

Inferring the Performance Diversity Trade-Off in University Admissions: Evidence from Cambridge

Debopam Julia
Bhattacharya Shvets

Abstract

Does increasing diversity in university-intake require sacrificing academic performance, and if so, by how much? We develop an empirical framework to explore this trade-off ex-post, using admissions data matched with post-admission academic outcomes. We propose a simple, theoretical model of admissions for a university that values both future academic performance and diversity, and faces capacity-constraints. We show that the implicit weight on equity vis-a-vis expected future performance in the university's objective-function is captured by the ratio of inter-group difference in the admission-rate and that in the post-entry academic performance of marginal entrants. The problem of identifying marginal entrants can be mitigated using performance data for students admitted from waitlists, leading to bounds for the relative weights. These bounds (a) hold irrespective of whether researchers observe all applicant characteristics known to admission officers and (b) require no information about rejected candidates, who are typically not followed up. We apply this idea to admissions data from Cambridge, using scores on blindly-marked post-admission exams as the performance metric. In mathematical subjects, where female enrolment is relatively low, we find robust evidence that improving gender-balance requires significant performance sacrifice, and conclude an implicit weight of at least 10-20% on gender-equity in the university's objective function. There is no evidence of such trade-off in equally competitive non-mathematical subjects and, contrary to popular perception, for applicants' school-type. Our methods and results illustrate a formal way to quantify ex-post efficiency costs of diversity in a context where societal objective encompasses both equity and efficiency.

Reference Details

CWPE	2238
Published	15 June 2022
Key Words	Affirmative action, Equity-efficiency trade-off, University Admission, Ex-post Evaluation, Marginal Admits, Waitlist Admission
JEL Codes	D61, J71, I23, I24
Website	www.econ.cam.ac.uk/cwpe

Inferring the Performance Diversity Trade-Off in University Admissions: Evidence from Cambridge*

Debopam Bhattacharya[†]

Julia Shvets

University of Cambridge

University of Cambridge

May 15, 2022

Abstract

Does increasing diversity in university-intake require sacrificing academic performance, and if so, by how much? We develop an empirical framework to explore this trade-off ex-post, using admissions-data matched with post-admission academic outcomes. We propose a simple, theoretical model of admissions for a university that values both future academic performance and diversity, and faces capacity-constraints. We show that the implicit weight on equity vis-a-vis expected future performance in the university's objective-function is captured by the ratio of inter-group difference in the admission-rate and that in the post-entry academic performance of marginal entrants. The problem of identifying marginal entrants can be mitigated using performance-data for students admitted from waitlists, leading to bounds for the relative weights. These bounds (a) hold irrespective of whether researchers observe all applicant characteristics known to admission officers and (b) require no information about rejected candidates, who are typically not followed up. We apply this idea to admissions data from Cambridge, using scores on blindly-marked post-admission exams as the performance metric. In mathematical subjects, where female enrolment is relatively low, we

*The opinions expressed in this article are solely those of the authors and do not reflect the views of the Department of Economics, the University of Cambridge and its colleges or the European Research Council.

[†]The authors acknowledge financial support from the European Research Council via grant number 681565, EDWEL. They also express deep gratitude to the Admissions Steering Committee, Dr Alexa Horner and the Central Admissions office at Cambridge for access to the data and to Renata Rabovic for her initial collaboration. We are grateful to Joe Altonji and Frank Kelly for their comments and to Amitabh Chandra and Douglas Staiger for early discussions on this topic.

find robust evidence that improving gender-balance requires significant performance sacrifice, and conclude an implicit weight of at least 10-20% on gender-equity in the university’s objective function. There is no evidence of such trade-off in equally competitive non-mathematical subjects and, contrary to popular perception, for applicants’ school-type. Our methods and results illustrate a formal way to quantify ex-post efficiency costs of diversity in a context where societal objective encompasses both equity and efficiency.

