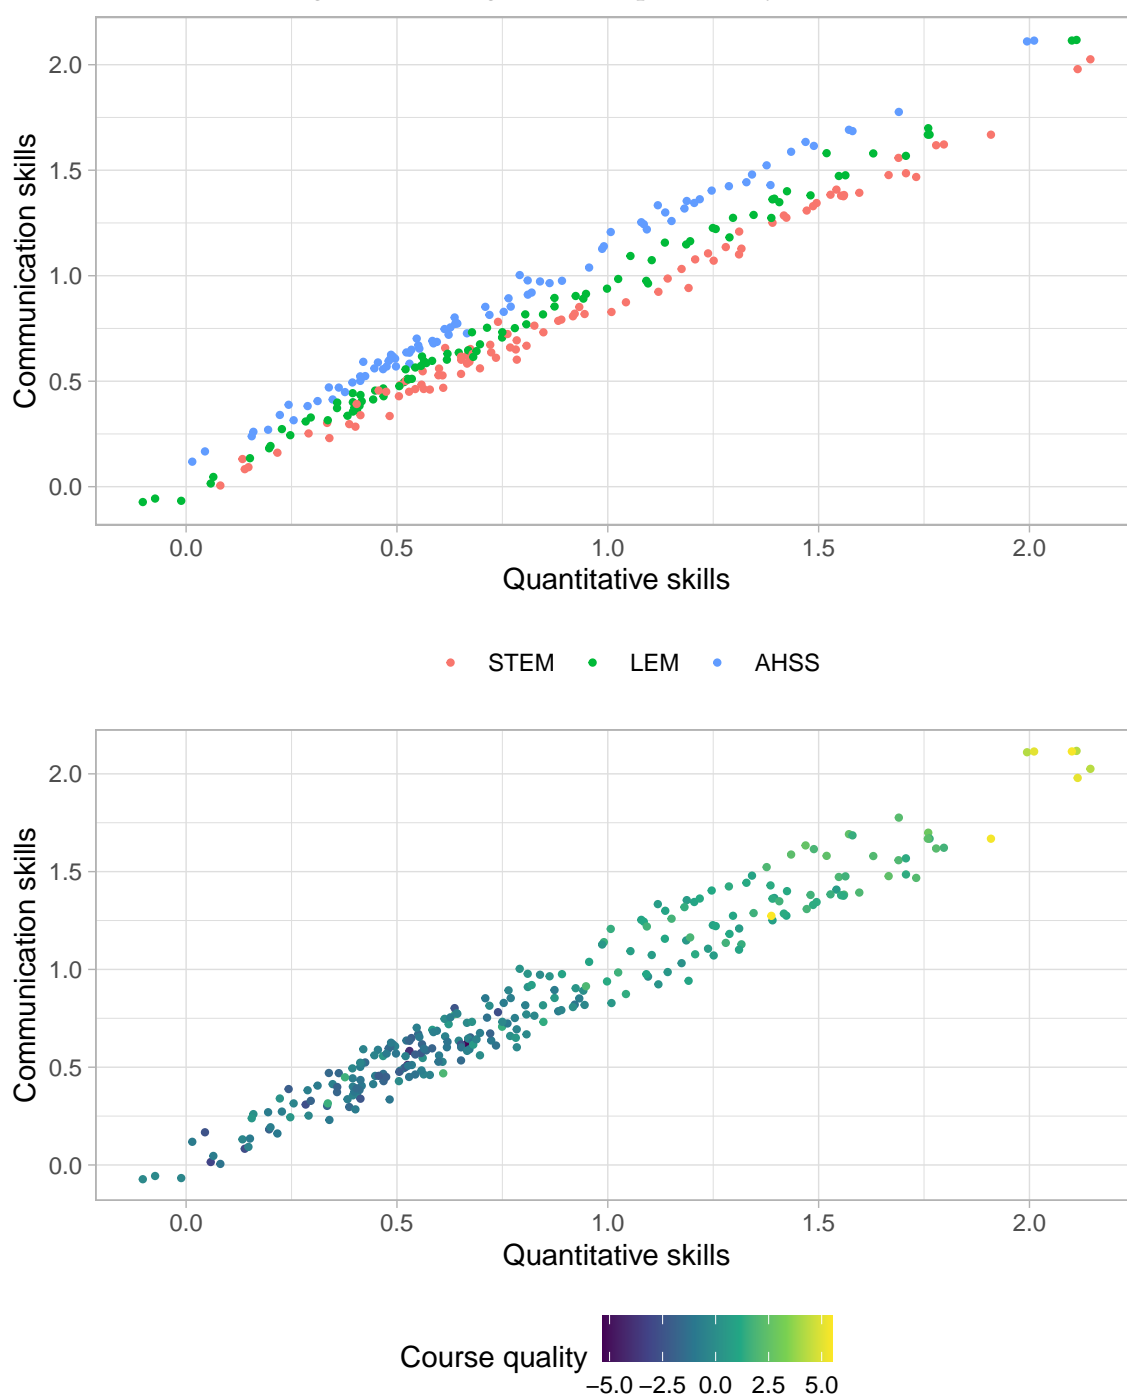
The NPD includes information on exam grades in specific subjects from national examinations taken at ages 16 and 18, as well as numerical scores in national Mathematics, English and Science examinations taken at age 11.

For the purposes of our model, we want to reduce the dimensionality of this prior attainment data while maintaining information on the relative subject strengths of each student. We do this by combining information from multiple exam results to construct two skill measures: one that captures 'quantitative' skills and another that captures 'communication' skills. The first variable puts more weight on test scores in subjects such as mathematics and science, while the second puts more weight on scores in English and humanities subjects. These two variables are highly correlated, but are not perfectly related.¹²

¹¹This allows us to avoid having to model the choice of whether to take up a spot immediately or to defer for a year. In practice people typically apply in the final year of school and universities take decisions on offers without knowing if people are going to defer or attend immediately.

¹²See Appendix B for more detail on the construction of these variables.

Figure 1: Average skill composition by course



Note: Based on 2005/06-2008/09 school leavers (we include more data here because the individual cohorts are noisier).

These skills measures are strongly related to subject choices at university, the quality of university attended, parental SES, and subsequent earnings. The top panel of 1 shows that despite the high correlation between the two measures of skills, relative advantages in one or the other are highly predictive of subject studied at university. Each point in the plot is an individual university-

subject group combination (e.g. STEM at the University of Manchester). The position of the points indicate the average quantitative and communication skills of enrollees in each course. In the top panel, the enrollees in STEM courses have higher average quantitative skills than the enrollees in AHSS courses, but lower average communication skills. LEM courses tend to lie in between.

The bottom panel of the figure shows the same set of points, but now the colour indicates course quality.¹³ This demonstrates the substantial amount of sorting on ability in the UK's higher education system, with the highest quality courses admitting the highest skilled individuals.

3.4 Socio-economic background

The NPD also includes information on parental SES, through the 'income deprivation affecting children index' (IDACI), which we base on where each student lived at age 16. For much of our analysis, we divide people into equally sized 'high' and 'low' SES groups based on this index, with the privately educated included in the high SES group.

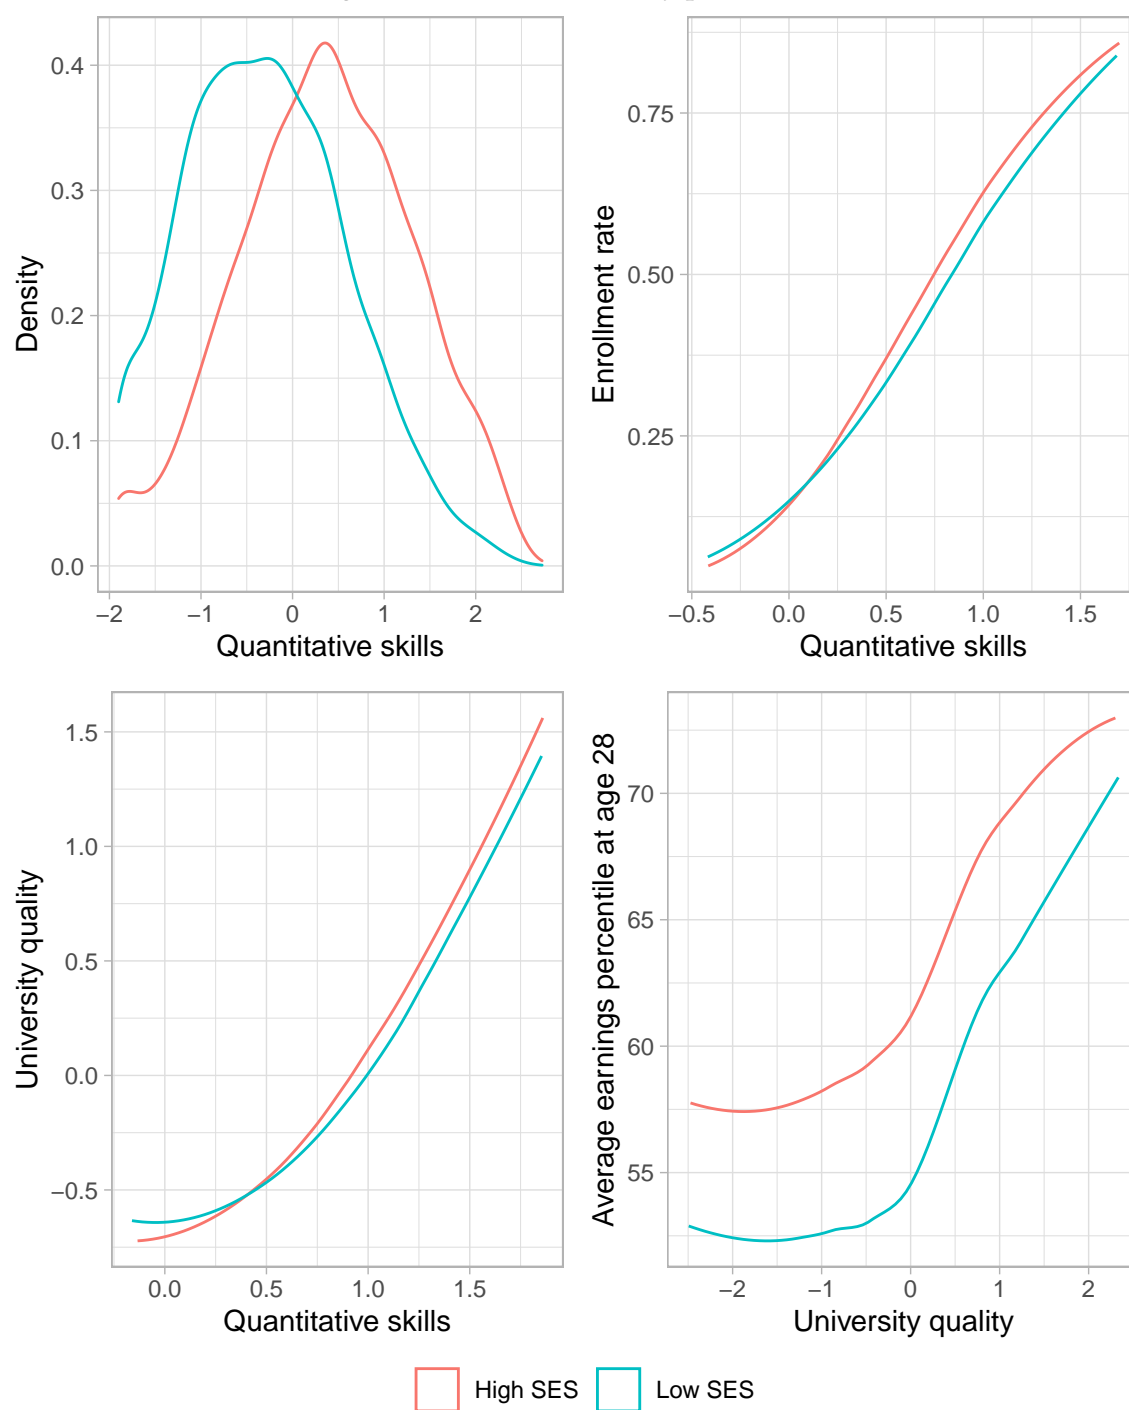
Figure 2 documents the relationship between parental SES, skills, higher education outcomes and earnings. The top-left plot shows that there are large differences in the skill distribution for high and low parental SES students. However, the top-right and bottom-left panels also show that for any given level of skills, low parental SES students are both less likely to attend university and, conditional on attending, are likely to attend a lower quality course. Policies that address these gaps could potentially increase intergenerational income mobility.¹⁴

Finally, the bottom-right plot shows that there is a gap in earnings between the high and low parental SES group, even when conditioning on course quality. This may capture a direct effect of SES on earnings, over and above the impact on higher education outcomes, but may also result from differences in higher education outcomes not captured in course quality. High SES students could, for instance, select into fields with higher returns or attend courses that are a better match for their skills. We explore these alternative explanations further in the next section.

