

CESifo AREA CONFERENCES 2021

Economics of Education

Munich, 3–4 September 2021

Matching in the Dark? Inequalities in student to degree match

Stuart Campbell, Lindsey Macmillan, Richard Murphy, and Gill Wyness



Matching in the Dark? Inequalities in student to degree match

Stuart Campbell^a, Lindsey Macmillan^a, Richard Murphy^{b,c} and Gill Wyness^{a,b}

July 2021

Abstract:

This paper examines inequalities in the match between student and degree quality using linked administrative data from schools, universities and tax authorities. We analyse two measures of match at the university-subject level: undergraduate enrollment qualifications, and graduate earnings. We find for both that disadvantaged students match to lower quality degrees across the entire distribution of achievement, in a setting with uniform fees and a generous financial aid system. While there are negligible gender gaps in academic match, high-attaining women systematically undermatch in terms of expected earnings, driven by subject choice. These inequalities in match are largest among the most undermatched.

JEL classification: I23, I24

Keywords: higher education, educational economics, college choice, mismatch, undermatch

Acknowledgements: We thank Paul Gregg, Sandra McNally, John Friedman, Peter Bergman and Leigh Linden for helpful comments, as well as seminar participants at Austin, Columbia, CEP, Cornell, ISER, and Queen's Belfast, workshop participants at Catanzaro, IAB, and York, and conference participants at APPAM, EALE, ESPE, RES and SOLE. We also thank the editor and anonymous reviewer for highly useful comments.

Wyness, Macmillan and Campbell acknowledge Nuffield Foundation funding (172585).

Correspondence: Richard.murphy@austin.utexas.edu

^a UCL Centre for Education Policy and Equalising Opportunities

^b Centre for Economic Performance, London School of Economics and Political Science

^c University of Texas at Austin,

1. Introduction

Increasing enrolments in higher education (HE) is a preoccupation of governments around the world. As a result, much academic research has been devoted to examining policies intended to increase participation by relaxing credit constraints (Carneiro & Heckman 2002, Lochner & Monge-Naranjo 2011, Murphy et al 2019), providing better information (Hoxby & Turner 2015, McGuigan et. al 2016, Dynarski et at, 2018,) or improving prior academic achievement (Avery, 2013, Chowdry et al, 2013). However, less attention has been given to the types of universities and degrees students enrol in once they decide to continue with their education.

This is an issue of critical importance given the recent evidence showing the gains from students being well-matched to their degrees, with students who over or undermatch underperforming (Arcidiacono and Lovenheim, 2016; Dillon and Smith, 2019 respectively). The existence of these complementarities between degree and student quality mean that matching students to degrees has large potential impacts on the aggregate returns to higher education for society.¹ A market in which there is mismatch would imply that there are inefficiencies.²

How efficient is the matching market in the higher education sector? And how equitable is the market - are some students systematically mismatching? This paper takes a step forward in answering these questions by using administrative data from all state schools, universities and the tax authority in England. From this we track the entire cohort of 140,000 students from school to university to construct measures of student to university-subject (henceforth described as “degree”) match. Using these measures, we document the extent of mismatch and the types of students that are systematically mismatching.

In order to measure mismatch we put forward a conceptual framework for how to define it. In a perfectly functioning market, there is no mismatch and each student will be matched to the course that will maximize their lifetime returns. Any market failures, such as credit constraints, search costs, and imperfect information will generate mismatch. The match of a student to a course will depend on their characteristics, their complementarities and student

¹ Even in the absence of penalties to overmatching, the existence of capacity constraints and undermatch penalties will deliver the same conclusion. Explicitly, if some students are overmatching, given the capacity constraints, this will result in some students undermatching.

² If, on the other hand, we assume there are no complementarities between students and courses, a well-functioning market would be one in which students attend the course with the highest return, given their characteristics (including attainment). Here, mismatch again implies the existence of inefficiencies, such as frictions in the application and enrolment process, or that students are badly informed about course quality. However, student preferences could also create the appearance of mismatch, which would be problematic since preferences cannot typically be ruled out.

preferences. Untangling unobservable underlying preferences from other market failures, such as imperfect information is a challenge. For example, high ability low-income female students may be qualified to attend high ranking (and high earning) STEM courses, but may choose to enrol in a lower ranking humanities course instead. This could be due to lack of perfect information about the costs and benefits, or because they actively prefer to attend the less selective university.

In practice, both unobservable preferences and market failures are responsible for student choices. Therefore, understanding whether a student is truly mismatched requires either i) perfect knowledge about their preferences, or ii) assumptions about the circumstances in which we believe someone is mismatched. In order to make progress in defining match (and mismatch) in the absence of information about underlying preferences, we take the stance of the social planner and assume that it is optimal for society that the highest quality students should attend the highest quality courses (Arcidiacono and Lovenheim, 2016; Dillon and Smith, 2019) and that students have no subject preferences. Any student not attending their highest quality course possible is therefore mismatched. In reality as some of the differences will be due to preferences, we must regard our estimates of mismatch as an upper bound. We explore the role of preferences and market failures as determinants of mismatch in the second half of the paper.

We create two measures of match. For both, we rank students nationally based on their end of secondary school qualifications. We also rank degrees nationally, first, according to the qualifications of the median student on each degree, and second, according to the median earnings of previous graduates on the degree. We create our measures of match by taking the difference between the percentile ranking of the student and the degree. The advantage of our approach is that it provides a transparent and continuous measure of match, which can be used in many settings. Moreover, defining match to degree based on potential earnings is a new addition to the literature, which allows us to shed light on previously undocumented large disparities in match.

We use these two measures to document socio-economic status (SES) and gender differences in match, taking three distinct approaches. First, we plot student qualification percentile against degree quality percentile for students by SES and gender. Plotting the raw data in this way imposes no functional form assumptions on the data, and presents the extent of match throughout the achievement distribution. Second, we estimate the average SES/gender match gaps conditional on individual characteristics and prior achievement across the entire distribution of achievement – an important shift in the literature, which has typically focused

on high achievers. Finally, we implement unconditional quantile regressions (UQR) across the distribution of match to determine whether these mean effects are masking larger SES and gender gradients for those who are very mismatched. The combination of these approaches allows us to explore hidden non-linearities across the entire academic achievement distribution, and reveals several important findings.

We find inefficiencies in the higher education market in the UK in the form of under and overmatch. We also find sizeable socio-economic inequalities in academic and earnings match across the achievement distribution, with low SES students consistently undermatching, attending degrees with lower attaining peers and lower potential earnings than their richer counterparts. These inequities in the market remain after conditioning on a set of individual demographics and a complete history of prior test scores. In the top quintile of the achievement distribution, disadvantaged students are 8 percentiles lower matched than their more advantaged counterparts. This corresponds to the difference between studying economics at the London School of Economics (ranked 5th in the Times Higher UK university rankings) versus Exeter (ranked 18th). The largest inequalities are not found at the top of the achievement distribution, but around the 90th percentile. These disadvantaged students are 9 to 11 percentiles lower matched than their more advantaged counterparts. We find little evidence that these gaps are driven by the subjects that people study at university. Even when they have similar prior achievement, and are studying similar degree subjects, low SES students study at lower ranked (in terms of both achievement and earnings) institutions. This is consistent with recent research from Chetty et al. (2020) which documents the ‘missing middle’ – middle class students with high test scores who are under-represented at the most selective US colleges.

The existence of mismatch in the system implies that there are inefficiencies at play. Research to date has highlighted the role of credit constraints, geographical isolation and information inequalities as the key drivers of mismatch, particularly among low SES students (Dillon and Smith, 2018; Hoxby and Avery, 2012). We can rule out the first two of these as explanations of mismatch in the system we study. Credit constraints play a minimal role in the UK; practically all university degrees charge the maximum tuition fees allowable – meaning there is no variation in fees across degrees, so poorer students cannot make a price-quality trade-off – and all students have access to income-contingent loans that cover the entirety of the tuition fees plus loans for living expenses³. This may explain why we observe less mismatch

³ Combined tuition fee and maintenance loan take-up was 86% in 2010/11 (when eligible amounts were similar to the cohort we study) (Bolton, 2019)

in the UK than in the US, where there is large variation in fees and such loans are not widely used⁴. We find that geography has little impact on the SES match gap. In our context, students only have to travel short distances to find a well-matched course, with the average distance to a well-matched course in England a little under 9 miles.⁵ On average low SES students attend colleges closer to home, but conditioning on distance to either university attended or a well-matched degree does not impact the match parameters significantly. That low SES students are systematically undermatching despite the lack of geographic or financial constraints, means that there are other structural or social factors at work that the research to date has not fully explored.

We show that a major determinant of SES inequalities in mismatch is high school attended. The SES match gap for students from the same school is reduced by up to 79 percent (down to 2 percentiles) with the inclusion of school fixed effects. This implies that factors correlated with high school such as peers, school resources, information inequalities including careers advice and guidance, and parental sorting play an important role in student match.

In addition to SES match, we provide facts about the previously undocumented gender inequalities in student to degree match. In contrast to the large SES gaps in academic match, we find only modest differences in academic match between males and females. Meaning that males and females with a given set of qualifications are enrolling in courses with similar entry standards. However, by stark contrast, we find sizeable gender gaps in earnings match. After accounting for prior test scores and demographics, high-attaining women attend degrees around 8 percentiles lower in potential earnings than men - this gap is the equivalent of £25,800 per year for those degrees at the top of the median earnings distribution. This highlights that women are attending degrees that are as academically competitive as their male peers, but that have substantially lower average earnings. We find that almost the entire of the gender gap in earnings can be accounted for by degree subject choice (rather than university attended), with women more likely to attend degrees such as Creative Arts and English – which are academically selective, but have typically lower earnings.

Our paper makes several key contributions to the emerging academic literature on the match between student achievement and college quality. Existing papers on mismatch, and the HE inequality literature in general have typically focused on high-achieving low-income students, using a binary measure of undermatch (Hoxby and Avery, 2012; Black, Cortes and

⁴ See Barr et al (2017) for a discussion.

⁵ The UK has 3.47 universities ranked in the QS top 1000 ranking, per 10,000 km² compared to 0.17 in the US.

Lincove, 2015) or examined mismatch at different points in the distribution (Dillon and Smith, 2017). In contrast, we create continuous measures of mismatch, and present estimates across the distribution of achievement. Our focus on the entire distribution of achievement is a non-trivial shift in this literature. Simply focusing on high achievers obscures more dramatic undermatch among those in the 70-90th percentiles of the skills distribution. Our results imply that efforts to reduce mismatch should be expanded beyond just the highest attainers.

The continuous nature of our measures in conjunction with our large dataset makes it possible to make a novel contribution to the literature by examining the nature of mismatch at its extremes through unconditional quartile regression (UQR). Standard OLS estimates of mismatch undersell the importance of SES and gender gaps for the most under-matched students. We estimate OLS conditional SES earnings gap of high achieving students to be 8 percentiles, but for the most under-matched the SES gaps they are as large as 26 percentiles. For gender, the earnings match-gap among the most qualified students is 16 percentiles (compared to 8 percentiles at the mean). Our focus on SES and gender gaps as opposed to race also sets us apart from the rest of the literature.

We are also the first to study mismatch on the basis of degree earnings potential. Previous studies have measured university quality based on entry qualifications of the students at that institution (Hoxby and Avery, 2012; Light and Strayer, 2000) or a composite of institution quality measures (Dillon and Smith, 2017; Smith, Pender and Howell, 2013). Measuring university quality on the basis of graduate earnings is important for understanding the role of match in intergenerational mobility. Our finding that talented low SES students are enrolling in degrees with lower returns undermines the potential for higher education to have a positive impact on social mobility.

A final contribution is that we can study mismatch at university subject (degree) level. All existing studies of mismatch have been unable to untangle the role of university subject/major as a factor in match. Moreover, we provide a framework for researchers to consider the role and identification of preferences in a system with market failures, which would otherwise be conflated. This, in conjunction with our new measure of earnings match, allows us to highlight large and undocumented gender match inequalities. Our finding that talented women enrol in subjects which command lower returns than equally talented men is relevant for the much-documented gender pay gap. It also implies that interventions aimed at improving match should be targeted at women as well as disadvantaged students.

The remainder of this paper proceeds as follows. Section 2 describes our institutional setting, the dataset, and the methods we use to create our measures of mismatch. Section 3

presents our results from the three approaches, while Section 4 presents robustness tests. Section 5 explores potential drivers of undermatch, and Section 6 concludes.