Keywords: Affirmative action, Equity-efficiency trade-off, University Admission, Ex-post Evaluation, Marginal Admits, Waitlist Admission

JEL codes: D61, J71, I23, I24

1 Introduction

From political appointments to workplace recruitment and promotions, the question of diversity is attracting increasing media attention and policy scrutiny around the world. An important example of this general trend is the issue of promoting socioeconomic diversity in admission to elite colleges, given its potential to reduce intra and inter generational economic inequality. At the same time, maintaining academic excellence remains a priority for such institutions. A natural economic approach to incorporating diversity concerns into choice problems would be to allow the decision-maker to pursue an objective function that incorporates both equity and efficiency goals.¹ Behind this formulation is the notion of a potential trade-off between the two. Hence, it is important to understand empirically whether the optimal decision requires trade-offs and if so, to measure their magnitude. However, such evidence is often missing from formal policy-analysis in the context of college admissions, and no empirical method currently exists, to our knowledge, that can be applied to standard admission data in order to measure them. In the present paper, we develop a methodological framework for doing this, and then apply it to anonymized micro-data on undergraduate admissions matched with post-entry academic performance from the University of Cambridge where admission is highly selective and faces intense political and media scrutiny, similar to other elite universities around the world.²

¹For example, Kleinberg et al 2018b consider such dual objectives in a college admissions problem in their work on using algorithms in policy.

²For example, a widely publicized Sutton Trust report in 2018 revealed that Oxford and Cambridge have, in recent years, recruited more students from eight specific schools than almost 3,000 other UK state schools put together. The majority of these eight schools are expensive private institutions. High-profile lawsuits have been fought in the United States on fairness of college admissions, e.g. Regents of the University of

Our analytical framework consists of two elements. The first is a simple theoretical model of admissions for a university that values both future academic performance and equity, and faces aggregate space constraints. The optimization of the university’s objective function leads to an expression for the underlying trade-off solely in terms of quantities potentially estimable entirely nonparametrically from the data. One quantity entering this calculation is the difference in expected performance of marginal entrants (i.e. the implicit ‘cutoffs’ for admission) across different demographic groups. Because admissions are typically based on many indicators, some of which, like confidential reference letters, are typically unavailable to researchers, the marginal entrants are impossible to identify directly. The second contribution of our analysis is to then show how one can use the post-enrolment performance of students entering from waitlists (or second-round clearing) to bound the performance difference between marginal candidates admitted from different groups. Using waitlists for admissions is common in many, if not most, other institutions; as such, our method suggests an approach that is applicable in all such settings. Finally, we show that from the reduced-form estimates of group differences in admission rates and (bounds on) ex-post performance of marginal entrants, one can infer the latent structural parameter representing the implicit weight of equity/diversity relative to academic performance in the decisionmaker’s objective function.

Upon applying the above methodology to the Cambridge data, we reach the following empirical conclusions: for gender, a diversity-performance trade-off exists in the mathematically intensive subjects, viz. Physical Sciences, Engineering, Economics and Mathematics, where the marginal male outperforms the marginal female by at least 0.15 of a standard deviation in university exams. These estimates, in turn, imply that the implicit weight on gender equity in these subjects is between 10-20 percent and that on expected future performance at most 80-90 percent. Interestingly, we find no such trade-off in other highly competitive but non-mathematical subjects such as Law, Biology and Medicine, nor for socioeconomic status as captured by type of school attended by the student. These findings show that the existence of a trade-off is context-dependent; furthermore, they stand in contrast to both the long-standing public perception of Oxford and Cambridge as being socially elitist and the universities’ own statements that their admissions are based solely on ‘academic profile and potential’ (cf. Cummings 2015 and Oxford and Cambridge admission statements).³

California v. Bakke 1978 and Students for Fair Admissions, Inc. v. President and Fellows of Harvard College 2014.

³For example, the official admission statement at Cambridge claims that it aims to “... offer admission to students of the highest intellectual potential, irrespective of social, racial, religious and financial

The measure of performance we use is the end-of-year exam results in the university. Although this may not capture the full benefit from admitting someone to university, it is a natural metric of ‘performance’ to use here as it is the university’s own assessment of students’ academic attainment. Further, for each paper constituting the exam, all students sit the same tests and their scripts are marked blindly, making their scores comparable with one another. Moreover, there is little doubt that exam results affect students’ future opportunities, with employers and post-graduate programs basing their selection on them. Note, however, that our method can potentially be applied to any other post-entry outcome of interest, e.g. post-college earnings or indices of well-being during and after college.