¹³This is constructed using principal components analysis, extracting the first component from five measures of course quality, including two measures of spending, the student-staff ratio, student satisfaction and research quality. See Appendix C for more details.

¹⁴The bottom left plot aligns with the result from Wyness et al. (2021) that poorer students are more likely to under-match at university than their wealthier peers.

Figure 2: Child outcomes by parental SES



Note: Based on 2005/06 school leavers. This shows quantitative skills here, but the pictures are very similar when we use communication skills instead.

3.5 Earnings

Table 3 by shows parameter estimates from the following earnings model, which we estimate using OLS:

$$\ln(y_i) = \alpha_{0F} + \alpha_{1F}X_i + \alpha_{2F}S_i^q + \alpha_{3F}S_i^c + \alpha_{4F}Q_jH_i + \alpha_{5F}Q_jH_iS_i^q + \alpha_{6F}Q_jH_iS_i^c + \epsilon_i \quad (1)$$

y_i is the annual earnings of individual i . X is a vector of background characteristics that includes sex, SES and private school status. S_i^q and S_i^c are quantitative and communication skills, respectively. Q_j is course quality, which has a direct effect and interacts with individual skills. This allows high skilled individuals to have larger returns to attending high quality courses than low skilled individuals, for instance. Q_j is also interacted with H_i , which is a dummy set equal to one if the individual enrolled in higher education, and zero otherwise. All of the parameters are allowed to take different values for individuals who did not attend university ($F = 0$) and, among those who do attend, take different values for each of the three fields that university attendees can study ($F \in \{1, 2, 3\}$). We estimate this equation using the 2005/06 school leaver cohort, with earnings measured in the 2016/17 tax year.

The estimates show that earnings are higher at age 28 for graduates in each of the three subject areas than they are for individuals who do not go to university, although earnings are lower for AHSS courses than for STEM and LEM courses. Female graduates earn around 15% less than male graduates, while female non-graduates earn about 35% less.¹⁵ There is a strong relationship between parental SES and children's earnings at age 28, even conditional on university, sex, and skills. These coefficients suggest that going from the bottom SES decile to the top decile is associated with 10-20% higher earnings, depending on field of study.¹⁶ There is also an earnings premium associated with being privately educated in the arts, while it is not significantly different to zero in STEM and LEM courses. Perhaps surprisingly, it is negative for those who do not attend university.

There are large positive returns to quantitative skills and smaller positive returns to commu-

¹⁵The coefficients in the table are in log points, so the percentage effects referenced in this paragraph are slightly larger than the estimates in the table.

¹⁶The SES coefficients are associated with a one unit increase in the SES variable, but the bulk of the SES distribution is over a narrower range.

nications skills, and strong returns to university quality.¹⁷ There is also some evidence of match effects for quantitative skills, as returns to quality are greater for those with higher quantitative skills. We include these OLS parameter estimates as moments in the estimation of our model.

Table 3: Wage parameters

	No Uni	STEM	LEM	AHSS
Intercept	9.901 (.002)	10.076 (.005)	10.144 (.007)	9.954 (.005)
Female	-.320 (.002)	-.154 (.004)	-.168 (.005)	-.119 (.004)
SES	.213 (.005)	.201 (.014)	.355 (.018)	.229 (.014)
Private	-.065 (.007)	.009 (.006)	.007 (.009)	.041 (.006)
S^m	.157 (.003)	.100 (.007)	.138 (.010)	.105 (.006)
S^c	.058 (.003)	.047 (.006)	.044 (.009)	.024 (.006)
Q		.032 (.002)	.044 (.004)	.029 (.003)
$Q * S^m$.016 (.004)	.040 (.006)	.017 (.004)
$Q * S^c$		-.013 (.004)	-.032 (.006)	-.015 (.004)

Note: Based on the 2005/06 school leavers, $N \approx 450,000$. Outcome variable is log of earnings in the 2016/17 tax year (approximately age 28).

4 Our model of sorting in higher education

We estimate a two-sided non-transferable utility matching model. Each prospective higher education student has a preference ordering over all of the degree courses (field-university combinations). Similarly, the degree courses have preferences over all of the students. Tuition fees are set exogenously and are uniform across programmes, so preferences and capacity constraints jointly determine the equilibrium match.

Students order degree courses based on a direct non-pecuniary value of attending each course as well the discounted present value of utility from consumption over their working life. The non-pecuniary value is intended to capture the direct costs and benefits of attending the course, relative to entering the labour market. These benefits depend on factors such as the match between the students own skills and the content of the course they are choosing, as well as the students' parental SES and how far away the course is from where they live. The present value of consump-

¹⁷Each of these variables - S^q , S^c and Q - have been normalised to have zero mean and a variance of one.

tion is based on the solution of a lifecycle problem in which individuals choose consumption and savings in each period, depending on their income and assets. Income in the life-cycle model depends on their higher education match. Not attending higher education is one of the options that students face, and they order it amongst their degree preferences as if it is a degree course.

On the other side of the market, universities care about their reputations, and therefore the skills of the students they admit. They are also allowed to care about the composition of their intake, and therefore the sex and parental SES of students. We assume that universities strictly prefer filling available spots to leaving them unfilled. In the rest of this section we provide more detail on the model and how assign students to universities.

4.1 Student preferences

Students are characterised by their skills and background characteristics. There are three skills: S^q , S^c and θ . S^q and S^c are the quantitative and communication skill measures introduced in the previous section, while θ is an unobserved skill, known to the agents in the model. θ affects university preferences for students and earnings in later life. This term could, for example, capture non-cognitive skills that are valued by universities, observable through students' applications, and positively impact future earnings. Their background characteristics, X_i , include students' parental SES (SES), an indicator for whether they attended a private secondary schools (*private*) and their sex (g). If student i is matched to course j , the student receives utility:

$$U_{ij} = u(X_i, S_i^q, S_i^c, \text{distance}_{ij}, \text{share}_{ij}, F_j, \eta_{ij}) + EV_{ij} \quad (2)$$

where distance_{ij} is the distance from individual i 's home at age 16 to course j ; share_{ij} is the share of individual i 's peers who chose to study in the same field¹⁸; F_j is field (namely STEM, LEM or AHSS); η_{ij} is a preference shock for course j ;¹⁹ and finally EV_{ij} is the expected value of life-cycle

¹⁸This is intended to capture the idea that seeing peers successfully follow a certain path might make an individual more likely to follow that path, for example due to improved information or confidence Altmejd et al. (2021).

¹⁹We set course preference shocks to be an additive, separable shock for a specific university and a specific field.

consumption for individual i attending course j . The specific functional form for $u(\cdot)$ is given by:

$$u = \beta_{F_j}^u + \beta_{F_j}^X X_i + \beta_{F_j}^q S_i^q + \beta_{F_j}^c S_i^c + \beta^d \text{distance}_{ij} + \beta^{ds} (\text{distance}_{ij} * SES) + \beta^{dp} (\text{distance}_{ij} * \text{private}) + \beta_{F_j}^{sh} \text{share}_{ij} + \eta_{ij} \quad (3)$$

There are 150 universities in the model and three subject areas, resulting in approximately 450 courses²⁰. We allow the parameters of the direct utility function $u(\cdot)$ to vary by field. For example, the relationship between direct utility from attending course j and quantitative skills S^q depends on whether j is a STEM course or an AHSS course.

If student i does not attend university, the student receives utility:

$$U_{i0} = EV_{i0}$$

4.2 Lifecycle model

After entering the labour market, whether directly from school or via higher education, agents earn labour income (y) each period and choose how much to consume (c) and how much to save in safe assets (a). Earnings and employment are determined by a stochastic processes that depends on higher education match, skills and background characteristics. Earnings are subject to deductions for income taxes (I) and student loan repayments (P). The solution to the lifecycle model is governed by the following equations:

$$V_{it}^j = \max_{c_{it}, a_{it+1}} \left[\ln(c_{it}) + \delta EV_{it+1}^j \right] \quad (4)$$

$$\text{st } a_{it+1} = Ra_{it} + d_{it}y_{it} - P(d_{it}y_{it}, l_{it}) - I(d_{it}y_{it}) - c_{it} \quad (5)$$

where δ is the discount factor, R is the interest rate on assets, and d is a dummy for working. The probability of being employed is a exogenous function of age and field of study. If the agent is unemployed, pre-tax earnings are set to zero. We assume individuals do not work if they are at university, that they attend for three years and that they consume all of their maintenance loan while they are at university, and as such, no individual has any assets when they leave education.

²⁰There are slightly fewer than 450 courses because not all universities offer degrees in all three areas

Pre-tax earnings of employed agent i at time t who studied field F at course j are given by:

$$\begin{aligned} \ln(y_{it}^e) = & \alpha_{0F} + \alpha_{1F}X_i + \alpha_{2F}S_i^q + \alpha_{3F}S_i^c + \alpha_{4F}Q_jH_i + \\ & \alpha_{5F}Q_jH_iS_i^q + \alpha_{6F}Q_jH_iS_i^c + \alpha_{7F}\ln(t+1) + \alpha_\theta\theta_i + \epsilon_{Fit} \end{aligned} \quad (6)$$

This specification mirrors equation (1). As above, X includes sex, ethnicity and parental SES; S^m and S^c are quantitative skills and communications skills respectively; Q is course quality; and H is a dummy for higher education attendance. The specification also includes an interaction between quality and skills to allow for match effects in higher education. All of the coefficients vary by $F \in \{\text{No University, STEM, LEM, AHSS}\}$ whether or not people attended higher education, and if they did, by field.

Whereas equation (1) was estimated on one cohort at a specific age, the full model includes earnings growth over the lifecycle. The age-profile of earnings differs across fields and is estimated outside of the model. This model also includes the unobserved skill θ , discussed above. Finally, the error term, ϵ_{Fit} , follows an AR(1) productivity process. Specifically, ϵ_{Fit} evolves according to $\epsilon_{Fit} = \rho_F\epsilon_{Fi,t-1} + \zeta_{Fit}$, where ζ is exogenous iid shock. The AR(1) process is characterised by three parameters: the persistence of productivity shocks (ρ_f), the variance of the iid shock ($\sigma_{\zeta,F}^2$) and the variance of $\epsilon_{i,t}$ at $t = 0$ ($\sigma_{0,F}^2$). All three parameters are allowed to vary by F .