2. Data and Methods

2.1 Institutional setting

We analyse inequalities of match in the UK context, which provides some perspective on the findings from the predominantly American literature. While other studies of mismatch have pointed to the role of finance as a potential driver (Hoxby and Avery, 2012), UK students face far fewer financial barriers. There are no upfront costs in the UK system - all college fees and living costs are covered by income-contingent loans which are repaid upon graduation once the graduate is earning over a certain level (Murphy, Scott-Clayton and Wyness, 2019). Moreover, there is little price variation between institutions or subjects meaning that students do not face a trade-off between quality and price which may cause them to mismatch. A final feature of the UK system is that it has a centralised applications system (the University and College Applications Service, or UCAS) which is easy to access and navigate and is used by the vast majority of university applicants. Students are provided with standardized information on all the degrees including typical grade requirements, and can apply for up to five degrees paying a single application fee of £24. Thus, the finding of substantial student to university mismatch even in a system with few financial barriers, relatively low costs, and streamlined application system is important, pointing to other possible reasons for this mismatch.

As in the US, students are still likely to face information constraints, however. The structure of the UK education system means that students make a number of crucial choices about their education path as early as age 13/14. At this age, students choose the types of qualifications and, crucially, subjects that they will study in their final two years of compulsory schooling, most often for 10 subject-specific General Certificates of Secondary Education (GCSEs). Those who stay on after the compulsory schooling age face another set of important decisions regarding their qualification and subject choices from age 16 to 18, most commonly in the form of 3 subject-specific Advanced Level qualifications (A-Levels). Finally, again unlike the US, students wanting to study for Bachelor's degrees then have to choose both an institution and subject (degree) at application stage. Such early subject specialisation, which begins at age 13/14, may be conducive to mismatch.

2.2 Data

We use individual-level administrative data on the population of state-school students in England for a single cohort. Our focus is on the cohort of young people who took their compulsory age 16 exams in 2006 and their non-compulsory exams two years later in 2008, at the end of secondary school. The grades from these exams are used to determine which university a student will be admitted to. The students enter university in the autumn of that year at age 18 (the traditional age for university entry in England) or 19 if they took a gap-year (around 25% of our sample - see Table 1). Our data cover students in all publicly funded English schools,⁶ and we combine this with information on the university degree attended by these students anywhere within the UK (England, Scotland, Wales, and Northern Ireland).

Our schools data come from the National Pupil Database (NPD), and include basic demographic information (gender, ethnicity, English as an additional language, special educational needs) alongside externally set and marked exam results at ages 11, 16, and 18. There is substantial attrition over this period of education in the English system, since many pupils leave at the end of compulsory education, after exams at age 16, and a smaller group leave at age 18 without going on to university. Our main interest is in the subgroup who go on to university, but we use information on the complete population of age 16 students to construct our measure of SES, as we describe in Section 2.3 below. Starting with a population of around 590,000 state-educated pupils in the 2006 cohort, we initially restrict the sample to all university students who went to a state-school, and on whom we have information on exam results at age 18, which results in a final sample of 138,969.

Our linked data on degree attended⁷ come from the Higher Education Statistics Agency (HESA). We use university entry information from 2008 and 2009, since a quarter of students in England delay university entry for one year after age 18 examinations. These data contain information on every student's degree in every higher education establishment in the UK. Our main estimates use a 23 subject classification to distinguish degrees within universities for both achievement- and earnings-based match. This classification distinguishes “Medicine & Dentistry” from “Nursing”, and “Economics” is separately classified from other Social Science

⁶ 93 percent of students attend publicly funded secondary schools in England (Table 2A, DfE, 2010)

⁷ Note that as in Dillon and Smith (2017) we observe a collapsed version of the student-degree match process, in that we only observe the degree that they attend, rather than where they apply.

disciplines.⁸ We also have access to a more detailed, 631 subject classification for academic-based match, which we use in robustness checks below.

Finally, we incorporate aggregated data on the earnings of outcomes of an earlier university cohort, which are based on tax records. These data come from the Longitudinal Education Outcomes (LEO) dataset, which is compiled from tax records by Her Majesties Revenue and Customs (HMRC) in the UK. The LEO dataset records the earnings of all students who were UK-domiciled on entry to university, and were in full or part-time work as employees in the UK five years after graduation.⁹

We use the median earnings outcomes 5 years after graduation (age 26) for the earliest available cohort, which is those who completed undergraduate degrees in 2009, the academic year before our cohort entered university.¹⁰ These data are available for all 23 subject categories at each university where a subject is offered. A full description of the data retrieval process, methodology, and data quality is provided in Department for Education (2017).

2.3 Measuring socio-economic status

To construct a measure of students' socio-economic status (SES) we follow Chowdry et al (2013). To obtain an indication of student's SES within the wider population, rather than limiting to the university going student population we use information on all state school students from when they were aged 16. This includes whether a student was eligible for free school meals at age 16 (around 15 percent of students), alongside a set of variables which describe the neighbourhood in which they live at that age. The neighbourhood characteristics are taken from the 2001 Census and are aggregated to the Lower Super Output Area level, which is a neighbourhood containing around 700 households or around 1,500 individuals. This provides information on the proportion of individuals in the neighbourhood that: 1) work in

⁸ The 23 subjects are: "Agriculture & Related Subjects", "Architecture, Building & Planning", "Biological Sciences (excluding Psychology)", "Business & Administrative Studies", "Combined", "Computer Science", "Creative Arts & Design", "Economics", "Education", "Engineering & Technology", "English Studies", "Historical & Philosophical Studies", "Languages (excluding English Studies)", "Law", "Mass Communications & Documentation", "Mathematical Sciences", "Medicine & Dentistry", "Nursing", "Physical Sciences", "Psychology", "Social Studies (excluding Economics)", "Subjects Allied to Medicine (excluding Nursing)", and "Veterinary Science".

⁹ Self-employed workers are not included, along with employees below the 'Lower Earnings Limit' which stood at £112 per week for this cohort. Self-employment is no higher than 10% in any subject group for the cohort in the data, and the proportion of UK domiciled graduates not matched to the HMRC dataset for any reason (including very low earnings, self-employment, emigration, or death) is less than 5% for this cohort (Department for Education, 2016).

¹⁰ A comparison of the ranking of earnings by institution and subject when individuals are age 29 illustrates that the ranking of courses are broadly stable at later ages (Belfield et al., 2018).

managerial or professional occupations; 2) hold an A-Level equivalent qualification or above; and 3) own their home. In addition, we also use the 2007 Index of Multiple Deprivation.

We combine these measures using principle components analysis to create a standardised index.¹¹ We use the whole population of state-school students at age 16 in the relevant cohort to construct the index, so throughout this paper “SES” refers to socio-economic position relative to the whole school-cohort population rather than relative to the university-attending sub-population. The final row of Table 1 illustrates that this results in 9 percent of our university-attending sample coming from the most disadvantaged families, and 34 percent from the least disadvantaged families.

Table 1 highlights the key characteristics of our sample by SES quintile. Women are overrepresented in higher education, making up 56 percent of the sample. A quarter of our sample took a gap year, with the least deprived families more likely to take a year out than the most deprived. There are only a small proportion of people with special educational needs in our sample as might be expected, with 6 percent of the most deprived families and 3 percent of the least deprived families being categorized in this way. Finally, there is a strong association between having English as an additional language, ethnic minority status, and low SES, with these groups accounting for a larger proportion of low SES families.

Table 1: Summary Statistics

| Student Characteristics | Quintile of SES | | | | | Gender | | Total |
|-----------------------------------|-----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|
| | 1 | 2 | 3 | 4 | 5 | Men | Women | |
| Ethnic minority | 0.41 (0.49) | 0.30 (0.46) | 0.18 (0.38) | 0.12 (0.33) | 0.10 (0.30) | 0.17 (0.17) | 0.18 (0.18) | 0.17 (0.38) |
| English as an Additional Language | 0.30 (0.46) | 0.21 (0.40) | 0.12 (0.32) | 0.08 (0.27) | 0.06 (0.24) | 0.12 (0.12) | 0.12 (0.12) | 0.12 (0.32) |
| Special Educational Needs | 0.06 (0.23) | 0.05 (0.21) | 0.04 (0.19) | 0.03 (0.18) | 0.03 (0.18) | 0.05 (0.05) | 0.03 (0.03) | 0.04 (0.19) |
| Gap year | 0.22 (0.41) | 0.23 (0.42) | 0.24 (0.43) | 0.25 (0.44) | 0.28 (0.45) | 0.26 (0.26) | 0.25 (0.25) | 0.25 (0.43) |
| A*-C in EBACCs | 0.23 (0.42) | 0.32 (0.47) | 0.40 (0.49) | 0.46 (0.50) | 0.54 (0.50) | 0.42 (0.42) | 0.45 (0.45) | 0.44 (0.50) |
| Sample composition | | | | | | | | |
| Proportion of sample | 0.09 | 0.14 | 0.19 | 0.24 | 0.34 | 0.44 | 0.56 | 1.00 |
| N | 12025 | 19074 | 25856 | 33859 | 47694 | 61348 | 77621 | 138969 |

Source: NPD-HESA. n=138,969. Notes: A*-C in EBACCs measures the percentage of students who achieve five or more grades A* to C in traditional academic GCSE subjects (English, Maths, Science, Geography or History, and a language). Quintile of SES is defined out of the entire age 16 student population.

¹¹ See Appendix Figure A1 for a comparison of this measure to an alternative measure of parental socio-economic status from a linked data source. Results are comparable when using this alternative measure to capture socio-economic status, or the free school meals indicator alone, or a measure of parental education.

2.4 Two measures of student-degree match

We are interested in the match between student quality and degree quality. We calculate student quality according to age 18 exam test scores. Note our intention is to measure student qualifications, rather than measure student ability, as student qualifications are the principal metric of which course the student can access.¹² We have two measures of degree quality, one based on the achievement of students on each degree, and one based on graduate earnings of previous cohorts of students on the degree, giving rise to two measures of student-degree match.

Each measure is calculated in three steps:

- (1) Calculate student quality: we rank individuals in the distribution of age 18 exam test scores based on their performance in their best three exams.¹³
- (2) Calculate degree quality: we rank each university-degree combination in a distribution of degree quality, based on either
 - (i) The median of the best three age 18 exam results of students on the degree (academic-based), or
 - (ii) The median earnings outcomes of an earlier cohort of students on the subject 5 years after graduation (earnings-based).

With regards to the calculation of student quality, as mentioned in Section 2.1, a distinctive feature of the UK education system is the importance of subject choices made in secondary education and at university. Our measure of individual quality is based on the best three exam results. A-levels are graded on a scale of A/B/C/D/E which are worth 270/240/210/180/150 QCA (Qualifications and Curriculum Authority) points respectively. Students typically study three A-levels in different subjects, and the majority of universities set their entry requirements according to this measure. However a further complication is that some subjects are considered by universities to be more rigorous than others. This can be explicit, for example by naming ‘facilitating’ or ‘preferred’ subjects, and other times implicit in the offers that universities make to potential students (Dilnot, 2018).

To account for these differences in universities’ subject preferences, we follow Kelly (1976) and Coe et al. (2008) in calculating a subject difficulty adjustment, using an iterative approach based on our samples’ performance in different combinations of age 18 exams. For

¹² In our conditional estimates of the match gap we control for student ability using the complete set of test scores from compulsory national examinations at age 11 and 16.

¹³We consider only the students who go on to university, so the relevant exam results distribution is that of university attendees. Some students take degrees that are equivalents to A-Levels. In these cases we calculate their A-Level equivalence scores.

example, if students who took the same set of subjects consistently scored higher in one of these subjects, that subject would be deemed easier and would be awarded less points. This is iterated over all students and subject groupings until the difficulty adjusted scores are equalised. This is explained in more detail in the Technical Appendix. Figure A2 illustrates the difficulty ratings calculated for each subject, with the most difficult subjects being mathematics and natural sciences. We use these difficulty-adjusted points when assigning students to percentiles of the achievement distribution¹⁴.

With regards to constructing our course quality metrics, it is important to account for courses being of different sizes. To address this, we assign the relevant course-level median student value to all students on that course. This is difficulty adjusted A-Level achievement, or graduate earnings five years after graduation for academic and earnings quality respectively. Respective course quality percentiles are then calculated using these adjusted student measures. This effectively weights the courses by enrolment as we are allocating course quality percentiles using the total number of seats available.¹⁵ Assigning the course percentile at the individual level ensures that there will be an equal degree of over and undermatch in the whole population.^{16,17}

As a final step we:

(3) Calculate match: We subtract the student's percentile in the exam results distribution from the percentile of their degree on the quality distribution.

We therefore have two continuous measures of match for each student, an academic-based measure and an earnings-based measure. The continuous nature of our outcomes allow us to analyse inequalities across the severity of mismatch, rather than relying on arbitrary thresholds to categorize students as matched or not. Both measures represent the distance of each student's chosen degree from their position in the achievement distribution. With both measures of

¹⁴ A parallel set of results using the un-diffculty adjusted rankings are also available in Appendix Table A2. All results are qualitatively similar.

¹⁵ Note that there is no variation in percentiles within a course, as any ties are assigned to the lower percentile. In practice, the attributions of ties makes no difference given that we have 1722 courses (university-subject combinations) and every course is less than one percent of the total seats available.

¹⁶ The first two rows of Appendix Table A2 present the total amount of over and undermatch for the academic and earnings match measures for the entire sample population. The table shows that total market-wide overmatch is virtually equal to total undermatch.