One important and challenging question which emerges from our results is *why* there is a trade-off with performance when it comes to gender diversity in mathematical subjects. There may be a number of competing explanations, and the data we have cannot distinguish perfectly between them. However, we present several pieces of evidence that provide further insight into this, including the possibility that encouraging girls to take more mathematics in secondary school and for them to apply to quantitative subjects may help narrow the gap in university.

Short Literature Review: The study of ‘discrimination’, in the sense of unequal treatment across demographic groups, has a long history in economics, starting with Becker 1957 and Arrow 1973. Compared to the large body of quantitative research on race and gender in the labour market (cf. Altonji and Blank 1999), and the more recent literature on discrimination in law enforcement and credit supply, rigorous evaluation of admissions at elite universities is less common. There is a moderately sized literature in both economics and sociology on the consequences of positive discrimination, e.g. race-based affirmative action in US colleges. Fryer and Loury 2015 provide an overview of the broader questions surrounding the nature and impact of affirmative action. Arcidiacono et al 2015 and Arcidiacono and Lovenheim 2016 summarize existing empirical work in Economics on affirmative action in (mostly US) college admissions with a focus on race. In the educational Sociology literature, Boliver and coauthors, in a series of papers (cf. Boliver, 2013), have analyzed ethnicity-based inequality in UK college admissions using descriptive methods. Bhattacharya et al 2017 tested meritocracy of admissions at a UK university using an assumption about the relationship between observables and unobservables. Sanders 2004 and Ayres and Brooks 2004 investigated the mismatch hypothesis in US law school admissions, Keith et al 1985 did this for medical schools, and Arcidiacono et al 2016 for undergraduate admissions at the considerations.” <https://www.undergraduate.study.cam.ac.uk/applying/decisions/admissions-policy>

University of California.

Our present paper contributes to this literature by developing a theoretical model of equity-efficiency trade-off that can be taken to standard admissions data. The main outcome of this exercise is that it provides a method to infer the weight that a university implicitly places on diversity relative to academic performance, as revealed by its actual admission decisions and subsequent student outcomes. We emphasize that we do not take any stance on what these weights ought to be; our goal is simply to measure what weights are implicitly being used.

Our paper also contributes to the broader empirical literature evaluating the efficiency of how organizations allocate valuable opportunities among competing candidates. Such issues have previously been investigated in the context of firms (Benson, Li and Shue 2019), academia and research (Laband and Piette 1994, Zinovyeva and Bagues 2015, Li 2017), sports (Massey and Thaler 2013) and courts (Arnold, Dobbie and Yang 2018)⁴. These papers use two methods to identify departures from efficiency – the first is based on data on performance of rejected candidates, and the second uses fluctuations in aggregate availability of opportunities which, when used as an instrument, can help identify marginal candidates.

By and large, these methodological approaches to estimating counterfactuals are not feasible when studying admission processes at educational institutions. In particular, universities do not typically collect data on outcomes of rejected applicants, making it difficult to estimate counterfactuals via matching or regression discontinuity type methods. Furthermore, it is unlikely that performance of rejected candidates who eventually enrolled at a different institution is a valid counterfactual for their outcomes had they entered the university under consideration. A second innovation in the present paper is to develop an alternative method for establishing whether there are departures from efficiency and their magnitude in settings where one cannot use the above methods but where, as often happens in educational contexts, the selection process involves multiple tiers.

Also, related to our study are Autor and Scarborough 2008 and Li, Raymond and Bergman 2020 both of which demonstrate that firms can improve equity without sacrificing productivity by changing their selection procedures, and Kleinberg et al 2018a and Kleinberg et al 2018b who build the case for using algorithms to improve selection in public organizations, with applications to judicial bail decisions and university admissions.