Student loans on graduation are equal to:

$$l_{i0} = \sum_{k=1}^3 (1 + R_l)^k (T + M(\text{SES}))$$

where T is tuition fees and M is borrowing for living costs (“maintenance”) that is dependent on parental SES. Student loans accumulate interest during study at the student loan interest rate R_l . Subsequently, loans evolve according to the following equation:

$$l_{it+1} = R_l l_{it} - P(d_{it}y_{it}, l_{it}) \quad (7)$$

$$l_{it} \geq 0 \quad (8)$$

until any outstanding loan is written off at the end of the repayment period.²¹ The student

²¹To keep the notation tractable, we have omitted the detail that after 2012 student loan interest rates varied depending on whether people were still at university, and on their income once they left university. We do capture this detail

loan repayment function, $P(\cdot)$, which is based on pre-tax income, is given by:

$$P(y_{it}, l_{it}) = \min\{R_l l_{it}, \max\{0, (y_{it} - \phi)\tau\}\}$$

where ϕ is the loan repayment threshold and τ is the repayment rate. Income tax, $I(\cdot)$, is a function of individual gross earnings. Agents are provided with a minimum income floor. Earnings above the income floor are taxed, with step-wise increases in the marginal rate of income tax at specific earnings thresholds. The marginal tax rates, earnings thresholds and income floor are all based on 2019 tax and benefit policy.

4.3 Course preferences

Courses are characterised by their field, $F_j \in \{\text{STEM}, \text{LEM}, \text{AHSS}\}$, and their quality, Q . If student i is matched to course j , the course receives utility:

$$W_{ji} = \gamma_{F_j}^q S_i^q + \gamma_{F_j}^c S_i^c + \gamma^X X_i + \gamma^\theta \theta_i + \eta_{F_j, i} \quad (9)$$

Courses are assumed to observe θ , which is unobservable to the researcher. The introduction of θ therefore allows the model to capture correlations between earnings and sorting outcomes. Importantly, course preferences are allowed to vary by field. This is an extension of Agarwal (2015), which assumes common supply side preferences. Allowing universities to have field-dependent preferences over students is logical as students with different skill sets will match more appropriately to different fields. However, course preferences only vary across fields based on observables; the shock, $\eta_{F_j, i}$, is common to all courses.

Finally, as mentioned above, we assume that courses always strictly prefer filling places to leaving them unfilled. This is consistent with a model of higher education with high fixed costs but low marginal costs of admitting additional students.

4.4 Solving the matching problem

To describe how we solve the model, we need to introduce some notation. $\mu : N \rightarrow J$ is the match function, which assigns students to courses. The inverse of the match function, $\mu^{-1}(j)$, is the set

in our simulation of the 2012 reforms, however.

of students matched to course j . From above, W_{ij} is the value of student i to course j . Let the minimum value of the student assigned to course j , $\min_{i' \in \mu^{-1}(j)} \{W_{i'j}\}$, be written as \bar{W}_j^k , where k is the algorithm iteration. The algorithm we use to solve the model is:

1. Set $\bar{W}_j^0 = -\infty$ for all j
2. Let all students select the course that they prefer given constraint $W_i > \bar{W}_j^k$
3. Count the number of students selecting each course, tot_j
4. For all courses with $tot_j > p_j$, where p_j is the capacity constraint:
 - Sort students who selected oversubscribed course j by W_i
 - Set \bar{W}_j^{k+1} to the W_i of the p_j th student
5. Repeat steps 2 to 4 until no universities have $tot_j > p_j$

This is a modified Gale-Shapley algorithm. In our setting, the resulting match is the unique stable equilibrium. Here, stability is equivalent to:

$$U_{i,j} > U_{i,\mu(i)} \implies \min_{i' \in \mu^{-1}(j)} \{W_{i'}\} > W_i$$

Stability guarantees that if student i prefers course j to their own course, then course j is at capacity and there is no student on course j that course j ranks lower than student i . We assume that the equilibrium match in the data is also stable.

In practice, students only initially apply to five courses, and there is likely to be some strategic application behaviour that violates stability. It is possible that some students failed to apply for a course that would have been willing to accept them and that they preferred to their final match. However, most courses publicise their approximate requirements for accepting applications in advance, giving students a good indication of their chances of receiving an offer. Five applications is also sufficient for students to include an insurance choice and still apply to four additional places that they would want to attend. Furthermore, if students receive no offers from the five courses they apply to, or have a place that they are unhappy with, there is a ‘clearing’ system which matches them to universities with unfilled spots just before the start of the academic year.

They can also reapply the following academic year. As a result, we think that assuming stability is a reasonable approximation in this market as violations are likely to be small.

5 Estimation

5.1 Minimum distance procedure

We use a simulated minimum distance estimator, solving for $\hat{\Theta}$ such that:

$$\hat{\Theta} = \min_{\Theta} (M - M(\Theta))' \mathbf{W} (M - M(\Theta))$$

Here M is a vector of data moments and $M(\Theta)$ is the vector of simulated moments from the model. \mathbf{W} is a $k \times k$ weight matrix, where k is the number of moments. We follow several papers (e.g., Blundell et al., 2016) and use the diagonal weight matrix where each of the elements is equal to the inverse of the variance of each moment in the data (approximated using a bootstrap). We then follow French and Jones (2004) in our computation of asymptotic standard errors.

Each higher education market consists of a single cohort of school leavers. For estimation, we use the four cohorts of school leavers from 2006 to 2009. As described above, this was a relatively benign period for English higher education policy. We estimate the life-cycle earnings parameters outside the model, including parameters defining life-cycle earnings growth (α_{7F} in equation 10), the probability of employment, and the parameters defining the persistent earnings shock. We set the discount rate at 5% per year.

5.2 Identification

A key insight from Agarwal (2015) and Diamond and Agarwal (2017) is that, in a many-to-one matching market, data on observed outcomes is sufficient to identify many parameters of interest. Application data is not strictly necessary. Nevertheless, we face several challenges with the identification of our model. First, we need to separate the preferences of the students from those of the programmes. Second, we need to identify returns to programmes in the labour market.²²

The subsections below describe the data features that assist in identification.

²²One point of divergence from Agarwal (2015) is the incorporation of lifetime earnings. We include several earnings moments, including the parameters from the OLS earnings regression in equation 1.

5.2.1 Excluded variables from earnings and university preferences

We include three types of instrumental variables (which we denote Z). Each of these are intended to shift student demand for courses, without directly affecting earnings or the preferences of courses for students (these are the exclusion restrictions). All instruments are defined at the student level. First, for each field, we include the average quality (Q_j) of courses in the relevant field located within 40km of the student's secondary school. Second, we include the distance from the student's secondary school to the nearest Russell Group university, nearest pre-1992 university nearest and nearest post-1992 university. These university groups are commonly used in the literature as broad proxies for university quality. Third, and again by field, we include the share of a students' peers studied in that field, conditional on attending university.²³

Table 4: Instrument power

	Study STEM	Study LEM	Study AHSS	Uni. Quality
<i>Average quality of local unis</i>				
STEM	-0.004* (0.001)	-0.006* (0.001)	-0.008* (0.001)	-0.315* (0.01)
LEM	-0.002* (0.001)	0.003* (0.001)	0.001 (0.001)	0.197* (0.011)
AHSS	0.007* (0.001)	0.005* (0.001)	0.008* (0.001)	0.214* (0.011)
<i>Distance to nearest</i>				
Russell Group Uni	-0.003* (0.000)	-0.001* (0.000)	-0.002* (0.000)	0.002 (0.002)
Pre-1992 Uni	-0.002* (0.000)	-0.003* (0.000)	-0.002* (0.000)	-0.025* (0.002)
Post-1992 Uni	-0.004* (0.000)	-0.003* (0.000)	-0.004* (0.000)	-0.012* (0.002)
<i>Share of peers choosing</i>				
STEM	0.005* (0.000)	-0.003* (0.000)	-0.004* (0.000)	-0.063* (0.003)
LEM	0.012* (0.000)	0.019* (0.000)	0.005* (0.000)	0.165* (0.003)
AHSS	-0.182* (0.000)	-0.147* (0.000)	0.11* (0.000)	0.518* (0.002)
Controls	✓	✓	✓	✓
F stat	689.7	1215.1	269.2	1793.6
N	2,268,051	2,268,051	2,268,051	660,986

Note: * indicates $p < 0.05$. The first stages here are estimated on four cohorts of data (the 2005/06-2008/09 school leavers). Controls include SES, a private indicator, sex and our two skill measures.

In the model, we utilise these instruments by including a direct measure of distance to the

²³We measure individual i 's peers as students in schools that are local to individual i (within 5km), excluding their own school. We exclude students from individual i 's own school due to concerns that the exclusion restrictions would be violated.

relevant course and the share of peers selecting the same field directly in the individual's utility function, as outlined in equation 2. We also include coefficient estimates from regressions of these instruments on sorting outcomes as moments in estimation.

Table 4 shows that these instruments shift educational choices. For the first three columns, the outcome variable is an indicator for studying in a particular field (including those who do not go to university in the regression), while for the final column the outcome is university quality (conditional on attending). In almost all cases the parameters are statistically significant at the 5% level, and the corresponding F statistics are large.

5.2.2 Sorting patterns

Sorting patterns reveal whether two different programmes (or two different students) are equally desirable (Diamond and Agarwal, 2017). If two different student types are equally desirable, then the distribution of programmes they match to should be similar. These patterns help us identify how programmes value student characteristics, and how students value programme characteristics. We therefore include moments on the covariance between individual level characteristics of enrollees and the characteristics of the matched course.

5.2.3 Many to one matching

Agarwal (2015) argue that the comparison of within-course and between-course variation can be informative about the value courses place on particular student characteristics. Comparisons of this type are only possible within many-to-one matching markets. For example, if courses highly value students' quantitative skills then variation in quantitative skill within courses should be small compared to variation between courses. There will be positive assortative matching between quantitative skills and course quality, and as a result little overlap in the skill distribution of students at high quality and low quality courses. By comparison, if universities do not value a characteristic, then the distribution of that characteristic at each course should be similar. In this case, between-course variation of the characteristic will be much smaller than the within-course variation. We therefore include moments describing within and between variation in student characteristics by field of study and the correlations between one's characteristics and those of

their peers on their course.

6 Parameter estimates and model fit

6.1 Parameter estimates

Tables 5, 6 and 7 report the model parameter estimates for student utility, earnings and university utility respectively. Table 5 shows the parameters relating to students' direct utility from attending university, which is measured relative to the outside option of entering the labour market. The intercept terms show that AHSS courses give people the highest direct utility while at university, while LEM courses give people the lowest direct utility. Women get less utility from all three subject areas than men do, and they particularly dislike LEM courses.

	STEM	LEM	AHSS		
Intercept	.023 (.002)	-.787 (.003)	1.827 (.002)	Distance	-1.265 (.001)
Female	-.050 (.002)	-.798 (.003)	-.035 (.002)	(Distance x SES)	.743 (.007)
SES	.796 (.002)	.431 (.003)	.783 (.013)	(Distance x Private)	.273 (.001)
Private	.378 (.011)	.360 (.004)	.371 (.003)	Field shock	1.007 (.001)
S^m	.010 (.002)	-.013 (.001)	-.007 (.002)	Uni shock	.569 (.001)
S^c	.011 (.002)	-.002 (.001)	-.027 (.002)		
Share	.099 (.001)	.214 (.001)	.150 (.001)		

Note: Parameters in the right hand panel are not allowed to vary by field. Standard errors, shown in the parentheses, are constructed following French and Jones (2004).

Direct utility increases strongly with parental SES for all subjects (particularly STEM and AHSS), and there is a further jump in utility for the privately educated for all fields. Students get more utility from studying in a given subject field if more of their school peers studied in that field. Finally, students have a strong disutility from attending a university that is further away, all else being equal, but the negative effect of distance is much larger for students with low parental SES. To contextualise these numbers, the parameter estimates imply that a student from the poorest SES decile would have to be paid around £2,000 a year in additional maintenance grants in order to fully compensate for travelling an extra 100km from home to attend university.

The equivalent figure for highest-decile SES students is just over £1,000, while it is only around £500 for a privately educated student.