¹⁷ An alternative way to do this would be to use only the median student's value, and then percentalise over all courses at the course level. This would be problematic since it would not account for the number of students on the course. Failing to account for this would render our student rankings incomparable with our course rankings. Note that the choice of method does not impact our results materially (Results available upon request). The reason for this is that in practice each course has a small share of students (<1%) and with small amounts of variation in course sizes.

match, a student at the 50th percentile of the A-Level distribution would be considered matched if they are enrolled on a degree at the median of the quality distribution. If a student attends a degree at a lower percentile than their own percentile in the student quality distribution, we consider them undermatched. If they attend a degree which ranks above their position in the student quality distribution, we consider them overmatched.

The academic-based measure of match measures whether students are enrolling in the degrees of the level of academic prestige that one might expect, given their qualifications. The earnings-based measure of match measures whether students are enrolling in degrees with the potential earnings that we might expect, given their qualifications. The latter is in the spirit of a classical human capital assumption, namely that students should expect earnings outcomes which are broadly comparable with their place in the achievement distribution. But it also has implications for social mobility, if low SES students are found to choose degrees with lower potential earnings.

We use raw graduate earnings for each course as a representative measure of course quality and complementarity with high quality students. Of course, these raw rankings of earnings may differ significantly from the rankings of causal value-added on earnings in higher education (Dale & Krueger, 2002; Kirkeboen et al. 2016; Mountjoy & Hickman, 2019; Hoxby & Stange, 2020). We argue, however, that for our purposes, ranking on the basis of raw graduate earnings is a preferable measure of course quality, and that ranking by value-added could actually be problematic.

An important underlying assumption of the mismatch literature is that there are complementarities between students and courses/universities. Some students will gain more from some courses than others will. High value-added courses generate a high return to the students that attend them, but it is unlikely that such courses would offer high value added for all students. Mountjoy and Hickman (2019) find high earnings courses tend to be attended by highly able students, whereas there is little relationship between a courses' value-added and its mean incoming SAT score. Indeed, high value-added courses are slightly more likely to be attended by students with lower SAT scores on entry (Department of Education 2015, Mountjoy and Hickman 2019). However, causal estimates of course value-added from natural experiments will be based on the marginal students that attend those courses. This means it is difficult to obtain value-added estimates of courses for students who would not attend those courses. Using value-added estimates based on marginal students has the potential to be unrepresentative of non-marginal students, and this is particularly important when considering mismatch across the distribution and could potentially lead to misleading results. For example,

would the courses estimated to have highest value-added, which predominantly serve low income and low prior achievement students, have just as high value-added for high-ability students who attend?

In addition to obtaining an appropriate measure of course quality for each student, in order to rank courses in a way in which we can make comparisons across students, we require a single index. The assumption that course rankings are the same for every student is possible for raw graduate earnings, but not in the case of value-added measures since as described above, the relationship between a course' value-added and will be dependent on the student. Courses could be consistently ranked in terms of value added if one was to assume there were homogeneous returns to courses across students. Such assumption however, would mean that there would be no gains from students matching to courses.

We use the earnings-potential measure of course quality as it is presentable in a single index, and we make the assumption that students with high levels of pre-university human capital would gain the most from attending courses with the other students who have high human capital as measured by their future earnings. Of course, graduate earnings will be a mixture of student inputs and value added.

Both the academic and earnings-based measures reflect different aspects of degree quality, and the same degree can be at quite different relative positions. For example, a degree which is positioned near the bottom of the achievement distribution, Computer Science at Southampton Solent, is considered high quality in terms of earnings, ranked at the 70th percentile. In contrast, English at Edinburgh is ranked at the 90th percentile on our academic-based quality measure, but is only ranked at the 35th percentile in terms of our earnings-based measure. The correlation between the two measures of 0.58. Appendix Figure 3 presents a scatterplot of each course by these degree quality measures.

2.5 Methods

To understand the nature of student matching we use three distinct methods to present the results. First, we show a simple plot of students' achievement decile against average degree quality for all students in that decile. If all students were perfectly matched to their degrees this line would be straight and at a 45-degree angle.¹⁸ The extent to which a point is above a 45-

¹⁸ This is feasible, even if at the extremes of the quality distribution it would require there to be no variation in student quality within a degree. For example the courses at the top percentile of the quality distribution would need all their students to be from the top percentile. There is evidence of near perfect matching in our data, for the top ranked degree of Maths at Cambridge more than half of the 92 students enrolled have 'best 3' A-grades in either Physics, Maths, Further Maths, or in Chemistry, Physics, Further Maths.

degree line indicates how overmatched these students are on average, and similarly the distance below the 45-degree line reveals the extent of undermatch. It imposes minimal assumptions beyond those involved in the creation of the metrics. Plotting this match-line for different types of students allows us to study inequalities in match at any point in the achievement distribution.

Second, we estimate SES and gender gaps in match, conditional on individual characteristics and achievement prior to age 18. Specifically, we estimate the following regressions:

$$M_{ia} = \beta_0 + \sum_{j=2}^5 \beta^j I(SES_i = j) + \gamma female_{ia} + \delta X_{ia} + \pi P_{ia} + \varepsilon_{ia}, \Delta a \quad (1)$$

Where M_i is our measure of match, $\widehat{\beta}^j$ represents our estimated SES gap in match, and $\hat{\gamma}$ is our estimated gender gap in match, conditional on background characteristics (X_i) and prior achievement at age 11 and 16 (P_i). Given that achievement is used to define match, there will be ceiling and floor effects; it would be impossible for the lowest ranked students to undermatch or the highest rank student to overmatch. We therefore estimate the models separately, first across deciles of achievement (a), before focusing on those in the top and bottom quintiles of achievement. Standard errors are clustered at the secondary school level.

While this approach estimates SES and gender gaps at the mean of our match outcomes, it is unlikely that all individuals of a given characteristic mismatch to the same extent. In fact one may be most interested in the students that under or overmatch the most, and how large the gaps are for these students. A key strength of our continuous measure of match is that it allows us to consider this for the first time. To this end, our third approach is to use unconditional quantile regression (UQR) to estimate SES and gender match across the distribution of match.

We use UQR instead of the more traditional Conditional Quantile Regression (CQR), because we do not want to alter the distribution of our match measure due to the inclusion of our covariates. Take the simple example of examining the gender gap in match using CQR. Using this approach would involve estimating the gender gap for men and women within their own (gender) conditional match distribution. If we consider females in the distribution of only females (the CQR), a relatively high attaining female attending a higher earning course would be well matched. However, if we consider the same female in the distribution of males and females (the UQR) that same female's position would change, since her position would be relative to both males and females. This is important because this changes the interpretation of the quantile regression estimates. For CQR, we would be estimating gender (or SES) gaps for those at that point in the conditional match distribution, whereas for UQR we are comparing gender (or SES) gaps at points in the unconditional distribution of match, regardless of the

covariates included. We are thus able to illustrate gaps in mismatch across the underlying match distribution.

One of the common ways to implement a UQR model is to use a Re-centred Influence Function (RIF) regression (Firpo et al., 2009), specifying our distributional statistic of interest as the quantiles of our match variable q_τ where τ is each decile from 1 to 9.

$$RIF(M_{ia}; q_\tau) = \beta_0^\tau + \sum_{j=2}^5 \beta^j I(SE S_i = j) + \gamma^\tau female_{ia} + \delta^\tau X_{ia} + \pi^\tau P_{ia} + \varepsilon_{ia}, \Delta a \quad (2)$$

Here, our coefficients $\widehat{\beta}^{\tau}$ and $\widehat{\gamma}^{\tau}$ illustrate the estimated SES and gender inequalities in match at the given decile, τ . Given that we estimate our models by achievement quintiles (a), for high-attainers this will estimate the SES and gender gaps from the most severely undermatched (10th percentile) to those who are matched (90th percentile). For low-attainers, this will estimate the SES and gender gaps for those who are matched (10th percentile) up to the most severely overmatched (90th percentile).

In the penultimate section of the paper, we explore the role of market failures and student preferences as potential drivers of SES and gender gaps in match. We consider a range of different mechanisms by including a series of additional covariates Y_{ia} in model (1) representing different sources of potential market failure and preferences. Here we are interested in how much these reduce our estimated SES and gender gaps.

$$M_{ia} = \beta_0 + \sum_{j=2}^5 \beta^j I(SE S_i = j) + \gamma female_{ia} + \vartheta Y_{ia} + \delta X_{ia} + \pi P_{ia} + \varepsilon_{ia}, \Delta a \quad (3)$$

To explore the role of transport costs, we estimate three alternative specifications for geography ($Y_{ia} = \varphi dist_{ia}$): one controlling for distance to university attended in kilometres, one controlling for distance to each of the nearest three universities to the student's home neighbourhood, along with the distance to all remaining universities (similar to Gibbons and Vignoles, 2012), and one controlling for distance to the nearest university at which individual i would be matched to any course.

To explore the role of school-level factors, and their relation to (mis)information, in driving SES and gender inequalities in match, we control for school fixed effects ($Y_{ia} = \omega school_{ias}$). We further investigate the role of possible school mechanisms by replacing these fixed effects with include a range of observable school factors, including the average achievement of the school, the provision and uptake of 'facilitating' or 'preferred' subjects, the proportion of students that attend university and high-status universities, and the proportion of

students who are high SES. The final set of analyses determines the extent to which student preferences are determining the match gaps. The role of subject preferences has the potential to be a large determinant of mismatch because, in the absence of direct information on student preferences, we have made the assumption that students have no subject preferences and consider the full range of subject and institution choices before them. In reality, subject preferences are likely to play a role and so some of the students who are defined to be mismatched may not be, from the student's perspective.

Our first approach to disentangle subject choice from our match measure is to include 23 subject categories in model (3) (where $Y_{ia} = \sum_{k=2}^{23} \sigma^k I(\text{Subject}_i = k)$). This accounts for the average mismatch of students studying a certain subject area e.g. students studying history being undermatched in terms of earnings. Any reduction in the gender or SES parameters represents the correlation between that characteristic and sorting into the subjects taken. Any gap remaining in this specification can be interpreted as likely institutional-driven inequalities, within subject of study.

Our second approach assumes students have very strong subject preferences such that they only consider degrees of the same subject as that which they went on to study. To do so we create a within-subject match measure, using the difference between adapted measures of student and course quality (this is described in full detail in Section 5.3). The new quality measures define student (course) quality on the basis of their position within the student (course) quality distribution for that subject. We repeat the non-parametric plots and specification (1) with this new match metric. Comparing the gap estimates assuming extreme subject attachment to those with no subject attachment reveals the extent of subject preference in these match gaps.


3. Results

Figure 1 shows the distribution of the two measures of student-degree match which result from this process. Both measures have peaks with students being well matched and are approximately symmetrical. The earnings-based measure is more dispersed than the academic-based measure. This reflects that there are observed academic-based entry requirements for enrolling on a degree. There are no such restrictions in terms of later earnings, and students are likely to be less well informed of the potential earnings of each degree.

Using the binary definition of mismatch from Dillon and Smith (2017) where mismatch is +/- 20 percentiles from the perfectly matched degree, 16 percent of our sample are

overmatched and 16 percent of our sample are undermatched using our academic-based measure. For our earnings-based measure 22 percent overmatch and 23 percent undermatch. Dillon and Smith (2017) find around 25% of students in the US are overmatched and 25% undermatched according to their composite college-input-quality measure¹⁹. This is most comparable to our academic-based measure of match, and while it would be problematic to draw strong conclusions from this, the comparison is suggestive that there is more mismatch in the US than in the UK.

Figure 1: Academic-based and earnings-based measures of student-degree match



Source: NPD-HESA, HMRC. n=138,969. Notes: Academic-based match defined by degrees' median student age 18 achievement percentile minus student's age 18 achievement percentile. Earnings-based match defined by degrees' median graduate earnings percentile minus student's age 18 achievement percentile.

Figure 1 further highlights the strength of our approach in being able to analyse the extent of inequalities in mismatch in the tails of the distribution. 3% of our sample are undermatched by over 50 percentiles using our academic-based measure, and 5% are undermatched by over 50 percentiles using our earnings-based measure. 1% overmatch by more than 50 percentiles using our academic-based measure, and 5% overmatch by more than

¹⁹ Dillon and Smith's college quality measure comprises 4 measures of quality – the mean SAT score (or ACT score converted to the SAT scale) of entering students, the percent of applicants rejected, the average salary of all faculty engaged in instruction, and the undergraduate faculty-student ratio.

50 percentiles using our earnings-based measure. In section 3.3 we will explore how SES and gender gaps vary for the most severely under- and over-matched.