⁴Related to this is a growing body of research comparing different methods of allocating these opportunities, including Autor and Scarborough 2008, Kleinberg et al 2018a, Kleinberg et al 2018b, Hoffman, Kahn and Li 2018, and Li, Raymond and Bergman 2020.

The rest of the paper is organized as follows: In section 2 we describe the context and admission process at the University of Cambridge which provides the rationale for our identification method. In section 3, we describe a simple theoretical model of a policy maker deciding how many applicants to admit from different groups so as to maximize an objective that includes both performance and diversity components. We show that the magnitude of the diversity-performance trade-off and can be identified empirically if the de-facto admission cutoffs the policy maker uses for different groups can be identified. Next, we develop an econometric approach, using the two-tier nature of admissions, to calculate bounds on these cutoffs, which addresses the problem that marginal candidates are not directly identifiable. In section 4, we describe our data. In section 5, we show the results from applying our methods to the Cambridge data. First, we identify the trade-offs and their magnitude, and estimate the policy maker’s implied weight on diversity, first for gender and then for socioeconomic status as captured by type of school attended. Then we present some robustness checks. Finally, we present some evidence to better understand our empirical findings, especially the source of the gender gap in performance in mathematically intensive subjects. Section 6 summarizes and concludes.

2 Empirical context

The University of Cambridge is one of the oldest and most prestigious higher education institutions in the world. Every year, around 20 thousand students apply for 3.5 thousand places in Cambridge’s undergraduate programme, making it among the most selective in the world. Roughly $2/3$ of applicants are from the UK, and $1/3$ from other countries.

We will first describe how the students are admitted to Cambridge and then how their performance is assessed once they are in Cambridge. Finally, we touch upon the significance in our context of the school attended by the UK applicants. All three are important for our empirical analysis.

2.1 Admission procedure

In contrast to the liberal arts approach popular in countries like the United States, the UK university system including Cambridge, requires prospective students to apply for a specific subject, e.g. Mathematics, Law, Engineering etc. Having chosen the subject they want to

study for their degree, a student applies to one of Cambridge’s 29 constituent colleges.⁵ The college conducts assessment for that applicant in that subject and makes the final decision on whether to admit the applicant. Each college has an approximately fixed number of places for each subject, which is agreed at the university level. For the purpose of this paper, we view the aggregate admission process as pertaining to the entire university. This is similar to other institutions, where different admission officers make decisions about individual applications but the general policy is set at the university level and so admission results are typically assessed at the level of the university as a whole.

The exact admission procedure at Cambridge varies slightly by subject, but the general procedure consists of the following steps. Students take an admission test, which is the same across all colleges. Those who perform well are then invited to a second assessment – usually an interview – and the rest are rejected.⁶ Those performing sufficiently well on both the second and first assessments and their school-leaving exams, are admitted directly by the college they had applied to. We will refer to these first-tier admissions as ‘direct admission’. A subset of the remaining students who are deemed relatively strong by the college they had applied to are placed in a common ‘pool’. Those colleges that have not filled their quota of places with the first-tier candidates admit one or more students from this pool. We refer to the candidates admitted in this way as the pooled or second-tier admits. About 20% of students in our sample are admitted via the second tier. This two-tier structure of the admission system plays an important part in our empirical strategy.

2.2 Cambridge exams

Most undergraduate degrees in Cambridge are three years long. Every year, students enrol in multiple courses and for each, they typically sit an exam once at the end of the year. For example, in first year, Economics students take five year-long courses, ending in five exams taken in the space of roughly 7-10 days in June. These exams are blindly marked with the only identifying information seen by the examiners is a 5-digit registration number assigned to each student specifically for exams that year. In the vast majority of cases, there are no other official assessments during the year. The final transcript at the end of the degree lists the student’s performance in each year separately: this includes all exam marks by course and an overall classification for this year.⁷ Note that there is no aggregation across years to

⁵In total, there are 31 colleges in Cambridge, but two admit only graduate students, and so are not relevant for this study.

⁶In a few subjects there is no test, so the decision to invite to interview is based on the student’s file.