Table 6 shows estimated earnings parameters from the model, comparing it directly to equivalent estimates from the OLS regression model presented earlier. The model estimates are generally quite similar to OLS but with some important differences. The model suggests that being from a higher SES background boosts your earnings if you do not attend university by about 50% more than the OLS estimate suggests. The model also finds that for people who do not attend university there is a small premium for the privately educated, which is the opposite sign to the OLS estimates (and is perhaps more in line with what we would expect to see here). The model finds no premium for the privately educated amongst those who do go to university, although the SES premium remains very strong.

Table 6: Earnings regression parameters


	No Uni		STEM		LEM		AHSS	
	Model	OLS	Model	OLS	Model	OLS	Model	OLS
Intercept	9.897 (.001)	9.901 (.002)	10.081 (.000)	10.076 (.005)	10.109 (.000)	10.144 (.007)	9.989 (.000)	9.954 (.005)
Female	-.309 (.001)	-.320 (.002)	-.149 (.000)	-.154 (.004)	-.186 (.000)	-.168 (.005)	-.132 (.000)	-.119 (.004)
SES	.324 (.001)	.213 (.005)	.199 (.001)	.201 (.014)	.209 (.001)	.355 (.018)	.232 (.001)	.229 (.014)
Private	.051 (.001)	-.065 (.007)	-.026 (.001)	.009 (.006)	-.005 (.001)	.007 (.009)	.000 (.001)	.041 (.006)
S^m	.104 (.001)	.157 (.003)	.085 (.001)	.100 (.007)	.125 (.000)	.138 (.010)	.109 (.001)	.105 (.006)
S^c	.054 (.001)	.058 (.003)	.045 (.001)	.047 (.006)	.031 (.000)	.044 (.009)	.008 (.000)	.024 (.006)
Q			.023 (.000)	.032 (.002)	.026 (.000)	.044 (.004)	.018 (.000)	.029 (.003)
$Q * S^m$.030 (.000)	.016 (.004)	.008 (.000)	.040 (.006)	.021 (.000)	.017 (.004)
$Q * S^c$			-.002 (.000)	-.013 (.004)	.008 (.000)	-.032 (.006)	.009 (.000)	-.015 (.004)
θ^*	.044 (.001)		.044 (.001)		.044 (.001)		.044 (.001)	

Note: α_θ is fixed to be constant across all fields. As θ is unobserved, this is omitted from the OLS regressions. The OLS regressions are as in Table 3. Standard errors, shown in the parentheses, are constructed following French and Jones (2004) for the model estimates.

The other important way the model estimates differ from the OLS estimates is through the estimated returns to quality. In general, OLS tends to overstate the returns to quality for average skilled students, while understating the returns to quality for high skilled students (at least for

STEM and AHSS). This can be seen from the fact that the estimates for the returns to quality are generally higher for OLS than for model (0.032 vs 0.023 for STEM; 0.044 vs 0.026 for LEM; 0.028 vs 0.018 for AHSS), while the interaction terms between quality and skills are generally larger in the model than for OLS.

Figure 3: Returns to quality by subject and skill level



Note: Based on parameter estimates from Table 6. Returns to quality are relative to attending a median-quality course. In the left hand panel we set quantitative and communications skill levels to zero, while in the right hand panel we set them both equal to one.

We present this more visually in Figure 3, which shows the returns to quality relative to the median for each subject area. For example, the OLS parameters imply that going from a median quality to a 99th percentile quality course in STEM for an average-skilled individual would increase earnings by around 18%, while the model estimates suggest the increase is closer to 13%. Meanwhile, the OLS estimates suggest that the equivalent quality move for a high-skilled individual would increase earnings by just over 20%, while the model suggests it is closer to 32%.

The model therefore implies stronger match effects, in that higher skilled individuals experience higher returns to quality. We also show in Appendix Figure A5 that match effects are particularly strong for quantitative skills in STEM, although there are strong returns to quality for individuals with high quantitative skills in all three subject areas. Conversely, there is no complementarity between communication skills and university quality in STEM, while there is in the other two subject areas.

These differences are likely related to the inclusion of the unobserved skill, θ , in the structural model. θ has a large effect on wages, with a one standard deviation increase in θ boosting annual earnings by around 4.5%. Selection into higher quality courses due to high unobserved θ will generate upward bias in the OLS estimates for the return to quality. Another interesting result from including the unobserved heterogeneity terms is that it flips the sign on the privately educated dummy for those who do not attend university. This means that while OLS suggests that being privately educated is negatively related to earnings amongst non-graduates, our model suggests that this result is driven by unobserved selection.

Table 7: University utility parameters


			STEM	LEM	AHSS
Female	-.062 (.001)	S^m	1.571 (.001)	.613 (.001)	.108 (.001)
SES	-.020 (.001)	S^c	.052 (.001)	.823 (.001)	1.435 (.001)
Private	.295 (.001)				
θ	.051 (.000)				

Note: Parameters in the left hand panel are not allowed to vary by field. Standard errors, shown in the parentheses, are constructed following French and Jones (2004).

Finally, Table 7 provides the university utility parameter estimates for the model. The estimates suggest that STEM courses value quantitative skills much more highly than communication skills, AHSS courses value communication skills more highly, while LEM courses value the two skill sets similarly. We show this visually in Figure 4. The top left hand panel shows the increase in standard deviations (SD's) of one skill required to compensate for the loss of one standard deviation in the other skill, all else equal. For STEM only a tiny increase in quantitative skills (0.03 SD's) is required to compensate for the loss of one standard deviation in communications skills. For AHSS, an

increase of just 0.07SD's in communications skills would fully compensate for the loss of one SD in quantitative skills. Meanwhile, LEM courses show a weak preference for communications skills over quantitative skills, as a 0.75 SD's increase in quantitative skills would fully compensate for the loss of 1 SD in communications skills.

Figure 4: Increase in skills (in SD's) required to compensate for the loss of...



Note: estimates based on the parameter estimates from Table 7.

The parameters also suggest that that universities are almost completely indifferent about an applicant's sex (just 0.04 SD's in quantitative skills are required to compensate for a student being female rather than male for STEM courses) or their SES if they state educated (the skill compensation required to compensate for a low versus high SES applicant is close to zero). We find these results reassuring, as universities are not allowed to select on sex, while SES is quite hard for them to observe, especially during this period.²⁴ Meanwhile, the model estimates do suggest that uni-

²⁴We are studying university entry in the mid to late 2000s, which is prior to introduction of Fair Access Agreements and Contextualised Admissions which took off from around 2010.

versities rank privately educated students more highly in their selection processes, all else being equal, although these preferences are still quite weak relative to the preferences for skills: in STEM, 0.18 SD increase in quantitative skills would fully compensate for the loss of a privately educated student. This could be reflecting a genuine preference for the privately educated (for example, because privately educated students might be easier to teach or be more likely to do well in the labour market), or it could be picking up another factor that universities care about that we are not explicitly modelling: for example, the privately educated might write much better personal statements, or might get better references letters from their teachers, both of which are included in their university applications. Finally, we note that universities also place a positive weight on unobserved skills that is quite small relative to the weights placed on observed skills.

6.2 Model fit

The model is able to replicate many of the important patterns in the data. As examples of this, Figure 5 shows university enrolment rates by SES, and Figure 6 shows earnings outcomes by SES. The model replicates enrolment patterns and the relationship between parental SES and child's income very well. It is also able to replicate the sorting patterns by subject and by university quality that we showed in Figure 1 extremely well (see Appendix Figures A3 and A4).

Figure 7 then explores how our model performs both in and out of sample. The points in the figure show the average quantitative skills by subject and university quality decile by year. We include four separate markets in estimation (namely, 2005/06-2008/09 school leavers) and show each of these four markets with circles in the plot (there are therefore $3 \times 4 \times 10$ circles). We additionally show equivalent estimates for 2009/10 school leavers, which is a cohort that was not included in estimation, which are given by the triangles (there are therefore $3 \times 1 \times 10$ triangles).

Figure 5: Relationship between parental SES and child's university enrolment






Figure 6: Relationship between parental SES and child's earnings



The y axis is the model and the x axis is the data, meaning deviations from the identity line reflect differences between the model and the data. Overall we observe a very strong correlation between the model and the data.²⁵ Furthermore, the out-of sample fit is certainly no worse than the in-sample fit, which means that our model is able to predict the match for markets other than those it is estimated on.



Note: Circles show in sample points (2005/06-2008/09 school leavers), triangles show out of sample points (2009/10 school leavers).

²⁵This plot closely resembles a plot of model fit produced in Agarwal (2015). This plot and his are visually very similar in terms of fit. We get a near-identical plot when we consider communications skills rather than quantitative skills.

6.3 Validation of the model through the 2012 reforms

As a further exploration of the ability of our model to simulate sorting patterns in the market, we test how well the model is able to replicate the effects of the large reforms to student loans that occurred in 2012. This is another out-of-sample test, as we estimate the model using earlier cohorts and then simulate the impact of the reform without explicitly including this in estimation. We then compare the model predicted impact to what we observe in the data.


As mentioned in Section 3, the headline component of the 2012 reforms was the large increase in tuition fees from £3,000 to £9,000 that applied to all English students. However, the contemporaneous changes to the income contingent loan repayment terms - and in particular, the raising of the loan repayment threshold - meant that the impact of higher tuition fees on the returns to higher education varied substantially depending on expected earnings. In fact, as shown by Belfield et al. (2017), the 2012 reforms had almost no effect on the total cost of higher education for lowest earning graduates but substantial effects on higher earning graduates. The reason is that lower earners did not pay back the full balance of their loan in the pre-2012 system, so increasing the principal of their loan did not increase their repayments. Instead, they slightly benefited from an increase in the repayment thresholds. High earning graduates, on the other hand, faced large increases in total repayments, as they repay most or all of their loans under both regimes, and therefore the increase in debt on graduation combined with the high interest rates resulted in large increases in repayments.²⁶

Drawing on the idea that those with the highest earnings potential were most affected by the reform, we investigate how the impact of the tuition fee increase varied by skills, which are highly predictive of later-life earnings. The blue dots in Figure 8 show our estimates of the change in participation rate through the reforms by quintile of quantitative and communication skills. For both skill types, they show drops in participation (relative to expectation) amongst higher skilled individuals of around two percentage points, with *increases* in participation further down the ability distribution.

One potential explanation for this pattern is that the government-imposed student numbers caps, coupled with tightly regulated prices, resulted in excess demand for higher education prior

²⁶Figure A6 in Appendix H shows estimated lifetime repayments pre and post reform.

Figure 8: 2012 reform effects by ability, model vs data




Note: The data shows the raw changes in participation between 2010 and 2012, adjusting for overall growth in participation rates.

to 2012. The increase in tuition fees decreased the returns to higher education for relatively high ability students and discouraged some from enrolling. Their vacated spots were filled by students further down the ability distribution. The biggest participation increases occurred among students in the fourth and fifth attainment deciles.²⁷ In previous cohorts, these students may not have been able to secure a place on an undergraduate course, or they may only have been admitted to courses in locations or fields that they disliked. The reforms have a limited impact on their returns to higher education, since their expected earnings are significantly lower than their higher attainment peers. As a result, the financing reforms have almost no impact on overall participation, but do result in a fairly big shift in the composition of students going to university.

The red dots show the simulated effects of the reform based on our model. While the model effects are slightly smaller in the middle of the distribution, it is able to replicate the overall pattern of the reform effects very well, which is reassuring as the moments are not included in the

²⁷We find that this result is extremely robust, including to more sophisticated approaches, such as including decile-specific time trends with post-treatment dummies and including additional control variables.

estimation of the model. We also show in Figure 9 that the model is able to replicate the effects amongst higher ability students by parental SES. This is particularly reassuring for the counterfactual policy simulations that we present in the following section, as our targeted interventions are focused on lower SES individuals.



Note: The data shows the raw changes in participation between 2010 and 2012, adjusting for overall growth in participation rates.

7 Policy experiments

We simulate five hypothetical counterfactual policy reforms, with overarching aim of understanding the implications of each of them for intergenerational mobility:²⁸

- **Policy 1:** No student loans for low SES students - tuition fees set to £0 and maintenance loans converted to grants (but held at current levels).
- **Policy 2:** Additional maintenance grants of £5,000 per year, on top of existing loans, for low SES (bottom 50%) students.