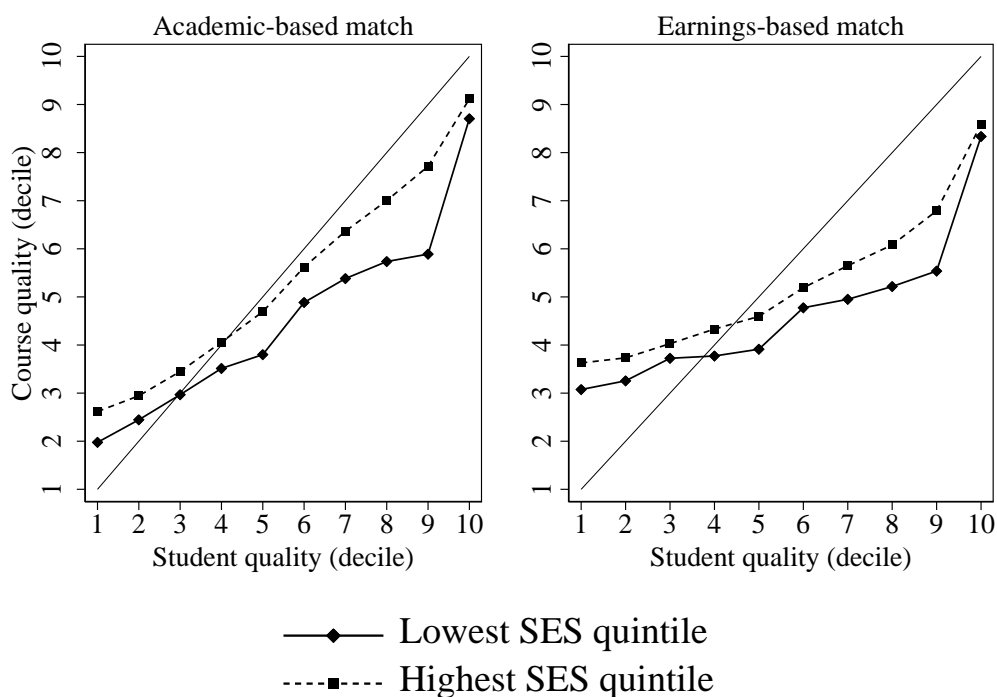
3.1 Inequalities in Match Gaps by Achievement

Figure 2 plots our two degree quality measures against student achievement for high and low SES pupils, illustrating raw gaps in points- and earnings-based match by SES across the achievement distribution (left and right panel respectively)²⁰. For both match measures we see that the relationship is approximately linear and is flatter than 45 degrees, meaning that low-attainers are more likely to overmatch and high-attainers are more likely to undermatch (as previously described, reflecting floor and ceiling effects). As would be expected given the distributions in Figure 2, we see that the earnings match curve is flatter than the points match, meaning that there is more mismatch in terms of earnings than points.

For both measures we see stark SES gaps in match. For every given percentile of individual achievement, high SES pupils attend higher ranked degrees than low SES pupils. The SES match gap increases in the top half of the distribution of student achievement, with the exception of the top decile of students where the gap is the smallest. As much of the previous literature on mismatch has focused on high-attaining low SES students (Hoxby and Avery, 2012; Black, Cortes and Lincove, 2015), this implies they may be underestimating the extent of mismatch, by failing to study those students for whom mismatch is largest, between the 70th and 90th percentiles of achievement. These patterns hold for both the points- and earnings-based match measures.

²⁰ Note for Figure 2, which plots the average match for High and Low SES students, this does not exhaust the population as this is just the highest and lowest SES quintiles and so the total areas between the lines and the 45 degree line need not cancel out.

Figure 2: SES match by student achievement




Source: NPD-HESA, HMRC n=138,969. Notes: The 45-degree line represents perfect matching throughout the achievement distribution. Student quality defined by their age 18 difficulty adjusted exam performance. Course quality is defined by two measures. Academic Quality defined by the course’s median student’s difficulty adjusted exam performance. Potential Earnings Quality defined by the course’s median student’s earnings five years after graduation. All quality measures are converted to percentiles, weighted by student enrolment, from which the deciles are obtained.

Figure 3 next plots gender gaps in match for our two degree quality measures²¹. Unlike our findings for SES, the findings differ across measures of match. For our academic-based match we observe almost no gender gap in match. In the bottom half of the achievement distribution men and women attend degrees that are equally academically selective, and in the top half men are enrolling in degrees with slightly higher peer achievement. By contrast, the earnings match measure highlights striking gender gaps. Men consistently attend degrees with graduate earnings around one decile higher than women across the distribution of achievement. This gender gap narrows in the top achievement decile, but even then, males with the same subject difficulty-adjusted achievement are still enrolling in degrees with higher median earnings. In the next section, we test the robustness of these gaps at the top and the bottom of the achievement distribution, by conditioning on characteristics and prior academic achievement.

²¹ As the categories in Figure 3 incorporate the entire population, the total amount of under and overmatch is equal. This is shown in Appendix Table A1.

Figure 3: Gender match by student achievement




Source: NPD-HESA, HMRC n=138,969. Notes: The 45-degree line represents perfect matching throughout the achievement distribution. Student quality defined by their age 18 difficulty adjusted exam performance. Course quality is defined by two measures. Academic Quality defined by the course's median student's difficulty adjusted exam performance. Potential Earnings Quality defined by the course's median student's earnings five years after graduation. All quality measures are converted to percentiles, weighted by student enrolment, from which the deciles are obtained.

3.2 Conditional Match Gaps

Figures 4 and 5 present estimates of the match gaps conditional on student characteristics and prior achievement up until age 16 (equation 1) across the distribution of achievement. Each point represents a separate regression for each achievement decile. Figure 4 plots the achievement gap between the lowest and highest SES quintile, and shows that conditional on demographics and prior achievement, the SES gap is increasing across the achievement distribution up to the ninth decile of achievement, where low SES students undermatch by 11 (9) percentiles more than high SES students for our academic- (earnings-) based measure of match. For top performing students the SES match gap reduces significantly to 3 (4) percentiles. This implies that there are factors in play that ensure that the very best students are well matched to degrees regardless of their level of disadvantage. The largest SES match gaps are found for the above average students, a group of students that have largely been passed over by the literature.

Figure 4:SES conditional match inequalities




Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 1, estimated for each decile of the student achievement distribution. Controls include dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Figure 5 shows a small conditional gradient in undermatch for high-attaining women, relative to men for academic-based match, with a more pronounced conditional gender gap across the entire distribution of achievement for our earnings-based match measure. Again, the top decile of attainers shows a smaller gender gap for this match measure, indicating that the highest attaining women are more similarly matched, relative to the highest attaining men.

Having explored the inequalities in match across the entire distribution, for the remainder of the paper we focus on those students in the top and bottom quintiles of achievement for brevity. Table 2 presents conditional estimates replicating Figures 4 and 5 for these quintiles. Each of the four columns represents a separate regression. The SES parameters represent the match gaps for each SES quintile relative to the highest.

Figure 5: Gender conditional match inequalities



Source: NPD-HESA. n=138,969. Notes: Each point represents the gender match gap from specification 1, estimated for each decile of the student achievement distribution. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Table 2: SES and gender conditional match gaps

| Achievement Quintile | Academic-based match | | Earnings-based match | |
|----------------------|----------------------|-----------------|----------------------|-----------------|
| | Lowest (1st) | Highest (5th) | Lowest (1st) | Highest (5th) |
| SES quintile | | | | |
| 1 | -2.73 (0.47) | -8.35 (0.80) | -5.95 (0.57) | -7.92 (0.88) |
| 2 | -2.57 (0.39) | -5.02 (0.51) | -3.05 (0.49) | -4.91 (0.57) |
| 3 | -1.19 (0.36) | -3.48 (0.35) | -1.78 (0.45) | -4.29 (0.45) |
| 4 | -1.13 (0.33) | -1.99 (0.27) | -0.73 (0.43) | -2.24 (0.34) |
| Women | 0.69 (0.24) | -2.43 (0.25) | -7.48 (0.32) | -8.06 (0.32) |
| Clusters | 2135 | 2005 | 2135 | 2005 |
| N | 27794 | 27786 | 27794 | 27786 |

Source: NPD-HESA. n=138,969. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. *** Significant at the 1% level, ** Significant at the 5% level, * Significant at the 10% level.

Our estimates for low-attainers have a positive constant, indicating that they overmatch on average while the for high-attainers, there is a negative constant indicating that they undermatch on average as shown in Figures 2 and 3. Therefore, the coefficients in columns 1 and 3 represent the SES or gender gap in overmatch, and the coefficients in columns 2 and 4 represent the SES or gender gap in undermatch.

The results in column 2 are similar to previous findings on mismatch that consider the extent of academic undermatch among high-attaining students. We find that there is an 8.4 percentile gap in match for those from the lowest SES quintile relative to those from the highest SES quintile using our academic-based measure of match. This is consistent with the findings in the literature (e.g Hoxby and Avery, 2012; Smith, Pender and Howell, 2013) that high-attaining disadvantaged pupils are more likely to undermatch than their more advantaged counterparts. This 8.4 percentile gap corresponds to the difference between studying economics at the London School of Economics (ranked 5th in the Times Higher UK university rankings) versus Exeter (ranked 18th). This could have real labour market consequences for the student; the median earnings difference five years after graduation between these two degrees is £13,200 per year. The extent of the match gap closes as the difference in the SES quintile narrows: the points gap between the highest SES quintile and the second, third and fourth quintiles are 5.0, 3.5 and 2.0 percentiles respectively.

Column 1 re-estimates these gaps for students from the lowest achievement quintile. Despite all students being in the lowest 20% in terms of achievement, the high SES students attend degrees with higher attaining peers. The mismatch gap is 2.7 percentiles, implying that low SES students overmatch by 2.7 percentiles less than high SES students. This SES match gap is about a quarter of the size of that for high attaining students.

Columns 3 and 4 estimate the earning-based match gaps. For high attaining students (column 4) the SES earnings-based match gaps are of the same magnitude as the academic-based gaps. However, for low attaining students the SES earnings-based gap is three times larger than the academic-based gap. Low-attaining low SES pupils undermatch in earnings by 6 percentiles more than their low-attaining high SES counterparts. This suggests that even among low attainers, more advantaged pupils are more likely to attend degrees with higher labour market rewards. These results represent the first key finding of the paper; that low SES students undermatch more, and overmatch less than high SES students, and this is true for both academic-based match, and earnings match.

The remaining parameter of interest is the coefficient for the female indicator variable, which shows the gender match gap. Conditional on student characteristics and prior achievement low-attaining women are 0.7 percentiles better matched in terms of academic-based match. For high-attaining students, women undermatch by 2.4 percentiles more than men (columns 1 & 2). In contrast, the gender gap is large when considering the earnings-based match in columns 3 & 4. These gaps are of a similar magnitude to the SES gaps, with both low- and high- attaining women undermatching by 7-8 percentiles more than men. This is in line with the raw plots of course quality enrolment by student achievement seen in Figure 4. This is the second of our key findings; that while women attend degrees that are almost as academically selective as men, at every point on the achievement distribution they attend degrees with substantially lower rewards on the labour market. We will return to potential drivers of these gaps, including preferences, in section 4.

3.3 Severity of match


While Table 2 illustrates SES and gender gaps in match for the mean level of match, it may be the case that these inequalities vary across the distribution of match. In particular, we are interested in the extent of these inequalities among cases where students are severely under or overmatched. Figures 6 and 7 explore this using unconditional quantile regression (equation 2) for our earnings-based match measure (see Appendix Figures A4 and A5 for academic-based match). We plot the SES and gender gaps from the 10th to the 90th percentile of the distribution of match for low- and high-attainers. Recall that low-attainers typically overmatch, whereas high-attainers typically undermatch. Therefore, for low attainers (left hand panel of figures 6 and 7), the x-axis runs from those who are matched (at the 10th percentile) to those who are severely overmatched (90th percentile). For high-attainers (right hand panels of figures 6 and 7) the x-axis runs from those who are severely undermatched (10th percentile) to matched (90th percentile). In each case, the estimates represent the earnings-based match gap between the top and bottom SES quintile or between genders. A negative value in Figure 6 represents the degree to which those from the lowest SES quintile are less overmatched (low-attainers), or more undermatched (high-attainers), compared with those from the top quintile. Similarly for Figure 7, a negative value represents the degree to which females are less overmatched (low-attainers) or more undermatched (high-attainers) compared with males.

Looking first at low-attainers (left hand panel of Figure 6) we see that the SES gap is very small for students who are well matched (10th percentile). However as we move along the distribution from matched to severely overmatched, the gap becomes more pronounced. The

implication of this negative gradient is that even within the group of low-attaining students who manage to significantly overreach themselves in terms of the degree they eventually access, students from richer backgrounds still manage to reach further – attending higher earning degrees - than poorer students.²²

A similar pattern is observed with the gender gap for low-attaining students (Figure 7, left hand panel), albeit that the gap is larger throughout the mismatch distribution. For well-matched students, females still attend degrees that are 3.3 percentiles less overmatched than males, and this gap increases to 11.6 percentiles for severely overmatched students.


Figure 6: SES gaps in severity of earnings-based match



Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 2, estimated for each decile of the match distribution. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

²² This only holds for our earnings measure of match. When we use our academic-based measure of match (see Appendix Figures A3 and A4), low-attaining low SES pupils overmatch to a similar extent to high SES pupils and women overmatch to a similar extent to men, when considering those who severely overmatch.

Figure 7: Gender gaps in severity of earnings-based match



Source: NPD-HESA. n=138,969. Notes: Each point represents the gender match gap from specification 2, estimated for each decile of the match distribution. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

For high-attainers, (right hand panels of Figures 6 and 7), the situation is reversed. The SES and gender gaps are largest for the most severely undermatched students (at the 10th percentile of match). In addition to the positive gradient, the size of the gaps are larger for the high attaining students. For the most severely undermatched students, low SES students undermatch 26 percentiles more than high SES students, and women undermatch 16 percentiles more than men. This holds also for the academic-based match measure (see Appendix Figures 4 and 5), with the SES gap among high-attaining students being 31 percentiles and women undermatching more than men by 9 percentiles. This suggests that even among high-attainers who are severely undermatched, low SES students and women are attending degrees that attract far lower financial rewards²³ than they could, compared to those from richer backgrounds and men.

²³ And are attending degrees that are substantially less academically selective

4. Robustness

The construction of our match measure requires us to make a number of decisions. In this section, we use the detailed and extensive nature of our dataset to test these robustness of our findings to these decisions by constructing alternative match measures. Table 3 and Appendix Table A2 presents estimates for the extent of the of match gap for high- and low-attaining students using alternative measures of match. This provides evidence that our results are robust to a number of alternative model choices, including alternative definitions of individual and course quality, and alternative approaches to weighting.

For our main measure of match, as described in Section 2, we adjust the points associated with each A-Level grade to account for the difficulty rating of each subject. One may be concerned with the possibility that certain groups of students choose different types of subjects at A-Level, in which case the difficulty adjusted measure may be endogenous to student SES. An alternative method of dealing with the potential endogeneity of A-level subject choice is to use earlier, broader measure of student achievement. In Columns 3-4 of Table 3 we therefore rank students based on their qualifications from compulsory education at age 16 (GCSE level). Typically students study 10 subjects at this level, and these qualifications are not the main feature of the university application process. We sum the scores across subjects for each student and then calculate their national percentile rank. As with our standard academic-based match measure, the degree ranks are calculated on the basis of the median student on each degree, replacing our standard measure of achievement with the scores from compulsory subjects at age 16. Columns 3 & 4 of Table 3 show that using this earlier measure of student quality makes little difference to our estimates. In Appendix Table 2 we also show that the estimates are also robust to not difficulty adjusting A-level subject (Columns 3 & 4) and using measures of student achievement from age 11 (Columns 5 & 6) which reduces the magnitude of the SES gap, but the gender gap is unaffected.²⁴

Our approach contrasts with much of the existing US literature in that we can observe match at the degree (subject*institution) rather than institution level. In Columns 5-6 of Table 3, for comparability with the existing literature, we condense our data to create a more comparable measure of match, by measuring university quality according to the median student at each university. This shows similar patterns in mismatch inequalities in SES, and gender for academic-based match, but reduced gender gradients in earnings-based match. We discuss this finding in depth in section 5.3.

²⁴ Appendix 1 contains the details for constructing these other measures.

Table 3: SES and gender conditional match gaps across alternative specifications

| Panel A | | Academic Match | | | | | |
|-------------------------------------|--|-----------------|-----------------|-----------------|-----------------|------------------|-----------------|
| | | Baseline | | GCSE-based | | University level | |
| Achievement quintile | | Lowest | Highest | Lowest | Highest | Lowest | Highest |
| SES quintile (ref 5 th) | | (1) | (2) | (3) | (4) | (5) | (6) |
| 1 | | -2.73 (0.47) | -8.35 (0.80) | -3.70 (0.80) | -8.12 (1.03) | -2.72 (0.49) | -9.19 (0.89) |
| 2 | | -2.57 (0.39) | -5.02 (0.51) | -2.97 (0.68) | -4.03 (0.75) | -2.78 (0.42) | -5.40 (0.55) |
| 3 | | -1.19 (0.36) | -3.48 (0.35) | -1.39 (0.61) | -3.10 (0.64) | -1.41 (0.37) | -3.93 (0.41) |
| 4 | | -1.13 (0.33) | -1.99 (0.27) | -1.63 (0.57) | -2.16 (0.53) | -1.53 (0.34) | -2.17 (0.31) |
| Women | | 0.69 (0.24) | -2.43 (0.25) | -3.23 (0.40) | -5.56 (0.53) | -2.40 (0.25) | -3.78 (0.29) |
| Clusters | | 2135 | 2005 | 2135 | 2005 | 2135 | 2005 |
| N | | 27794 | 27786 | 27794 | 27786 | 27794 | 27786 |

| Panel B | | Earnings Match | | | | | |
|-------------------------------------|--|-----------------|-----------------|-----------------|------------------|-----------------|------------------|
| Achievement quintile | | Lowest | Highest | Lowest | Highest | Lowest | Highest |
| SES quintile (ref 5 th) | | | | | | | |
| 1 | | -5.95 (0.57) | -7.92 (0.88) | -7.31 (0.93) | -6.07 (1.10) | -7.09 (0.62) | -10.16 (0.94) |
| 2 | | -3.05 (0.49) | -4.91 (0.57) | -4.35 (0.79) | -3.31 (0.82) | -2.74 (0.50) | -5.55 (0.61) |
| 3 | | -1.78 (0.45) | -4.29 (0.45) | -2.22 (0.70) | -3.59 (0.69) | -1.55 (0.44) | -4.24 (0.47) |
| 4 | | -0.73 (0.43) | -2.24 (0.34) | -1.86 (0.64) | -2.30 (0.58) | -0.86 (0.41) | -2.45 (0.35) |
| Women | | -7.48 (0.32) | -8.06 (0.32) | -9.98 (0.47) | -10.97 (0.59) | -1.99 (0.28) | -3.91 (0.32) |
| Clusters | | 2135 | 2005 | 2135 | 2005 | 2135 | 2005 |
| n | | 27794 | 27786 | 27794 | 27786 | 27794 | 27786 |

Source: NPD-HESA. n=138,969. All specifications control for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. The cubic in age 16 exam results is omitted from the GCSE-based regressions. SES baseline is category 5 (highest quintile).

When calculating the percentile rank of the degree we weight degrees by the number of students in our administrative data. However, in some cases this will not include all students, as our data does not contain students that went to a private secondary school or are international students. Columns 7 and 8 of Appendix Table 2 therefore recalculate the degree percentile ranks using the actual number of students on the degree. This is not used for our main measure because we do not have data on the qualifications of these students. Therefore, for consistency we weight and rank degrees according to our population data. This makes very little difference to our estimates.

Throughout our analysis, we analyse match based on degree-level measures constructed from 23 broad subject levels across every institution. The reason for using 23 broad subject levels is that our earnings-based measure of ‘quality’ is only available at this level and we choose to keep the match measures consistent in this way. However for our academic-based measure, we have access to four digit JACS (Joint Academic Coding System) degree codes, which separately classify around 1,300 different university subjects. For example, we can separately identify those who are studying ‘Economics’ from those who are studying ‘Applied Economics’, and those who are studying ‘History by period’ from those who study ‘History by topic’. In Columns 9 and 10 of Appendix Table 2 we use our more detailed university*subject groupings for our academic-based measure of match to show that our results are robust to using the disaggregated subject categories.

In summary, adopting almost all of these alternative measures of match does not result in any substantial changes to our main findings – low SES students are more undermatched and less overmatched than high SES students in terms of both academic-based and earnings-based match, and high-attaining women are more undermatched than men in terms of earnings-based match. We return to describe why gender inequalities in earnings-based match are less pronounced at the university level in section 5.3.

5. Mechanisms

Given our definition of match, mismatch will be driven by market failures or student preferences. In this section we use various proxies for potential market failures and subject preferences to explore their role as possible drivers for these SES and gender inequalities (again, for simplicity concentrating on earnings-based match)²⁵. We use various measures of geography as proxies for transport/search costs and credit constraints, school characteristics as indicators of (mis)information, and subject choice as an indicator for underlying preferences. For each of these three potential sets of drivers, we condition on additional measures separately to investigate whether our SES and gender gradients in earnings-based match are reduced by the inclusion of these variables. Figures 8 and 9 present the SES and gender gap coefficients after each characteristic is separately added relative to the baseline conditional SES and gender gap reported from Table 2.

²⁵ See Appendix Figures A5 and A6 for academic-based match

Figure 8: SES gaps in earnings-based match, conditional on geography, subject, and schools





Figure 9: Gender gaps in earnings-based match, conditional on geography, subject, and schools




Source: NPD-HESA. n=138,969. Notes: Each point represents the gender match gap from specification 3, estimated for the top and bottom quintiles of the achievement distribution. The baseline controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

5.1 Geography

Geography is often highlighted as a key driver of match (Hoxby and Avery, 2012). Distance to universities could represent a number of market failures. Students living far from a university may be less informed (imperfect information), face a higher search costs from assessing a potential match, and higher credit constraints if they are unable to continue residing with their parents or guardians. A simple plot of distance to university attended by SES, as presented in Figure 10, shows that there is a substantial SES gap in distance travelled to university. In particular, low SES students are far more likely to be found at universities close to their home location.²⁶ If the SES gap in match is driven by geography, with high SES students travelling further in order to achieve a better match then conditioning on distance to university would reduce the gap. However, we find that the inclusion of distance to university attended has no impact on the SES gap for high- or low-attaining students (second column in Figure 8).²⁷ This implies that low SES students undermatch to degrees regardless of distance. However, the distance to university attended is endogenous to the students' choice, therefore we test for the impact of distance using two other pre-determined geographic characteristics.

Figure 10: Distribution of distance travelled to university for low and high SES students



Source: NPD-HESA. n=138,969. Distance to university is calculated using the centre of the student's neighbourhood, defined at the Lower Super Output Area level. This chart includes all students.

²⁶ Here, we define their home location using the centre of the student's neighbourhood, defined at the Lower Super Output Area level. This chart includes all students, but the results are similar if we restrict to high-attaining students only.

²⁷ Other measures of distance to university attended, such as log or quadratics specifications produce similar results.

The first pre-determined measure of proximity relates to the size of the choice set of universities close to the student's home. For this we calculate the total distance to the nearest three universities. If a student lives in an area with several institutions this should improve the probability of a good match, and if high SES students are more likely to be located in such areas this could contribute to the gap. We find this to be the case for low-attaining students, where the inclusion of this term reduces the SES gap by almost half from 5.9 to 3.4 percentiles (see Figure 8). In contrast the inclusion of this parameter has little impact on high-attaining students, suggesting these students are less reliant on universities in their local area. This implies that neither distance nor number of options available are driving the SES gap among high achieving students.

The second pre-determined measure of higher education proximity is the distance student i has to travel to be matched to a degree course, where we define match as a student attending a degree whose quality percentile is within 20 percentage points of their achievement percentile. As with the distance to university attended with find that conditioning on geography in this dimension has little impact on the SES gap for either low or high attaining students (Figure 8). Therefore, we conclude that for students in England, distance to university and its associated market failures does not differentially impact the match to degrees for high or low SES students. The same is true for the gender gap: geographical factors do not reduce the gender gap in match (Figure 9), suggesting that males are not attending better matched degrees because they are travelling further.

5.2 School characteristics

The secondary school that a student attends has the potential to have a large influence on course selection. This could occur through peers, school advisors (imperfect information), or even facilitating subjects offered (long-run credit constraints). In Figures 8 and 9 we included an indicator for each secondary school. This school fixed effect will account for all school-level factors associated with the school, including information, peers, geography, and sorting to schools. The inclusion of school fixed effects greatly reduces the SES gap for both high- and low-attaining students, decreasing the SES gap by 73 and 79 percent respectively. This implies that much of the SES gap in match corresponds to these students attending different types of schools.²⁸ Low and high SES students from the same secondary school tend

²⁸ We are not inferring that the school is causing the (mis)match, as it is will be reflective of students sorting to different schools.

to match to degrees in a more similar manner. However, a significant SES gap still remains, with high-attaining low SES students enrolling in degrees with lower earnings potential than high SES students, by around 2 percentiles.²⁹ We attempt to unpack the school fixed effect by considering six different measures that could potentially relate to how well students from a school would match to university degrees. This is illustrated in Figure 11 (and for achievement-based match in Figure A8).

The first three measures relate to the academic achievement of the school, including school-level achievement, and the provision and uptake of qualifications in subject areas that are ‘preferred’ by universities. We see that the importance of these factors is different for high and low achieving students. For high attaining students, school-level achievement explains around 25 percent of the SES gap in match, while the uptake of ‘preferred’ subjects accounts for an additional percentage point. Students who attend schools where a higher proportion of the student body take more academically challenging qualifications are better matched to courses at university. This could reflect having access to more information about universities.