²⁸For our baseline we use the 2009 higher education cohort assuming they faced the post-2012 higher education system. We hold everything fixed across the reforms, including the number of places available on each university course. The policies evaluated here are not necessarily reflective of actual government policy considerations.

- **Policy 3:** As for Policy 2, but only if attending a high quality university.
- **Policy 4:** As for Policy 2, but only if studying high earning (STEM and LEM) courses.
- **Policy 5:** A ‘10% rule’ of preferential admission to those who graduated in the top 10% of their high school class (conditional on being low SES).

The outcomes from these simulations are summarised in Table 8 and Table 9, which show effects on participation decisions and on longer-run outcomes, respectively. Table 8 shows effects on overall participation, the subject mix (conditional on attendance) and the shares attending high quality courses (also conditional on attendance). It splits the population in the bottom and top half of the SES distribution in the top panel, and does the same for ‘high ability’ students (as measured by being in the top 20% for average communication and quantitative skills) in the bottom panel.²⁹

Policy 1, which involves free fees and maintenance support for SES students, results in a small increase in overall participation amongst poorer students, by 0.6ppts. Much of this increase is driven by higher ability students, for whom participation increases by 2.6 percentage points. Since we do not allow overall students to increase across the reforms, there is a corresponding crowding-out of higher SES students, for whom participation drops by 0.6 percentage points. However, the drop amongst high ability, high SES students is small, meaning the policy results in higher ability students from poorer backgrounds displacing predominantly lower ability students from wealthier backgrounds.

Nevertheless, the top two panels of Table 9 shows the implications of this for intergenerational mobility are small. The first set of estimates show that Policy 1 would reduce the overall earnings gap between high and low SES students relative to the baseline, but only by a small amount, from 29.3% to 28.5%. Between very low SES (bottom quintile) and very high SES (top quintile), it would reduce the earnings gap from 56.5% to 55.7%. The next set of rows consider a similar definition of mobility to that used in Chetty et al. (2020), namely entry to the top 20% of the earnings distribution in adulthood.³⁰ We estimate that Policy 1 would narrow the overall gap

²⁹30% of those we define as high ability are ‘low SES’ (i.e. from the bottom 50% of the parental SES index). 9% are ‘very low SES’ (i.e. from the bottom 20%).

³⁰The estimates here are not directly comparable to those used by Chetty et al. (2020), as they focus only on the set of people who attend university, while we consider policies that affect the decision of whether to attend as well as where to attend.

in entry by just 1.2%. Finally, average earnings of low and high SES individuals would also not change much. In the second panel we repeat the estimates for high ability students, and again the impacts of Policy 1 are small: earnings gaps decline by only 1 or 2 percentage points, while the gap in entry to the top 20% of the earnings distribution declines by 9%. Earnings gains for high ability, low SES students would outweigh earnings losses for high ability, high SES students, but the overall effects on average earnings are also small. The final set of estimates in the bottom panel of Table 9 show that these small gains in terms of mobility would come with a large associated cost of around £1 billion per cohort of students.³¹ Almost all of this is due to long run losses in revenue from student loan repayments (there is a small up front cost as up-front government grants are slightly larger for poorer students).

Policy 2, which increases the cash support available to students while at university, has similar impacts to Policy 1. It is slightly more successful at increasing the participation of low SES students, especially high ability low SES students, but the effects on earnings gaps and mobility rates are similarly small. The long run cost of the policy is also high at £950 million per cohort. Unlike Policy 1, this is driven by higher upfront costs, with additional spending of £5,000 per low SES student, per year. In fact, there is a long-run benefit to the taxpayer in terms of tax and loan repayment receipts - driven by the increase in average ability amongst students under this policy - but this is minor compared to the upfront cost.

Policy 3 tweaks Policy 2 so that the same sized grants are now only available to people studying on a 'high quality' course. The intention of this is to investigate whether grants could be used to address the issue of gaps in academic undermatch that have been documented in recent research (Wyness et al., 2021). And although the policy has almost no effect on overall participation (highlighting that Policy 2 was drawing poorer students into lower-quality institutions), it does indeed increase the share of poorer students attending high quality institutions from 13.6% to 16.4% overall and from 28.7% to 34.4% amongst high ability students. This means that the policy almost completely removes the gaps in academic undermatch between high ability, high and low SES students (as around 35.9% of high ability high SES student attend high quality courses).

³¹This is approximately 15% of the per-cohort long run cost to government of student loans in 2022 (according to the Institute for Fiscal Studies Student Finance Calculator, using the ONS indicators and a discount rate of 3.5%).

Table 8: Counterfactual policy reforms, participation effects

	Baseline (2012 entrants)	Policy 1 (Free fees & maint.)	Policy 2 (Extra grants)	Policy 3 (High qual grants)	Policy 4 (Subj. grants)	Policy 5 (Percent plan)
<i>Low SES (bottom 50%) students</i>						
Overall	23.8	24.4	25.0	23.8	24.5	25.2
STEM Attend	43.1	43.7	43.4	43.8	54.5	43.1
LEM Attend	20.5	21.0	21.1	20.1	27.9	19.8
AHSS Attend	36.4	35.3	35.5	36.1	17.5	37.1
High Qual Attend	13.6	13.6	12.9	16.4	13.3	27.7
<i>High SES (top 50%) students</i>						
Overall	40.6	40.0	39.3	40.5	39.9	39.2
STEM Attend	42.4	42.1	42.2	42.0	35.4	42.4
LEM Attend	17.2	16.9	16.8	17.5	12.6	17.6
AHSS Attend	40.3	41.1	41.0	40.5	52.0	40.0
High Qual Attend	24.1	24.2	24.8	22.5	24.4	15.4
<i>High-ability, low SES students</i>						
Overall	79.9	82.5	84.9	80.1	83.8	89.1
STEM Attend	44.9	45.8	45.5	46.8	61.9	45.1
LEM Attend	17.2	17.8	17.4	16.2	25.6	16.4
AHSS Attend	37.9	36.4	37.1	37.0	12.6	38.4
High Qual Attend	28.7	28.7	27.1	34.4	27.3	53.7
<i>High-ability, high SES students</i>						
Overall	81.7	81.4	80.9	81.6	81.5	79.7
STEM Attend	43.8	43.7	43.7	43.4	38.9	43.3
LEM Attend	15.8	15.7	15.7	16.0	13.2	16.2
AHSS Attend	40.4	40.6	40.6	40.6	47.8	40.5
High Qual Attend	38.3	38.1	38.7	35.9	38.4	24.9

Note: All estimates are from model simulations. High ability includes those in the top 20% of skills (combined). STEM, LEM and AHSS are the subject areas, while High Qual indicates attending high quality (top quintile) degree.

Again, however, there is very little impact in terms of intergenerational mobility - in fact, this is the policy with the smallest impact in terms of average earnings gaps and gaps in entry to the top of the earnings distribution. This suggests that policies which focus on reducing undermatch may not have much long-run impact.³²

Policy 4, which instead targets the grants at those who choose to study STEM courses, has a larger effect. Overall participation of poorer students increases by half a percentage point, but the major effect is on the subject margin, with a big shift towards STEM (43.1% to 54.5%) and LEM (20.5% to 27.9%) subjects amongst poorer students, with corresponding reductions in the shares studying AHSS courses and in higher SES students studying STEM and LEM (as we hold student

³²We also consider a more complex definition of undermatch that follows Wyness et al. (2021) and focuses on very high ability individuals, and we draw the same conclusions.

numbers fixed at the course level). The effects for high ability students are larger still. In terms of longer run outcomes in Table 9, we start to see some impacts. The overall earnings gap between low and high SES students falls from 29.3% to 25.9%, while the gap between very low and very high SES students falls from 56.5% to 53.1%. The mobility rate gap narrows by 4.7%. The effects for high ability students are much starker, however: the policy almost completely removes the gap in earnings between high ability low and high SES students, and reduces the earning gap between very low and very high SES students from 21.4% to 13.4%. It also closes the mobility rate gap amongst high ability students by around a quarter. Average earnings of high ability low SES students increase by 5.3% relative to the baseline, while the cost for high ability high SES students is a drop in average earnings of 1.1% relative to the baseline. The upfront cost of the policy is quite high at £780 million per cohort (around 10% of the per-cohort cost of student loans), although this is offset slightly by boosts to tax and loan revenues. Importantly, however, when contrasted with Policy 3, these results suggest that crossing subject margins is more important than focussing on moving poorer students up the university quality distribution (in fact, the undermatch gap *increases* as a result of this policy). It also emphasises the importance of preferences in determining people's education paths, as we see that poorer students would be better off taking STEM and LEM subjects, and would potentially be accepted to study them, but a large share are choosing not to take that path.³³

Finally, for Policy 5 we move away from demand-side reforms and consider a policy that forces universities to give priority admissions to low SES students who are in the top 10% of their secondary school class.³⁴ Although this policy does not have a dramatic impact on the overall participation rate of low SES students, it substantially increases participation of high ability, low SES students (from 79.9% to 89.1%) and the share of low SES students attending high quality courses (13.6% to 27.7% overall, and from 28.7% to 53.7% amongst higher ability students).³⁵ This policy

³³In Appendix Tables A4 and A5, we show estimates of the effects of Policy 4 when it is further conditioned to so only high ability students are eligible. Costs of this policy would be much lower at just under £300 million, while still achieving most of the benefits in terms of mobility improvements. This simulation also reveals that improvements in tax revenues from the subject grant policy is driven by higher ability students switching fields.

³⁴In Appendix Tables A4 and A5 we also show the impacts of an "unconditional" percent plan that involves forces priority admission for students in the top 10% of their class, without requiring students to be low SES in order to be eligible. The impacts on mobility are non-negligible, but quite a lot smaller.

³⁵While these are large effects, we find that the composition of students on any individual course is not completely transformed (see Appendix Figure A7). For example, the share of eligible students on the highest quality courses increases from around 5% to around 30% as a result of the policy.

Table 9: Counterfactual policy reforms, long run effects

	Baseline (2012 entrants)	Policy 1 (Free fees & maint.)	Policy 2 (Extra grants)	Policy 3 (High qual grants)	Policy 4 (Subj. grants)	Policy 5 (Percent plan)
<i>Earnings gap (%), between:</i>						
Low SES and High SES	29.3	28.5	28.0	28.9	25.9	24.5
V. Low SES and V. High SES	56.5	55.7	55.3	55.8	53.1	50.2
<i>Mobility rate (entry to top 20% of earnings distn.)</i>						
V. Low SES to top 20% (earns)	11.8	11.8	11.9	11.8	12.0	12.3
V. High SES to top 20% (earns)	28.1	28.0	27.9	28.1	27.6	27.5
Narrowing of gap (%), rel. to baseline		1.2	2.3	0.5	4.7	7.2
<i>Av. earnings change, rel. to baseline (%):</i>						
Low SES		0.3	0.6	0.2	1.5	2.1
High SES		-0.3	-0.4	-0.1	-1.1	-1.7
<i>High-ability earnings gap (%), between:</i>						
Low SES and High SES	7.8	6.5	6.0	6.4	1.2	-4.9
V. Low SES and V. High SES	21.4	19.7	19.7	19.3	13.4	5.9
<i>High-ability mobility rate (entry to top 20%)</i>						
V. Low SES to top 20% (earns)	31.7	32.3	32.2	32.2	33.3	35.9
V. High SES to top 20% (earns)	40.2	40.1	40.0	40.1	39.7	38.9
Narrowing of gap (%), rel. to baseline		8.9	8.1	7.2	24.6	65.3
<i>High-ability earnings change, rel. to baseline (%):</i>						
Low SES		1.0	1.5	0.9	5.3	9.5
High SES		-0.2	-0.2	-0.4	-1.1	-3.4
<i>Fiscal costs (£millions, rel. to baseline)</i>						
Upfront grant cost (A)		6.5	971.1	105.9	780.1	15.7
Long run lost tax/loan receipts (B)		987.9	-18.7	0.7	-18.7	58.7
Total long run cost (=A+B)		994.4	952.5	106.6	761.4	74.4

Note: Earnings outcomes are based on earnings at age 40 (overall mean £24500; mean for high ability £27750; 2019 prices). Low and High SES are bottom 50% and top 50% respectively (as elsewhere in the paper), while V Low and V High are bottom 20% and top 20% respectively. For the fiscal calculations we assume a cohort size of 530,000 and government discount rate of 3.5%, following UK government accounting conventions (noting this is different to the 5% rate applied in the model, as the fiscal calculation is from the point of view of government, which we assume to discount at a lower rate than individual students). Also see notes to Table 8.

also has the largest impacts in terms of mobility. Again the effects are most pronounced when focussing on high ability students: the policy reduces the gap in earnings between high and low SES, high ability students from 7.8% to -4.9%, and the gap between very low and very high SES, high ability students from 21.4% to just 5.9%. It also reduces gap in entry to the top 20% of the earnings distribution by two-thirds and increases average earnings of high ability, low SES students by almost 10% (while the cost to high ability, high SES student is 3.4%).

The long run cost of this policy would be around £75 million per cohort. Some of this is due

to increased upfront spending on grants (due to the increase in participation of poorer students), but mostly it is due to efficiency losses: that is, tax and student loan receipts are lower as a result of the policy, highlighting that replacing richer students with poorer students in top universities would lower average earnings. However, it is reasonable to assert that this cost is small, at only around 1% of the long run per-cohort costs of student loans under the baseline system, and the lowest cost of all the policies considered.

8 Conclusion

In this paper we develop and estimate a novel empirical matching model of sorting in the UK's higher education system. The model is able to replicate sorting patterns in the data extremely well, both for the markets included in estimation and for markets excluded from the estimation. It also does an excellent job of replicating descriptive patterns observed from the 2012 reforms.

The ability of our model to emulate these patterns enables us to confidently simulate counterfactual policies. We do this with the overarching aim of trying use higher education policy levers to boost intergenerational mobility. In general, we find that the scope for the higher education system to substantially improve intergenerational mobility statistics for entire cohorts is limited. However, when we zoom in on students who are higher ability, we do see substantial effects in some cases.³⁶ Contrary to some of the perceived wisdom in higher education policy, we find that these substantial effects do not come from policies that cut students loans for poorer students, or from policies that increase cash support for poorer students attending university or attending high-status universities. Instead, our estimates suggests that of the potential set of policies that focus on the demand side of the market, targeting cash support at poorer students conditional on studying higher earning subjects would have the greatest impacts. However, our largest impacts in terms of boosting mobility, come not from a policy targeting the demand side, but rather from an aggressive reform to the supply side.³⁷ Our simulated effects of forcing universities to give priority admission to poorer students that score in the top 10% of their secondary school class suggest

³⁶This result is not inconsistent with Chetty et al. (2020), which focusses on mobility only amongst the set of people who go to university.

³⁷There would potentially be even larger gains from interacting the supply side policy with the subject grants policy, although we do not model this.

substantial improvements in terms of intergenerational mobility, at least conditional on the set of people who leave school with reasonably high levels of attainment.³⁸ Although this policy would involve a decline in the overall efficiency of the match between students and universities, we estimate that the fiscal implications of this would be small.

We therefore conclude that - at least for the period we study - gaps in skills between those from wealthier and poorer backgrounds, and, crucially, strong university preferences for those skills, are a key barrier for intergenerational mobility. Policies which focus on reducing skill gaps should remain central. However, we also find evidence to suggest that higher education policy, even holding gaps in skills and educational content on the courses fixed, can be highly effective at boosting mobility. This is promising for policymakers. Further exploration of policy levers to encourage poorer students to apply to subjects that lead to better earnings outcomes is would be welcome. Finally, universities are likely to respond to our recommendation to consider reforms to admissions policies by referring to the contextualised admissions policies that have been scaled up in recent years. Assessment of the efficacy of these policies and exploration of whether there is scope to do more with them is therefore a crucial next step.

References

- Agarwal, Nikhil**, “An empirical model of the medical match,” *American Economic Review*, 2015, 105 (7), 1939–78.
- Altmejd, Adam, Andrés Barrios-Fernández, Marin Drlje, Joshua Goodman, Michael Hurwitz, Dejan Kovac, Christine Mulhern, Christopher Neilson, and Jonathan Smith**, “O brother, where start thou? Sibling spillovers on college and major choice in four countries,” *The Quarterly Journal of Economics*, 2021, 136 (3), 1831–1886.
- Andrews, Rodney J, Scott A Imberman, and Michael F Lovenheim**, “Risky Business? The Effect of Majoring in Business on Earnings and Educational Attainment,” Technical Report, National Bureau of Economic Research 2017.

³⁸Although an extreme case, this policy is quite similar in spirit to contextualised admission policies, whereby universities give preferential admission to poorer students.

- Anelli, Massimo**, “Returns to elite university education: a quasi-experimental analysis,” *Journal of the European Economic Association*, 2018.
- Arcidiacono, Peter**, “Ability sorting and the returns to college major,” *Journal of Econometrics*, 2004, 121 (1), 343–375.
- , “Affirmative action in higher education: How do admission and financial aid rules affect future earnings?,” *Econometrica*, 2005, 73 (5), 1477–1524.
- Azmat, Ghazala and Ștefania Simion**, “Charging for Higher Education: Estimating the Impact on Inequality and Student Outcomes,” *The BE Journal of Economic Analysis & Policy*, 2020, 1 (ahead-of-print).
- Belfield, Chris, Jack Britton, Franz Buscha, Lorraine Dearden, Matt Dickson, Laura Van Der Erve, Luke Sibieta, Anna Vignoles, Ian Walker, and Yu Zhu**, “The relative labour market returns to different degrees,” 2018.
- , —, —, **Lorraine Dearden, and Laura van der Erve**, “Higher education funding in England: Past, present and options for the future,” *IFS Briefing Note BN211*, 2017.
- Black, Dan A and Jeffrey A Smith**, “Estimating the returns to college quality with multiple proxies for quality,” *Journal of Labor Economics*, 2006, 24 (3), 701–728.
- Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw**, “Female labor supply, human capital, and welfare reform,” *Econometrica*, 2016, 84 (5), 1705–1753.
- Britton, Jack, Laura van der Erve, Chris Belfield, Anna Vignoles, Matt Dickson, Yu Zhu, Ian Walker, Lorraine Dearden, Luke Sibieta, and Franz Buscha**, “How much does degree choice matter?,” *IFS Working Paper w21/24*, 2021.
- Broecke, Stijn**, “University selectivity and earnings: Evidence from UK data on applications and admissions to university,” *Economics of Education Review*, 2012, 31 (3), 96–107.
- Chetty, Raj, John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan**, “Income Segregation and Intergenerational Mobility Across Colleges in the United States,” *Quarterly Journal of Economics*, 2020.

- Chevalier, Arnaud**, "Subject choice and earnings of UK graduates," *Economics of Education Review*, 2011, 30 (6), 1187–1201.
- Chowdry, Haroon, Lorraine Dearden, Wenchao Jin, and Barnaby Lloyd**, "Fees and student support under the new higher education funding regime: what are different universities doing?," *Institute for Fiscal Studies Briefing Note*, 2012.
- Corak, Miles**, "Income inequality, equality of opportunity, and intergenerational mobility," *Journal of Economic Perspectives*, 2013, 27 (3), 79–102.
- Dale, Stacy B and Alan B Krueger**, "Estimating the effects of college characteristics over the career using administrative earnings data," *Journal of human resources*, 2014, 49 (2), 323–358.
- Dale, Stacy Berg and Alan B Krueger**, "Estimating the payoff to attending a more selective college: An application of selection on observables and unobservables," *The Quarterly Journal of Economics*, 2002, 117 (4), 1491–1527.
- Dearden, Lorraine, Emla Fitzsimons, and Gill Wyness**, "Money for nothing: Estimating the impact of student aid on participation in higher education," *Economics of Education Review*, 2014, 43, 66–78.
- Delavande, Adeline and Basit Zafar**, "University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints," *Journal of Political Economy*, 2019, 127 (5), 2343–2393.
- Denning, Jeffrey T, Benjamin M Marx, and Lesley J Turner**, "ProPelled: The effects of grants on graduation, earnings, and welfare," *American Economic Journal: Applied Economics*, 2019, 11 (3), 193–224.
- Diamond, William and Nikhil Agarwal**, "Latent indices in assortative matching models," *Quantitative Economics*, 2017, 8 (3), 685–728.
- Dillon, Eleanor Wiske and Jeffrey Andrew Smith**, "The consequences of academic match between students and colleges," *Journal of Human Resources*, 2020, 55 (3), 767–808.
- Epple, Dennis, Richard Romano, and Holger Sieg**, "Admission, tuition, and financial aid policies in the market for higher education," *Econometrica*, 2006, 74 (4), 885–928.

- French, Eric and John Bailey Jones**, “On the distribution and dynamics of health care costs,” *Journal of Applied Econometrics*, 2004, 19 (6), 705–721.
- Fu, Chao**, “Equilibrium tuition, applications, admissions, and enrollment in the college market,” *Journal of Political Economy*, 2014, 122 (2), 225–281.
- Hastings, Justine, Christopher A Neilson, and Seth D Zimmerman**, “The effects of earnings disclosure on college enrollment decisions,” Technical Report, National Bureau of Economic Research 2018.
- Hastings, Justine S, Christopher A Neilson, and Seth D Zimmerman**, “Are some degrees worth more than others? Evidence from college admission cutoffs in Chile,” Technical Report, National Bureau of Economic Research 2013.
- Hübner, Malte**, “Do tuition fees affect enrollment behavior? Evidence from a ‘natural experiment’ in Germany,” *Economics of Education Review*, 2012, 31 (6), 949–960.
- Kapor, Adam**, “Distributional effects of race-blind affirmative action,” *Job Market Paper*, Yale University, 2020.
- Keane, Michael P and Kenneth I Wolpin**, “The career decisions of young men,” *Journal of political Economy*, 1997, 105 (3), 473–522.
- and —, “The effect of parental transfers and borrowing constraints on educational attainment,” *International Economic Review*, 2001, 42 (4), 1051–1103.
- Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad**, “Field of study, earnings, and self-selection,” *The Quarterly Journal of Economics*, 2016, 131 (3), 1057–1111.
- Marx, Benjamin M and Lesley J Turner**, “Borrowing trouble? human capital investment with option costs and implications for the effectiveness of grant aid,” *American Economic Journal: Applied Economics*, 2018, 10 (2), 163–201.
- Mountjoy, Jack and Brent Hickman**, “The Returns to College (s): Estimating Value-Added and Match Effects in Higher Education,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020.

- Neill, Christine**, "Tuition fees and the demand for university places," *Economics of Education Review*, 2009, 28 (5), 561–570.
- Sá, Filipa**, "The effect of university fees on applications, attendance and course choice: evidence from a natural experiment in the uk," *Economica*, 2019, 86 (343), 607–634.
- Scott-Clayton, Judith and Basit Zafar**, "Financial aid, debt management, and socioeconomic outcomes: Post-college effects of merit-based aid," *Journal of Public Economics*, 2019, 170, 68–82.
- Skron dal, Anders and Sophia Rabe-Hesketh**, *Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models*, Chapman and Hall/CRC, 2004.
- Wiswall, Matthew and Basit Zafar**, "Determinants of college major choice: Identification using an information experiment," *The Review of Economic Studies*, 2014, 82 (2), 791–824.
- Wyness, G, S Campbell, L Macmillan, and R Murphy**, "Matching in the dark? Inequalities in student to degree match," *Journal of Labor Economics*, 2021.
- Zimmerman, Seth D**, "Elite colleges and upward mobility to top jobs and top incomes," *American Economic Review*, 2019, 109 (1), 1–47.