The next two measures are more direct measures of how much information about universities students at the school may be exposed to; the proportion of the school attending university, and the proportion of the school attending high-status (Russell Group) universities. These measures are more related to information about how universities close the SES match gap further for high achieving students.

For low achieving students none of these school characteristics significantly reduces the SES match gap. The final school characteristic we consider is SES mix, defined as the proportion of the school from the top SES quintile. This has the largest impact on the SES gap for low attainers reducing it by around 30 percent (2 percentage points). Low attaining low SES students have better matches to courses as the proportion of high SES students in their school increases.

Combining all six observable school characteristics accounts for around half of the SES gap for both low and high attainers, but cannot fully explain the total ‘effect’ of schools from the fixed effects model. This leaves scope for other unobservable features of schools, such as the information and advice offered by teachers and councillors, and also the sorting to schools related to student unobservables.

²⁹ Note from Figure 9 that school factors have no impact on the gender gap in match, which is expected as males and females are equally represented in most schools.

Figure 11: SES gaps in earnings-based match, conditional on school factors



Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 3, estimated for the top and bottom quintiles of the achievement distribution. The baseline uses our main specification. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level. “Sch results” is the mean age 18 exam score at the school, “Fac sub prop” is facilitating subjects as a proportion of A-level subjects available, “Fac sub uptake” is facilitating subjects as a proportion of all A-levels taken, “Sch prop uni” is proportion of school students going to university, “Sch prop rg uni” is proportion of school students going to a Russell Group university, “Sch prop high SES” is the proportion of school students in the top SES quintile.

5.3 Preferences

The role of subject preferences has the potential to be a large determinant of mismatch because, in the absence of direct information on student preferences, we have made the assumption that a student is matched to a course when both the student and course are from the same quality percentile, regardless of subject preferences. Implicit in this is the extreme assumption that students have no subject preferences and consider the full range of subject and institution choices before them. In reality, subject preferences are likely to play a role and so some of the students who are defined to be mismatched may not be, from the student’s perspective. Table 3 has already given us some indication of the role of subject variation in our findings by presenting results at the aggregate university level. In this section we take two approaches to determine the role of preferences in mismatch, using different sets of assumptions.

First, we condition on subject studied, allowing us to compare the extent of mismatch between institutions accounting for the average mismatch of that subject nationally. If we observe a reduction in the SES or gender parameters, this implies that some of the respective match gaps are driven by sorting into the subjects taken e.g. females being more likely to study subjects that are more undermatched on average. Any gap remaining implies that this group attend lower quality institutions, regardless of the subjects they study.

Figure 8 shows that the inclusion of subject fixed effects does not substantially impact the SES parameters (e.g. for high attaining students it drops by around 1 percentile), meaning that little of the SES inequalities in match are driven by the subjects that people study at university³⁰. The implication is that even when they have similar qualifications, and are studying similar degree subjects, low SES students study at institutions with lower average earnings. The same conclusion can be drawn using the achievement-based measure of match (Appendix Figure A6), i.e. even when they have similar prior achievement, and are studying similar degree subjects, low SES students study at less academically selective universities.

In contrast, Figure 9 shows that subject studied is the only factor that reduces the gender gap in match. For low-attainers, the gender gap reduces from 7.5 to 1 percentile, when conditioning on subject studied³¹. This implies that low-attaining women attend degrees in lower earning subjects compared to their male counterparts. For high-attaining students conditioning on subject studied reduces the gender gap to 2 percentiles. This suggests that the majority, but not all, of the gender gap among high-attaining students is driven by revealed subject preferences.

Our main measure of match considers all university-subject pairs in a single spectrum, and so a match would occur regardless of subject studied. This is consistent with taking the role of the social planner and therefore assumes that students have no subject preferences. For example, we may consider a student to be matched for history at Bristol, but with a marginal increase in qualifications would then be matched for maths at Exeter. We now consider a second approach to match, making an alternate strong assumption, that students only consider one subject, and choose from a range of institutions within that subject. For example, a student may decide that they only want to study physics and so will only consider physics courses at

³⁰ Similarly, columns 5 and 6 of Table 3 suggested that aggregating our measures up from course to university level had little impact on SES inequalities.

³¹ This was also implicit in the reduced gender gradient in Table 3 column 5 and 6, when measuring match at the university rather than course level.

different institutions. Here, we should only consider matches to physics courses. Of course, it is likely that the reality likely exists somewhere between these two extremes.

In order to explore the match-gaps assuming strong subject preferences we create a within-subject match. Instead of ranking students on the basis of their prior achievement nationally against all other students, we rank students within the group of students who go on to study the same subject at university. For the academic course quality measure, we then recalculate the quality index of each course, using the median students rank within the given subject. For the earnings course quality measure, we generate potential earnings percentiles by subject. As before, we define match as the difference between a student's quality percentile and the course quality percentile.

For example, consider a student at the 80th percentile nationally across all subjects (our main measure), who is at the 98th percentile among students who go on to study music. This student would then be considered well matched if they enrol in the music course at the 98th percentile of music courses, being either the music course whose median student is at 98th percentile of entry qualifications of students that studied music (academic match), or the music course at the 98th percentile of graduate earners among students who took music (earnings match).

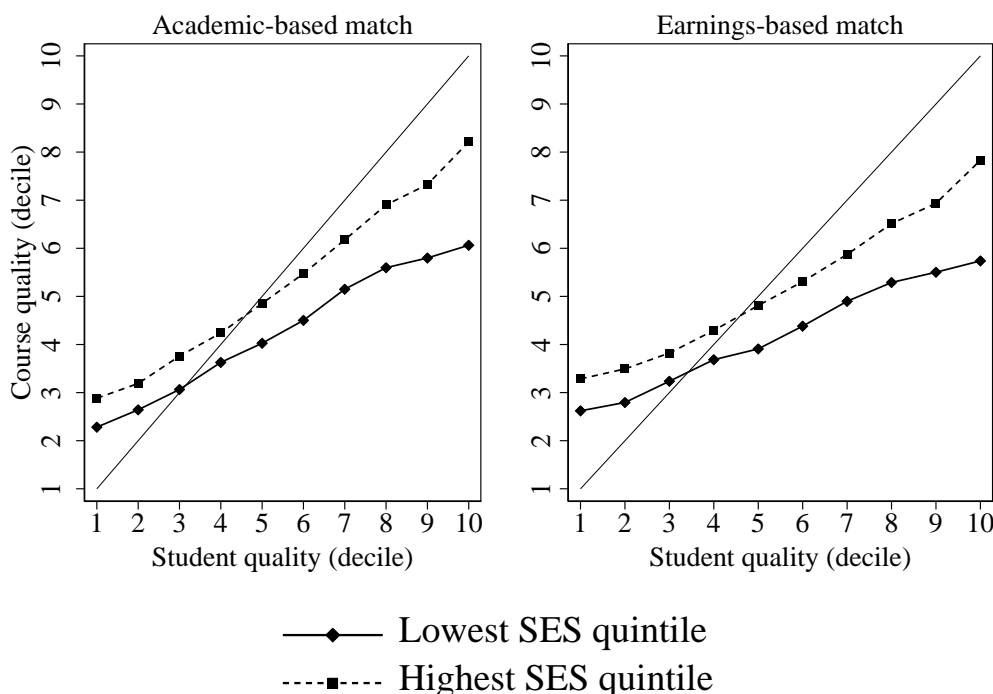
We replicate our non-parametric plots of the match gap throughout the student quality distribution (Figures 2 & 3) using these new measures of course and student quality in Figures 12 and 13 for SES and gender respectively.

Similar to the estimates of match when conditioning on subject taken (Figures 8 & 9), accounting for subject studied in this way, we find that students from low SES backgrounds consistently undermatch, academically and in terms of graduate earnings, throughout the ability distribution within subject. In other words, for students that choose to study a specific subject at university, students from the lowest SES quintile consistently attend universities with lower quality peers in terms of qualifications or lower future earnings. These gaps grow with prior academic achievement. At the top of the distribution low-SES students are undermatching by around 20 percentiles for either match measure.

One possible reason for increase in the gap may be due to these students missing out on places on the most competitive courses on the basis of their predicted grades. Students in the UK are given conditional offers based on predicted grades, and Murphy and Wyness (2020) show that high achieving low SES students are more underpredicted relative to their final

grades compares to their high SES peers, and that they are then less likely to enroll in highly selective courses.

Figure 12: Student and Course Quality Percentile Within Subject, by SES




Source: NPD-HESA, HMRC n=138,969. Notes: The 45-degree line represents perfect matching throughout the achievement distribution. Student quality defined by their percentile within in age 18 difficulty adjusted exam performance distribution for their undergraduate subject choice. Course quality is defined by two measures. Academic Quality defined by their percentile in the courses’ median student’s prior achievement distribution among students studying that undergraduate subject. Potential Earnings Quality defined by their percentile in the courses’ median student earnings five years after graduation distribution among students studying that undergraduate subject. All quality measures are weighted by student enrolment, from which the deciles are obtained.

Appendix Table A3 presents the corresponding conditional match gaps using these within subject measures of quality. For students in the highest achievement quintile the lowest SES group undermatch by 9.3 (10.5) percentiles compared to the highest SES group, in terms of academic (earnings) match (Column 4 (8)). Notably, these within subject match gaps are marginally larger than the gaps without accounting for subject preferences, implying that subject preferences are not driving the SES match gap.

With regards to the gender gap, we continue to find no differences in academic match for lower attaining students (Figure 13). For higher attaining students the gender gap within subject is marginally larger than the gap when there are no subject preferences (3.3 percentiles compared to 2.4 percentiles in our main models – Appendix Table A3 Columns 2 and 4). The

situation is different with regards to earnings match. When not accounting for subject preferences women enrol in courses around 8 percentage points lower in future earnings quality ranking compared to males across the achievement distribution (Figure 3). Accounting for subject preferences however accounts for all of the gender earnings match gap in the lower half of the achievement distribution (Figure 13). Women from the bottom half of the ability distribution are just as well matched as males. For higher achieving students, the potential earnings gap is much reduced, but still exists. Specifically, conditional on prior achievement, high achieving women attend courses with 2.7 percentiles lower earnings potential than men of the same achievement level and subject studied (Appendix Table A3, Column 8). This means that that even when taking preferences into account high achieving females are undermatching relative to men.

Figure 13: Student and Course Quality Percentile Within Subject, by Gender



Source: NPD-HESA, HMRC n=138,969. Notes: The 45-degree line represents perfect matching throughout the achievement distribution. Student quality defined by their percentile within in age 18 difficulty adjusted exam performance distribution for their undergraduate subject choice. Course quality is defined by two measures. Academic Quality defined by their percentile in the courses' median student's prior achievement distribution among students studying that undergraduate subject. Potential Earnings Quality defined by their percentile in the courses' median student earnings five years after graduation distribution among students studying that undergraduate subject. All quality measures are weighted by student enrolment, from which the deciles are obtained.

There are three points to note from these results:

First, the fact that the results from our within-subject measure of match (which makes the strong assumption that students consider university options within one subject) are very similar to the results from our main measure of match after controlling for subject (which reveals the extent of sorting to subjects that are under/overmatched), tells us that students have strong subject preferences (as opposed to none).

Second, subject choice is an important driver of the gender gap in earnings-based match, while, by contrast, we find that subject of study has no impact on the (admittedly small) gender gap in academic match. This implies that women have preferences for subjects that are as equally academically selective as men, but which command lower earnings in the labour market. For example, highly qualified women may choose to study English at a selective institution, while men may choose a degree with an equally high entry requirement, but with higher potential earnings such as a STEM degree (Belfield et al, 2018). This is in line with the STEM literature (Card and Payne, 2017) which finds significant gender gaps in STEM entry. Critically, we are showing that this not driven by differences in the grades or subjects of the qualifications taken by men and women prior to university entry.

Third, conditional on major chosen, females with high qualifications are still attending institutions with lower earnings potential. We find small but significant gender gaps in earnings-based match, even conditional on subject studied.

6. Conclusions

Given the proven complementarities between student and course quality, the type of degree that students enrol in is an important factor in how much society benefits from higher education. Therefore from a social planner's point of view, ideally all students would be well matched to their degree. The existence of mismatch in the UK higher education system then implies that inefficiencies exist in the market. We document large inequalities in student-to-degree (university-subject) match using detailed administrative data from schools, universities and tax records, on some 140,000 students. We create two measures of match, one based on the academic achievement of students, and a new measure of match, characterising university degrees by the median earnings of graduating students.

We find a significant proportion of students are mismatched to the degree they attend. While a direct comparison with other studies is not possible, our results imply that there may be less mismatch in the UK than the US. This may be attributable to the UK's relatively

generous financial system (students are eligible for maintenance loans, and fees are fully covered with income-contingent loans³²), and the fact that there is almost no price variation between degrees, meaning poorer students cannot make a price-quality trade-off. The UK's centralised applications and admissions system, UCAS, which allows students to easily apply to up to 5 university degrees for a very small fee may also be a factor in helping UK students to match well to their degrees.

Yet despite these important features, we still find significant SES inequalities in match. Low SES students are more likely to undermatch and less likely to overmatch on academic-based match. This finding has been documented in previous papers in this area (Dillon and Smith, 2017; Smith, Pender and Howell, 2013). However our earnings-based measure of match shows that not only do disadvantaged students attend less academically selective degrees but they also enrol in lower earning degrees across the achievement distribution. This novel finding has important societal implications: if low SES students are attending degrees with lower returns, this will impact their future potential earnings, and undermine the potential for higher education to have a positive impact on social mobility. We find a key role for secondary school attended in accounting for our SES disparities in match, with the inclusion of school effects eliminating three-quarters of the gap. The academic rigour of schools attended explains part of this story for high-attaining pupils, while the social mix explains part of this story for low-attaining pupils. But around half of the school fixed effects are still unexplained, pointing to the role of other unobserved factors associated with secondary school such as peers, parental sorting, and information provided by the school as likely key drivers for improving student-to-degree match.

Our earnings-match measure also highlights important gender gaps in match. In particular we find that women tend to choose degrees that are as academically selective as men, but with lower associated earnings. For both high-attaining and low-attaining women, subject choice plays a key role. But for high-attaining women, small gaps remain after controlling for subject studied: even where they enrol in a similar field as men, they still appear to study at institutions with lower average graduate earnings. This finding has implications for the gender pay gap, suggesting that higher education plays an important role in this much studied issue.

Recent studies have shown that providing information to low SES students can successfully improve match (McNally, 2016; Dynarski et al, 2018). Similarly, studies have

³² Barr et al (2019) considers the UK system more favourable to those in place in the US

shown that altering individuals' self-perception can impact their choice of field of study, in turn closing gender gaps in subject choice (Owen, 2020; Shan, 2020; Saltiel, 2020). These studies provide suggestive evidence that the link between gender and field of study is not entirely driven by preferences, or that differences in preferences can arise through socialisation (Lordan and Pischke, 2016), and can be malleable through intervention 'upstream'.

In practice it is likely that both preferences and market failures (such as imperfect information) are responsible for the mismatch we observe. However, untangling one from the other is an empirical challenge that is common in the literature. Ultimately, we cannot definitively disentangle preferences from market failures meaning we cannot say how much of our mismatch is to do with preferences and how much is to do with market failures. Women may be well-informed on the earnings potential of subjects, but simply prefer not to study male dominated majors for rational reasons. This could be due to societal norms, or because women prefer to avoid being in male dominated industries. Similarly, it may be the case that low SES students prefer to attend less academically challenging institutions even when their achievement levels suggest they are academically prepared. Thus, we must regard our estimates of mismatch (and those in the wider literature) as an upper bound, even when considering the within subject measures.

Regardless, providing information, advice and guidance in a targeted way that tries to break down existing barriers in terms of both understanding and perceptions, can only result in more informed choices.

References


- Avery, C. (2013). *Evaluation of the college possible program: Results from a randomized controlled trial* (No. w19562). National Bureau of Economic Research.
- Barr, N., Chapman, B., Dearden, L., & Dynarski, S. (2017). Getting student financing right in the US: lessons from Australia and England. *Centre for Global Higher Education*.
- Barr, N., Chapman, B., Dearden, L., & Dynarski, S. (2019). The US college loans system: Lessons from Australia and England. *Economics of Education Review*, 71, 32-48.
- Belfield, C., Britton, J., Buscha, F., Dearden, L., Dickson, M., Van Der Erve, L., Sibieta, L., Vignoles, A., Walker, I. & Zhu, Y. (2018). *The relative labour market returns to different degrees*. Institute for Fiscal Studies.
- Black, S. E., Cortes, K. E., & Lincove, J. A. (2015). Academic undermatching of high-achieving minority students: Evidence from race-neutral and holistic admissions policies. *American Economic Review*, 105(5), 604-10.
- Bolton, P (2019). Student loan statistics. Briefing Paper Number 1079, House of Commons Library, London
- Card, D., & Payne, A. A. (2017). *High school choices and the gender gap in STEM* (No. w23769). National Bureau of Economic Research.
- Carneiro, P., & Heckman, J. J. (2002). The evidence on credit constraints in post-secondary schooling. *The Economic Journal*, 112(482), 705-734.
- Chetty, R. Friedman, J., Saez, E., Turner, N. and Yagan, D. (2020). Income Segregation and Intergenerational Mobility Across Colleges in the United States *Quarterly Journal of Economics*, 135 (3), 1567-1633.
- Chowdry, H., Crawford, C., Dearden, L., Goodman, A., & Vignoles, A. (2013). Widening participation in higher education: analysis using linked administrative data. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 176(2), 431-457.
- Coe, R., Searle, J., Barmby, P., Jones, K., & Higgins, S. (2008). *Relative difficulty of examinations in different subjects*. CEM Centre, Durham University.
- Dale, S. B. and A. B. Krueger (2002). Estimating the Payoff to Attending a More Selective College. *Quarterly Journal of Economics* 117 (4), 1491-1527.
- Department of Education (DfE). (2010). Schools, pupils and their characteristics. January 2010.
- Department of Education (DfE). (2016). Employment and earnings outcomes of higher education graduates: experimental statistics using the Longitudinal Education Outcomes (LEO) data: further breakdowns. December 2016.
- Department of Education (DfE). (2017). Employment and earnings outcomes of higher education graduates by subject and institution: experimental statistics using the Longitudinal Education Outcomes (LEO) data. June 2017.
- Dillon, E. W., & Smith, J. A. (2017). Determinants of the match between student ability and college quality. *Journal of Labor Economics*, 35(1), 45-66.

- Dillon, E. W., & Smith, J. A. (2018). *The consequences of academic match between students and colleges* (No. w25069). National Bureau of Economic Research.
- Dilnot, C. (2018). The relationship between A-level subject choice and league table score of university attended: the 'facilitating', the 'less suitable', and the counter-intuitive. *Oxford Review of Education*, 44(1), 118-137.
- Dynarski, S., Libassi, C. J., Michelmore, K., & Owen, S. (2018). *Closing the gap: The effect of a targeted, tuition-free promise on college choices of high-achieving, low-income students* (No. w25349). National Bureau of Economic Research.
- Firpo, S., Fortin, N. M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77(3), 953-973.
- Gibbons, S., & Vignoles, A. (2012). Geography, choice and participation in higher education in England. *Regional science and urban economics*, 42(1-2), 98-113.
- Hoxby, C. M., & Avery, C. (2012). *The missing "one-offs": The hidden supply of high-achieving, low income students* (No. w18586). National Bureau of Economic Research.
- Hoxby, C. M., & Stange, K. (Eds.). (2020). *Productivity in higher education*. University of Chicago Press.
- Hoxby, C. M., & Turner, S. (2015). What high-achieving low-income students know about college. *American Economic Review*, 105(5), 514-17.
- Kelly, A. (1976). A study of the comparability of external examinations in different subjects. *Research in Education*, 16(1), 37-63.
- Kirkeboen, L. J., E. Leuven, and M. Mogstad (2016). Field of Study, Earnings, and Self-Selection. *Quarterly Journal of Economics* 131 (3), 1057-1111.
- Light, A., & Strayer, W. (2000). Determinants of College Completion. *Journal of Human Resources*, 35(2), 299-332.
- Lochner, L.J., & Monge-Naranjo, A. (2011). The Nature of Credit Constraints and Human Capital, *American Economic Review*, 101(6), 2487-2529
- McGuigan, M., McNally, S., & Wyness, G. (2016). Student awareness of costs and benefits of educational decisions: Effects of an information campaign. *Journal of Human Capital*, 10(4), 482-519.
- McNally, S. (2016). How important is career information and advice? *IZA World of Labor*.
- Mountjoy, J., & Hickman, B. (2020). The Returns to College (s): Estimating Value-Added and Match Effects in Higher Education. University of Chicago, Becker Friedman Institute for Economics Working Paper, (2020-08).
- Murphy, R., Scott-Clayton, J., & Wyness, G. (2019). The end of free college in England: Implications for enrolments, equity, and quality. *Economics of Education Review*, 71, 7-22.
- Ofqual (2017). *Native speakers in A-Level modern foreign languages*. Office of Qualifications and Examinations Regulation.
- Oreopoulos, P., & Dunn, R. (2013). Information and college access: Evidence from a randomized field experiment. *The Scandinavian Journal of Economics*, 115(1), 3-26.

- Sanders, M., Burgess, S., Chande, R., Dilnot, C., Kozman, E., & Macmillan, L. (2018). Role models, mentoring and university applications-evidence from a crossover randomised controlled trial in the United Kingdom. *Widening Participation and Lifelong Learning*, 20(4), 57-80.
- Smith, J., Pender, M., & Howell, J. (2013). The full extent of student-college academic undermatch. *Economics of Education Review*, 32, 247-261.
- U.S. Department of Education (2015). The College Scorecard. <https://collegescorecard.ed.gov/data/>


Appendix

Figure A1: Comparison of SES index with HESA NS-SEC



Source: NPD-HESA. n=138,969. Compares our SES index with the 'National Statistics-Socio-Economic Classification' (NS-SEC) which is available for around 80% of university attendees in our sample. The NS-SEC measure available in the HESA data is a fairly noisy categorical indicator of SES, since it relies on a mapping from the parental occupation which each student enters into their university application form, and has a relatively high level of non-response. Still, it is reassuring that our continuous measure of SES places the categories of the NS-SEC in a plausible ranking.


Figure A2: A-Level subject difficulty correction



| Subject | Adjustment | Subject | Adjustment | Subject | Adjustment | Subject | Adjustment |
|--------------------------------|------------|--------------------|------------|--------------------------------|------------|------------------------|------------|
| Persian | -50.53 | Dance | -15.23 | PE | -3.39 | IT | 5.54 |
| BTEC | -45.88 | A&D: 3D design | -14.84 | RE | -3.34 | Ancient history | 5.74 |
| Bengali | -39.74 | Drama | -14.09 | Psychology (sci) | -2.90 | Spanish | 7.12 |
| Urdu | -38.78 | Fine Art | -13.82 | Geography | -2.53 | Accounting | 7.25 |
| Panjabi | -38.28 | Art and Design | -13.64 | Business studies and economics | -2.44 | Mathematics (pure) | 9.42 |
| Turkish | -35.66 | Portuguese | -13.26 | English literature | -1.85 | Greek | 9.67 |
| Polish | -34.43 | Performing | -11.43 | Psychology (soc) | -0.09 | Computer studies | 10.43 |
| Film | -24.43 | Dutch | -10.04 | Home Economics | -0.02 | Logic/philosophy | 10.65 |
| Communication | -24.31 | D&T: Food | -9.27 | Mathematics (statistics) | 0.75 | Mathematics | 10.85 |
| Russian | -23.44 | Business | -9.20 | Government and Politics | 0.92 | French | 12.43 |
| Modern Greek | -22.60 | Development | -9.05 | Other classical languages | 0.93 | Music | 12.72 |
| A&D: Photography | -22.21 | D&T: Production | -8.68 | Archaeology | 1.08 | German | 13.83 |
| A&D: Graphics | -20.78 | Electronics | -7.99 | Modern Hebrew | 1.64 | Environmental Science | 14.24 |
| A&D: Crit & Contextual studies | -20.43 | D&T: Systems | -7.97 | Italian | 1.67 | Latin | 14.33 |
| Media, film, and TV | -20.30 | English | -7.74 | Gujarati | 3.21 | Science | 18.39 |
| A&D: Textiles | -19.89 | Chinese | -7.51 | Music technology | 3.40 | Biology | 20.41 |
| Arabic | -19.35 | Law | -5.85 | Classical civilisation | 4.67 | Mathematics (further) | 23.12 |
| Vocational double award | -16.45 | English language | -5.50 | History | 4.71 | Chemistry | 23.94 |
| Sociology | -16.27 | Geology | -4.35 | History of art | 5.42 | Physics | 25.05 |
| Japanese | -16.26 | Vocational A-Level | -3.70 | Economics | 5.51 | Biology (human) | 25.32 |
| | | | | | | Additional mathematics | 27.74 |


Source: NPD-HESA. n=138,969. Units are QCA (Qualifications and Curriculum Authority) points, where one A-Level grade is 30 points. Adjustment method discussed in Technical Appendix.

Figure A3: Academic- and earnings-based measures of match




Source: NPD-HESA. n=1,722. Notes: Each point represents a university degree, plotted against our two quality measures: Median graduate earnings percentile; Median student entry qualification percentile.