Appendix

A Subject groupings

Table A1: Subject groupings

Subject	Subject group	CAH2 code and description
Agriculture	STEM	(CAH06-01) agriculture, food and related studies
Allied to med	STEM	(CAH02-03) subjects allied to medicine not otherwise specified
Architecture	STEM	(CAH13-01) architecture, building and planning
Biosciences	STEM	(CAH03-01) biosciences
Business	LEM	(CAH17-01) business and management
Chemistry	STEM	(CAH07-02) chemistry
Comms	Other	(CAH18-01) communications and media
Computing	STEM	(CAH11-01) computing
Creative arts	Other	(CAH21-01) creative arts and design
Economics	LEM	(CAH15-02) economics
Education	Other	(CAH22-01) education and teaching
Engineering	STEM	(CAH10-01) engineering
English	Other	(CAH19-01) English studies
Geography	STEM	(CAH12-01) geographical and environmental studies
History	Other	(CAH20-01) history and archaeology
Languages	Other	(CAH19-03) languages, linguistics and classics
Law	LEM	(CAH16-01) law
Maths	STEM	(CAH09-01) mathematical sciences
Medicine	STEM	(CAH01-01) medicine and dentistry
Nursing	STEM	(CAH02-01) nursing
Pharmacology	STEM	(CAH02-02) pharmacology, toxicology and pharmacy
Philosophy	Other	(CAH20-02) philosophy and religious studies
Physics	STEM	(CAH07-01) physics and astronomy
Physsci	STEM	(CAH07-03) physical, material and forensic sciences
Politics	Other	(CAH15-03) politics
Psychology	STEM	(CAH04-01) psychology
Social care	Other	(CAH15-04) health and social care
Sociology	Other	(CAH15-01) sociology, social policy and anthropology
Sportsci	STEM	(CAH03-02) sport and exercise sciences
Technology	STEM	(CAH10-02) technology
Vetsci	STEM	(CAH05-01) veterinary sciences

B Individual skill measurement

We estimate two skill variables, summarising each individuals quantitative and communication skills. These skill measures aren't directly observed in our data but we assume that exam results, which we do observe, are noisy indicators of these underlying skills. We first estimate a set of structural equations capturing the relationship between the latent skills and observable exam results. Then, given each individuals exam results, we predict their underlying skills. For more details on the approach applied, see Skrondal and Rabe-Hesketh (2004).

The skill measures are based on two sets of exam results. The first set of exams are the Key

Stage 2 *Standard Assessment Tests* (SATs). These exams are taken by all English students at age 11. In our data, we can observe the numerical score in three exams:

- Mathematics
- Reading
- Writing

We supplement these measures with *General Certificate of Secondary Education* (GCSE) grades. GCSEs are sat at age 16. Most students sit between 8 and 11 GCSEs across a range of subjects. The only compulsory subjects are Mathematics, English Language and Science, but most schools require that students also sit English Literature, at least one GCSE in a humanities subject (Geography or History), one in modern foreign languages subject (such as French or Spanish) and one in an arts subject (such as Music or Drama). We observe the letter grade (A* to G) for each GCSE. We use the following exams when estimating skills:

- English Language
- English Literature
- Mathematics
- Science
- Humanities

where humanities is measured as the maximum of grades in Geography and History. We choose this relatively small subset of GCSE subjects to minimise missing data, which would complicate both estimation and prediction.

We model each grade as a function of (latent) quantitative and communication skill:


$$M_i^k = \zeta_0^k + \zeta_m^k S_i^q + \zeta_c^k S_i^c + \epsilon_i^k$$

where k indexes the measure, S_i^q is quantitative skill and S_i^c is communication skill. S_i^q and S_i^c are assumed to be jointly normally distributed. ϵ_i^k is unobserved and assumed to be uncorrelated

across measures. We also assume that the variance and covariance of the underlying latent variables are the same for all cohorts, but allow the parameters of each measurement equation (ζ_0^k , ζ_m^k , ζ_c^k and σ_{ϵ^k}) to differ between cohorts. This allows us to capture grade inflation.

To identify the system of measurement equations we need to make some additional exclusion restrictions. We exclude quantitative skill from the measurement equation for English Language or English Literature GCSEs, and the communication skill from the measurement equation for Mathematics GCSE. Parameters are estimated using full information maximum likelihood. Once we have estimated the parameters, we predict the skill measures for each individual using the empirical Bayesian modal approach detailed in Skrondal and Rabe-Hesketh (2004).

Our method allows for our measures to be highly correlated with one another: as see in Figure A1, those with high quantitative skills generally also have high communication skills, and vice versa.



Note: Based on the 2005/06 school leavers. We have added a random jitter to the plot to remove disclosivity risks, without losing meaning.

C Course quality measurement

Course quality is a summary measure derived using the following variables:

- **Academic services spend:** Expenditure on all academic services (such as library and computing facilities) divided by the number of full-time equivalent students.
- **Facilities spend:** Expenditure on all student facilities (such as sports, careers services, counselling) divided by the number of full-time equivalent students.
- **Student-staff ratio:** The total number of undergraduate and postgraduate students divided by the number of academic staff.
- **Student satisfaction:** Average across all satisfaction scores as measured in the National Student Survey.
- **Research quality:** Score of the relevant department in the Research Excellence Framework.

Data is collected from the *Complete University Guide*. Academic services spend, facilities spend and student-staff ratio are measured at the university level. Research quality and student satisfaction are both measured at the course level, varying across fields within each university. Fields within the *Complete University Guide* data are more disaggregated than the three broad fields that we consider in our analysis (STEM, LEM and AHSS). To create measures defined for our broad fields, we take a weighted average over the disaggregated fields using number of undergraduates enrolled in the relevant field at the relevant university in 2010. Finally, we take the first principal component of these five variables as our measure of course quality. The weights assigned to each of the variables (after standardisation) are reported below:

Variable	Weight in 1st principal component
Academic services spend	0.483
Facilities spend	0.397
Student-staff ratio	-0.545
Student satisfaction	0.195
Research quality	0.523

D Life-cycle estimation

Several parameters governing the life-cycle earnings profiles associated with each field are estimated prior to the main estimation exercise. As discussed in Section 4, the earnings equation in

the model is:

$$\begin{aligned} \ln(y_{i,t}^e) = & \alpha_{0F} + \alpha_{1F}X_i + \alpha_{2F}S_i^q + \alpha_{3F}S_i^c + \alpha_{4F}Q_jH_i + \\ & \alpha_{5F}Q_jH_iS_i^q + \alpha_{6F}Q_jH_iS_i^c + \alpha_{7F}\ln(t+1) + \alpha_\theta\theta_i + \epsilon_{i,t} \end{aligned}$$

The unobserved shock, $\epsilon_{i,t}$, follows an AR(1) process:

$$\epsilon_{i,t} = \rho_F\epsilon_{i,t-1} + \zeta_{F,i,t}$$

where $\zeta_{F,i,t}$ is an exogenous iid shock. Each period, individuals can either be employed ($l = 1$) or unemployed ($l = 0$). Agents have an exogenous probability of receiving a job offer in any given period, v_F , which depends on their field of study. If no job offer is received, pre-tax earnings are zero:

$$y_{i,t} = ly_{i,t}^e$$

$$\Pr(l = 1) = v_F$$

Prior to estimation of the matching model, the following life-cycle earnings parameters are estimated:

- life-cycle average earnings growth, α_{7F}
- average employment rate, v_F
- persistence of productivity shock, ρ_F
- variance of the iid shock, $\sigma_{\zeta,F}^2$

These parameters all vary by field and gender.

Given our assumptions, the growth of earnings over the life-cycle (α_{7F}) can be estimated consistently outside the model. Crucially, while we allow for selection into fields on earnings levels, we assume that there is no selection into fields on earnings growth. We also assume that employment is random conditional on field. We can therefore consistently estimate α_{7F} using OLS on the

sample of employed graduates for each field and gender:

$$\ln(y_{i,t}^e) = \beta_{0F} + \beta_{1F} \ln(t+1) + \omega_{i,t}$$

where β_{1F} is a consistent estimator for α_{7F} . Similarly, we can estimate v_F as the average employment rate for graduates of each field and gender. The only information that we require for this exercise is gender, field of study, employment and earnings for the employed. Since the earnings panel in our main sample is truncated at age 30, we supplement our sample using additional data documented in detail in Britton et al. (2020). In brief, we extend the sample by incorporating individuals with a HESA record but no school record (up to age 40) and individuals drawn from the Labour Force Survey (up to retirement). This allows us to ensure that the estimated earnings growth and employment rates are representative of earnings growth and employment over the whole of working life.

Once we have estimated the growth rate of earnings over working life, we can construct the estimated first difference of the residual:

$$\Delta \hat{\epsilon}_{i,t} = \Delta \ln(y_{i,t}^e) - \hat{\alpha}_{7F} \Delta \ln(t+1)$$

Since $\epsilon_{F,i,t}$ is AR(1) and assuming that the variance of $\epsilon_{F,i,t}$ is zero at the beginning of working life, we can derive:

$$\begin{aligned} V(\Delta \epsilon_{i,t}) &= \frac{\sigma_{\xi,F}^2 (1 - \rho_F)^2 (1 - \rho_F^{2(t-1)})}{1 - \rho_F^2} + \sigma_{\xi,F}^2 \\ C(\Delta \epsilon_{F,i,t}, \Delta \epsilon_{F,i,t-k}) &= \rho_F^k \frac{\sigma_{\xi,F}^2 (1 - \rho_F)^2 (1 - \rho_F^{2(t-1)})}{1 - \rho_F^2} - (1 - \rho_F) \rho_F^{k-1} \sigma_{\xi,F}^2 \\ &= \rho_F^k V(\Delta \epsilon_{F,i,t-k}) - \rho_F^{k-1} \sigma_{\xi,F}^2 \end{aligned}$$

We estimate ρ_F and $\sigma_{\xi,F}^2$ using a minimum distance estimator matching these moments. We construct the moments separately for each field and gender using the sub-sample of individuals who are aged between 25 and 30 and are employed in both the current and prior period (and for whom we therefore can construct $\Delta \epsilon_{f,i,t}$). We match all elements of the auto-covariance matrix for $\Delta \epsilon_{f,i,t}$, which contains 21 moments. Standard errors for the auto-covariance matrix are estimated using

bootstrapping, and we weight according to the inverse diagonal of the variance of the moments.

Table D reports the estimates for all of the life-cycle parameters. These values are used as inputs to the matching model.

Table A2: Estimates for life-cycle parameters

Parameter	Men				Women			
	No Uni	STEM	LEM	AHSS	No Uni	STEM	LEM	AHSS
α_{7F}	0.255	0.494	0.466	0.468	0.102	0.311	0.268	0.253
v_F	0.814	0.814	0.804	0.794	0.834	0.846	0.895	0.840
ρ_F	0.746	0.824	0.812	0.801	0.796	0.789	0.812	0.806
$\sigma_{\xi,F}^2$	0.122	0.140	0.124	0.145	0.116	0.144	0.121	0.139

E Moments for estimation of the matching model

Each of the subsections below describes a category of moments that we match when estimating the full model. The relevant variables are defined in detail within the main body of the paper. We briefly summarise these definitions below:

- X_i is a vector of individual characteristics, including gender, socio-economic status and a private school indicator
- S_i^q, S_i^c are skill measures, discussed in Appendix B
- Z_i is a vector of individual instrumental variables, discussed in Section 5.2
- Q_j is course quality, discussed in Appendix C
- H_i is an indicator for enrollment in university
- $y_{i,t}$ is earnings at age t
- $\mu(\cdot)$ is the matching function, such that $\mu(i)$ indicates the course attended by individual i and $\mu^{-1}(j)$ indicates the set of individuals that attend course j

In addition, define the following variables:

- H_{iF} is an indicator for enrollment in field F

- $d_i, \mu(i)$ is the distance between the school attended by individual i and the course attended $\mu(i)$
- \bar{y}_i is average earnings between the ages of 25 and 30
- N_j is the number of individuals enrolled at university
- N_F is the number of individuals enrolled at university in field F

All moments are estimated for each of the four cohorts included in the estimation data. We then solve the model separately for each of these four cohorts and calculate the corresponding simulated moments.