Figure A4: SES gaps in severity of academic-based match




Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 2, estimated for each decile of the match distribution. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Figure A5: Gender gaps in severity of academic-based match




Source: NPD-HESA. n=138,969. Notes: Each point represents the gender match gap between groups 1 and 5 from specification 2, estimated for each decile of the match distribution. Controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Figure A6: SES gaps in academic-based match, conditional on geography, subject, and schools




Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 3, estimated for the top and bottom quintiles of the achievement distribution. The baseline controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Figure A7: Gender gaps in academic-based match, conditional on geography, subject, and schools



Source: NPD-HESA. n=138,969. Notes: Each point represents the gender match gap between groups 1 and 5 from specification 3, estimated for the top and bottom quintiles of the achievement distribution. The baseline controls are dummies for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level.

Figure A8: SES gaps in academic-based match, conditional on school characteristics



Source: NPD-HESA. n=138,969. Notes: Each point represents the SES match gap between groups 1 and 5 from specification 3, estimated for the top and bottom quintiles of the achievement distribution. The baseline uses our main specification. We present the 95 percent confidence intervals, with standard errors clustered at the secondary school level. “Sch results” is the mean age 18 exam score at the school, “Fac sub prop” is facilitating subjects as a proportion of A-level subjects available, “Fac sub uptake” is facilitating subjects as a proportion of all A-levels taken, “Sch prop uni” is proportion of school students going to university, “Sch prop rg uni” is proportion of school students going to a Russell Group university, “Sch prop high SES” is the proportion of school students in the top SES quintile.

Table A1: Sum of total overmatch and undermatch

| | Total Overmatch | Total Undermatch | Prop. to reallocate for balance |
|---------------------------------------|--------------------|---------------------|---------------------------------------|
| Overall Academic Match | 1,171,280 | -1,176,518 | 0.11% |
| Overall Earnings Match | 1,574,769 | -1,620,956 | 0.72% |
| Figure 3 Gender Academic Match | | | |
| Women | 634,155 | -710,391 | -2.84% |
| Men | 537,125 | -466,127 | 3.54% |
| Figure 3 Gender Earnings Match | | | |
| Women | 702,692 | -1,087,532 | -10.75% |
| Men | 872,077 | -533,424 | 12.05% |

Source: NPD-HESA, HMRC n=138,969. Match is defined as the percentile difference in quality between the student and the course they enrolled on. Student quality is defined by their age 18 difficulty adjusted exam performance. Course quality is defined by two measures. Academic Quality defined by the course's median student's difficulty adjusted exam performance. Potential Earnings Quality defined by the course's median student's earnings five years after graduation. Columns 1 and 2 present the total over(under)match as the product of the number of individuals and their extent of over(under)match. Column 3 presents the proportion of mismatch required to balance overmatch and undermatch.

Table A2: SES and Gender Conditional Match Gaps across Additional Alternative Specifications

| Panel A: Academic-based match | | | | | | | | | | | | |
|-------------------------------------|-----------------|-----------------|--------------------|------------------|-----------------|-----------------|---------------------|-----------------|-------------------|-----------------|-----------------|-----------------|
| Achievement quintile | Baseline | | No diff adjustment | | KS2-based | | Degree size weights | | 631 subject level | | Combined honors | |
| | Lowest | Highest | Lowest | Highest | Lowest | Highest | Lowest | Highest | Lowest | Highest | Lowest | Highest |
| SES quintile (ref 5 th) | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) |
| 1 | -2.73 (0.47) | -8.35 (0.80) | -2.09 (0.44) | -8.90 (0.84) | -1.65 (0.71) | -3.59 (0.93) | -2.83 (0.48) | -7.93 (0.76) | -2.32 (0.42) | -7.34 (0.77) | -2.71 (0.46) | -7.78 (0.79) |
| 2 | -2.57 (0.39) | -5.02 (0.51) | -1.89 (0.37) | -5.45 (0.54) | -1.90 (0.59) | -1.66 (0.61) | -2.64 (0.39) | -4.73 (0.48) | -2.26 (0.36) | -4.37 (0.48) | -2.25 (0.38) | -5.00 (0.50) |
| 3 | -1.19 (0.36) | -3.48 (0.35) | -0.55 (0.36) | -3.77 (0.38) | -1.86 (0.54) | -2.60 (0.48) | -1.24 (0.36) | -3.27 (0.33) | -1.24 (0.35) | -3.03 (0.34) | -1.16 (0.36) | -3.33 (0.35) |
| 4 | -1.13 (0.33) | -1.99 (0.27) | -0.98 (0.33) | -2.28 (0.30) | -1.17 (0.52) | -1.52 (0.42) | -1.16 (0.33) | -1.85 (0.26) | -0.91 (0.32) | -1.79 (0.26) | -1.02 (0.32) | -1.92 (0.27) |
| Women | 0.69 (0.24) | -2.43 (0.25) | 1.49 (0.22) | -3.31 (0.27) | 0.81 (0.38) | -3.82 (0.36) | 0.72 (0.24) | -2.26 (0.24) | 0.98 (0.23) | -2.51 (0.24) | 0.89 (0.24) | -2.26 (0.24) |
| Clusters | 2135 | 2005 | 2135 | 2005 | 2121 | 1983 | 2135 | 2005 | 2135 | 2005 | 2135 | 2005 |
| n | 27794 | 27786 | 27794 | 27786 | 26554 | 26580 | 27794 | 27786 | 27794 | 27786 | 27794 | 27786 |
| Panel B: Earnings-based match | | | | | | | | | | | | |
| Achievement quintile | Lowest | Highest | Lowest | Highest | Lowest | Highest | Lowest | Highest | | | | |
| SES quintile (ref 5 th) | | | | | | | | | | | | |
| 1 | -5.95 (0.57) | -7.92 (0.88) | -5.47 (0.58) | -7.97 (0.95) | -6.66 (0.76) | -0.42 (1.12) | -5.79 (0.56) | -7.76 (0.86) | | | | |
| 2 | -3.05 (0.49) | -4.91 (0.57) | -2.64 (0.51) | -4.87 (0.61) | -3.72 (0.70) | 0.19 (0.71) | -2.98 (0.48) | -4.84 (0.56) | | | | |
| 3 | -1.78 (0.45) | -4.29 (0.45) | -1.45 (0.46) | -4.18 (0.47) | -2.48 (0.61) | -1.74 (0.58) | -1.73 (0.44) | -4.21 (0.44) | | | | |
| 4 | -0.73 (0.43) | -2.24 (0.34) | -0.55 (0.45) | -2.27 (0.36) | -0.64 (0.60) | -0.94 (0.47) | -0.70 (0.42) | -2.21 (0.33) | | | | |
| Women | -7.48 (0.32) | -8.06 (0.32) | -8.30 (0.33) | -10.07 (0.34) | -2.38 (0.42) | -6.94 (0.43) | -7.32 (0.31) | -7.94 (0.32) | | | | |
| Clusters | 2135 | 2005 | 2135 | 2005 | 2121 | 1983 | 2135 | 2005 | | | | |
| n | 27794 | 27786 | 27794 | 27786 | 26554 | 26580 | 27794 | 27786 | | | | |

Source: NPD-HESA. N=138,969. All specifications control for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. These prior achievement controls are omitted for the KS2-based regressions.

Table A3: SES and Gender Conditional Match Gaps Using Within Subject Measures

| Achievement Quintile | Academic Match | | | | Earnings Match | | | |
|-------------------------|----------------------------------|-----------------|--|-----------------|----------------------------------|-----------------|--|------------------|
| | Main Measure (No Preferences) | | Within-subject (Subject Preference) | | Main Measure (No Preferences) | | Within-subject (Subject Preference) | |
| | Lowest (1) | Highest (2) | Lowest (3) | Highest (4) | Lowest (5) | Highest (6) | Lowest (7) | Highest (8) |
| SES quintile | | | | | | | | |
| 1 | -2.73 (0.47) | -8.35 (0.8) | -3.54 (0.54) | -9.29 (0.92) | -5.95 (0.57) | -7.92 (0.88) | -8.00 (0.70) | -10.52 (0.96) |
| 2 | -2.57 (0.39) | -5.02 (0.51) | -3.21 (0.48) | -5.38 (0.59) | -3.05 (0.49) | -4.91 (0.57) | -3.01 (0.57) | -5.74 (0.64) |
| 3 | -1.19 (0.36) | -3.48 (0.35) | -1.63 (0.43) | -3.91 (0.45) | -1.78 (0.45) | -4.29 (0.45) | -1.61 (0.49) | -4.33 (0.50) |
| 4 | -1.13 (0.33) | -1.99 (0.27) | -1.91 (0.40) | -2.19 (0.37) | -0.73 (0.43) | -2.24 (0.34) | -1.02 (0.46) | -2.50 (0.40) |
| Women | 0.69 (0.24) | -2.43 (0.25) | -0.36 (0.27) | -3.34 (0.34) | -7.48 (0.32) | -8.06 (0.32) | 0.68 (0.31) | -2.65 (0.37) |
| Clusters | 2135 | 2005 | 2135 | 2005 | 2135 | 2005 | 2135 | 2005 |
| N | 27794 | 27786 | 27794 | 27786 | 27794 | 27786 | 27794 | 27786 |
| Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

Source: NPD-HESA, HMRC n=138,969. Notes: All specifications control for ethnicity, English as an Additional Language, Special Educational Needs, and gap year before college, and cubics in age 11 and age 16 exam results. SES baseline is category 5 (highest quintile) while the other quintiles are included in the model but not reported here. Match is defined by the difference between student and course quality percentiles. Student quality defined by their percentile within in age 18 difficulty adjusted exam performance distribution for their undergraduate subject choice. Course quality is defined by two measures. Academic Quality defined by their percentile in the courses' median student's prior achievement distribution among students studying that undergraduate subject. Potential Earnings Quality defined by their percentile in the courses' median student earnings five years after graduation distribution among students studying that undergraduate subject. All quality measures are weighted by student enrolment, from which the deciles are obtained.

Technical Appendix

A.1 Difficulty Adjustments

Our approach to accounting for varying difficulty in age 18 examinations across different subjects follows Kelly (1976) and Coe et al. (2008). We calculate difficulty scores for all subjects based on the full set of examination results for the population of students taking age 18 examinations in 2008. To calculate the difficulty of each subject, we subtract each participating student's score in that subject from her average score across all other subjects, and sum the result across all participating students. This gives us a difficulty correction factor for each subject. As Kelly (1976) and Coe et al. (2008) note, these initial difficulty correction factors are likely to underestimate the variation in difficulty across subjects, since students who take 'hard' subjects tend to combine them with other 'hard' subjects, and those who take 'easy' subjects tend to combine them with other 'easy' subjects. We therefore 'correct' each student's score using the initial difficulty correction factors and repeat the process. We do this ten times. With each repetition, the difficulty correction get smaller, and after ten times, the required adjustments have effectively shrunk to zero.

Figure A2 and the accompanying table show the total difficulty correction factor applied to each subject. The units are "QCA points", and 30 points represents one A-Level grade. The 'easiest' subject, Persian, with a total difficulty correction of -51, is therefore 2.6 grades easier than the most difficult subject, Additional Mathematics, which has a total difficulty correction of 28. Intuitively, this means that students who take Persian tend to score higher in that subject than they do in others, while those who take Additional Mathematics tend to score lower.

NB: In the case of Persian and other minority languages, it may be that many students who take the subject already have some understanding of the language, which decreases the perceived difficulty of the subject using our method (see Ofqual, 2017). Only a small numbers of students take these minority language A-Levels – none of A-Level Persian, Bengali, Urdu, Panjabi, Turkish, Polish, or Russian, is taken by more than 100 students in our estimation sample.

A.2 Alternate Match Metrics

Appendix Table A2 presents a number of alternative specifications to supplement further the robustness checks discussed in Section 4. As discussed, we may be concerned that different students choose to study difference A-Level subjects, which are not of equal value. To attempt to reduce the effect of this on our estimated SES and gender gaps, we adjust our

total points by a subject difficulty rating. Our first additional robustness specification in columns 3 and 4 presents the SES and gender gradients without this adjustment, to show that it is making little difference to our estimates.

A further alternative method of dealing with the potential endogeneity of A-Level subject choice is to alternative measures of student achievement to rank students and degrees (for academic-based match). We discuss the use of compulsory exams at age 16 in Section 4. Here we extend this even further to rank students based on their qualifications at age 11 (Key Stage 2). At age 11 all students in England take the same three exams in English, maths and science, which completely removes the issue of student choice. We sum the scores across subjects for each student and then calculate their national percentile rank. The degree ranks are in turn calculated on the basis of the median students using these measures. The results are similar for this very early ranking of student achievement, with a slight reduction in the SES gradient for high-attainers.

Finally, columns 11 and 12 consider an alternative way to specific students on combined honours degrees. About 10% of our sample are doing these types of degrees, where each subject studied falls into more than one group of our 23 subject classification. In our baseline results we assign these students their highest weighted subject, or if weighting is equal they are given the first subject listed. In these results, we assign them according to their highest 2 weighted subjects, so there are 117 different categories including single and combined honours. Note it is not possible to carry out this test for our earnings-based measure because we do not observe later earnings for these combined honours degrees. Our results are again very similar using this alternative specification.