Auxiliary earnings regressions

For all earnings moments, we strip out the time trend in earnings by field and gender, which we estimate prior to estimating the main model (see Appendix D):

$$\tilde{y}_{i,t} = y_{i,t} - \hat{\alpha}_{7F} \ln(t+1)$$

We include auxiliary regressions by field using the sample of **all employed individuals** relating earnings to individual covariates and the characteristics of the course they are matched to.

$$\tilde{y}_{i,t} = \beta_{0F} + \beta_{1F}X_i + \beta_{2F}S_i^q + \beta_{3F}S_i^c + \beta_{4F}Q_jH_i + \beta_{5F}Q_jH_iS_i^q + \beta_{6F}Q_jH_iS_i^c$$

This specification is chosen to mirror the earnings equation in the model, but it does not account for unobserved heterogeneity.

In addition, we include an auxiliary regression using the sample of **all employed individuals** relating earnings to individual characteristics, skills and instruments.

$$\tilde{y}_{i,t} = \beta_0 + \beta_1X_i + \beta_2S_i^q + \beta_3S_i^c + \beta_4Z_i$$

Auxiliary HE outcome regressions

We include an auxiliary regression using the sample of **all university enrollees** relating quality of university to individual characteristics, skills and instruments.

$$Q_{\mu(i)} = \beta_0 + \beta_1 X_i + \beta_2 S_i^q + \beta_3 S_i^c + \beta_4 Z_i$$

We also include auxiliary regressions using the sample of **all individuals** relating enrollment in each specific field with individual characteristics, skills and instruments.

$$H_{iF} = \beta_0 + \beta_1 X_i + \beta_2 S_i^q + \beta_3 S_i^c + \beta_4 Z_i$$

Covariance between individual and course characteristics

We include the covariance between individual characteristics, skills and instruments and the field and quality of the matched course. We calculate the covariances using the sample **all enrollees**.

$$\frac{1}{N_J} \sum_i (k_i - \bar{k})(l_{\mu(i)} - \bar{l})$$

for all $k_i \in \{X_i, S_i^q, S_i^c, Z_i, \bar{y}_i\}$, $l_{\mu(i)} \in \{H_{\mu(i)F}, Q_{\mu(i)}\}$.

Variance by field

We include the variance of individual characteristics, skills and earnings by field, including those who didn't enrol in higher education as a separate field. We calculate the variances using the sample of **all individuals**.

$$\frac{1}{N_F} \sum_i H_{iF} (k_i - \bar{k}_F)^2$$

for all $k_i \in \{X_i, S_i^q, S_i^c, \bar{y}_i\}$ and $F \in \{\text{No Uni}, \text{STEM}, \text{LEM}, \text{AHSS}\}$.

Within-course variance by field

We include the within-course variance of individual characteristics, skills and earnings by field. We calculate the variance using data on **enrolees only**, taking out the course specific mean of each characteristic. Let $\bar{k}_{\mu(i)}$ denote the mean of individual characteristic k at the course $\mu(i) = j$ that i

is matched to.

$$\frac{1}{N_F} \sum_i H_{iF} (k_i - \bar{k}_{i,\mu(i)})^2$$

for all $k_i \in \{X_i, S_i^q, S_i^c, \bar{y}_i\}$ and $F \in \{\text{STEM}, \text{LEM}, \text{AHSS}\}$.

Leave-one-out covariance by field

We include the covariance by field between individual characteristics, skills and earnings and the leave-one-out mean of the same variables at the course that the individual is matched to. We calculate the covariance using data on **enrolees only**. Let $\bar{k}_{i,\mu(i)}^{-1}$ denote the leave-one-out mean of individual characteristic k at the course $\mu(i)$ that i is matched to.

$$\frac{1}{N_F} \sum_i H_{iF} (k_i - \bar{k})(\bar{k}_{i,\mu(i)}^{-1} - \bar{k})$$

for all $k_i \in \{X_i, S_i^q, S_i^c, \bar{y}_i\}$ and $F \in \{\text{STEM}, \text{LEM}, \text{AHSS}\}$. Above, I have used the fact that the mean of the leave-one-out mean is equal to the sample mean \bar{k} .

Distance regression

We include an auxiliary regression using the sample of **all university enrolees** relating the distance between each students school and university course to their individual characteristics and skills.

$$d_{i,\mu(i)} = \beta_0 + \beta_1 X_i + \beta_2 S_i^q + \beta_3 S_i^c$$

F Student number caps

Higher Education Funding Council for England (HEFCE) provides information on the student number controls in place for each English university from 2010 to 2014. These controls applied to new full-time UK and EU undergraduate students and post-graduate students undertaking teacher training. Universities that exceeded the cap had their teaching grants reduced in an effort to prevent over-recruitment and control government expenditure on the HE sector.

A number of exemptions from the student number controls were in place at the beginning of the period studied. For instance, students who recently completed a full-time foundation (two-

year) degree were exempt from the caps, as were any students who were studying for a qualification of an equivalent or lower level to one that they already possessed. Students who transferred courses within an institution were generally not counted against the student number controls, unless they were previously studying part-time. Some degrees which were necessary training for particular public sector professions, such as medicine or nursing, were covered in a separate cap system.

Additional exemptions were introduced in 2012 and 2013. Students with grades AAB at A-level, or other entry qualifications which are equivalent to or higher than such A-level grades, were removed from the cap in 2012. This exemption was extended to students with ABB grades in 2013. In both cases, all institutions had their cap reduced by the number of AAB or ABB students that they recruited in the previous year, to maintain the same overall level of tightness in the HE market. Finally, in 2013, HEFCE introduced a "flexibility range", allowing universities to exceed their cap by 3% without receiving any penalties. This effectively increased the caps by 3% for all universities.

While it is relatively straightforward to identify first-year UK and EU domiciled HE students for each university using HESA data, it is less straightforward to replicate the system of exemptions. We cannot, for instance, distinguish a student who has newly arrived in an institution from one who has transferred from another course, nor can we identify students who have previously completed a foundation course or an equivalent qualification. As such, we can only measure the number of students counted against the cap each year with some error.

Universities did not, in general, perfectly match their cap. The majority of offers are made over 6 months before students are due to arrive on campus and are conditional on attaining specific grades at A-level. Universities need to estimate the proportion of offers that are accepted, the proportion of accepted offers whose conditions are met, and the proportion of students who, after accepting the offer and meeting the conditions, actually turn up to enrol. The clearing system, which provides a secondary market after grades are realized, can help universities fill spare places if they find that fewer offers than expected have been accepted, but cannot be used to compensate for students failing to enrol.

Table A3 shows both our best estimate for the tightness of the overall market from 2010 to 2014.


This is calculated as the total number of new undergraduate enrollees divided by the total number of places available through the student control system. We have adjusted the places available to account for the exemption of high-attaining students and the introduction of the flexibility range. While there is some variation year on year, the overall student number control closely matches the total number of new enrollees each year.

Table A3: Student number caps

Year	Total new undergraduates	Total student number control	Average tightness
2010	285,225	297,972	0.957
2011	300,795	295,432	1.018
2012	258,705	269,504	0.959
2013	284,845	278,685	1.022
2014	293,390	278,890	1.052

Figure A2 plots the estimated proportion of cap filled for individual universities in 2010 relative to their rank in the Guardian University Guide in the same year. As noted above, the tightness of individual institutions is measured with significant error, and we are therefore likely to be overstating the extent to which individual institutions miss or exceed their cap. However, the figure demonstrates that, on average, universities are filling all their places and there is no visible relationship between the measured tightness of the university and its ranking.

Figure A2: University ranking versus estimated proportion of cap filled



G Model fit and parameter estimates

Figure A3: Relationship between quant and comms skills within different courses




Figure A4: Relationship between quant and comms skills within different courses





Figure A5: Earnings returns to quality by skills and subject



H 2012 reforms

Figure A6: Lifetime student loan repayments pre and post 2012 reform



Source: Taken from Belfield et al. (2017), with permission. Assumes maximum loan uptake and holds lifetime earnings fixed across the two systems.

I CF Reforms additional results

Figure A7: Share of percent-plan eligible students by course




Table A4: Additional counterfactual reforms, participation effects

	Baseline (2012 entrants)	Policy 4 alt. (Subj + ab grants)	Policy 5 alt. (Uncond. Perc. Plan)
<i>Low SES (bottom 50%) students</i>			
Overall	23.8	24.1	24.3
STEM Attend	43.1	48.2	42.9
LEM Attend	20.5	23.2	20.0
AHSS Attend	36.4	28.5	37.1
High Qual Attend	13.6	13.4	17.3
<i>High SES (top 50%) students</i>			
Overall	40.6	40.2	40.0
STEM Attend	42.4	39.4	42.6
LEM Attend	17.2	15.6	17.5
AHSS Attend	40.3	45.1	39.9
High Qual Attend	24.1	24.3	22.0
<i>High-ability, low SES students</i>			
Overall	79.9	85.3	87.7
STEM Attend	44.9	62.8	44.0
LEM Attend	17.2	27.4	16.9
AHSS Attend	37.9	9.8	39.1
High Qual Attend	28.7	27.1	35.1
<i>High-ability, high SES students</i>			
Overall	81.7	81.5	83.9
STEM Attend	43.8	41.2	43.3
LEM Attend	15.8	14.4	16.0
AHSS Attend	40.4	44.4	40.7
High Qual Attend	38.3	38.4	35.4

Note: See Table 8. Policy 4 alt. repeats Policy 4 from above (grants that are conditioned on subject and SES) but also conditions eligibility on ability (individuals must be 'high ability' (i.e. in the top 20% of the skills distribution) in order to qualify. Policy 5 alt. repeats Policy 5 above (the percent plan) but does not condition eligibility on SES.

Table A5: Additional counterfactual reforms, long run effects

	Baseline (2012 entrants)	Policy 4 alt. (Subj + ab grants)	Policy 5 alt. (Uncond. Perc. Plan)
<i>Earnings gap (%), between:</i>			
Low SES and High SES	29.3	27.7	28.1
V. Low SES and V. High SES	56.5	54.9	54.0
<i>Mobility rate (entry to top 20% of earnings distn.)</i>			
V. Low SES to top 20% (earns)	11.8	11.8	12.0
V. High SES to top 20% (earns)	28.1	27.9	27.8
Narrowing of gap (%), rel. to baseline		1.9	3.3
<i>Av. earnings change, rel. to baseline (%):</i>			
Low SES		1.5	0.8
High SES		-1.1	-0.5
<i>High-ability earnings gap (%), between:</i>			
Low SES and High SES	7.8	0.9	4.0
V. Low SES and V. High SES	21.4	13.6	15.0
<i>High-ability mobility rate (entry to top 20%)</i>			
V. Low SES to top 20% (earns)	31.7	33.7	33.8
V. High SES to top 20% (earns)	40.2	39.9	40.4
Narrowing of gap (%), rel. to baseline		26.4	22.8
<i>High-ability earnings change, rel. to baseline (%):</i>			
Low SES		5.3	6.1
High SES		-1.1	-0.6
<i>Fiscal costs (£millions, rel. to baseline)</i>			
Upfront grant cost (A)		349.5	5.8
Long run lost tax/loan receipts (B)		-49.0	-88.6
Total long run cost (=A+B)		300.4	-82.8

Note: See Table 9 for notes on what is included in the Table. See Table A4 for explanation of the policies.