⁷The classes are first, upper second, lower second, third, where first class is the best etc.

produce a final indicator, unlike, for example, an aggregate GPA in US universities. Thus, for a typical student, three sets of exam performance data are available, one for each year of study.⁸

2.3 Schools

Universities in the UK, including Cambridge, are primarily funded and regulated by the government. Much of this regulation is focussed on widening access to higher education. One of the key measures used by the UK government to assess the socioeconomic background of students is the type of school attended prior to university. For this purpose, all UK schools are divided into two groups: state-funded schools that are free to attend, and independent, fee-charging schools⁹. The latter typically enrol children from higher socioeconomic status households, and these children tend to be overrepresented at the top UK Universities relative to their proportion in the population. As a result, one of the key widening participation targets that the UK government sets for the universities pertains to the proportion of students from UK state-funded schools. Currently, the government’s target for Cambridge University is 64% and is set to increase to 69% by 2024 (University of Cambridge 2018, 2020).

3 Formal framework and identification

This section sets up the theoretical framework and derives the conditions which we then take to the data in section 5 to identify empirically whether there are equity-performance trade-offs and estimate their magnitude. We first provide some graphical intuition to understand the questions of equity and efficiency, the latter meaning post-entry exam performance in our context. Consider figure 1 which plots the distribution of expected performance of two groups of applicants, termed Red and Blue. The top panel corresponds to the admission decision which uses the same cutoff for the expected performance in both groups which is equivalent to equalizing the performance of the marginal candidates from the two groups.¹⁰ This, however, leads to different numbers of students entering from each group, because the right tail of the blue distribution is thicker than that of the red distribution. This corresponds

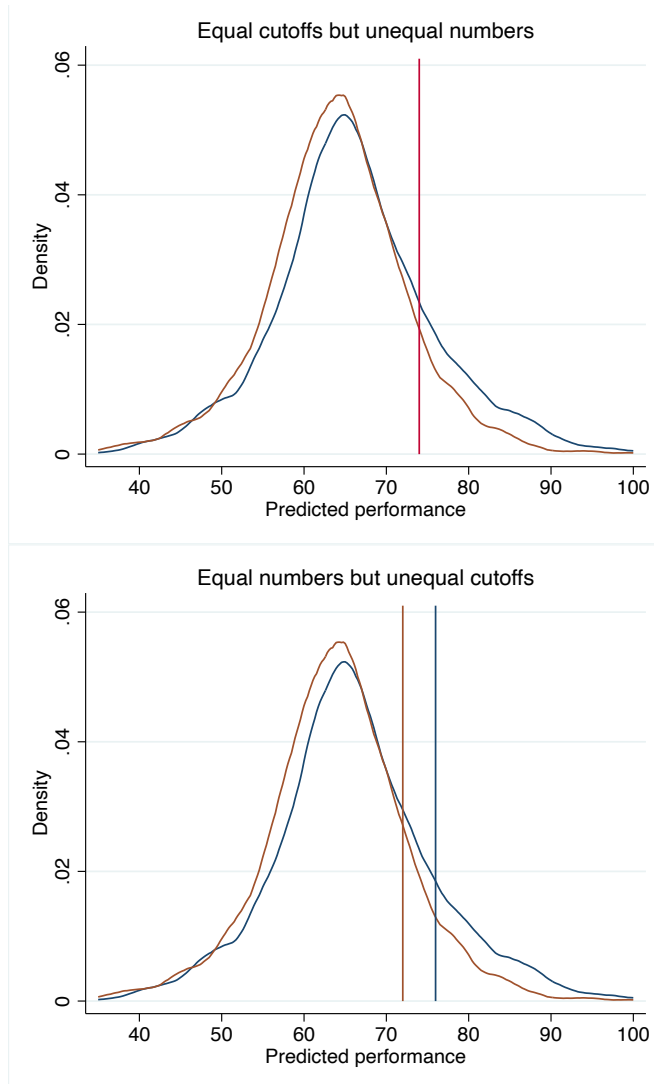
⁸In our data, notable exceptions are Engineering, Mathematics and Medicine that (can) last longer. We discuss this in the section 4 where we describe our data.

⁹In total, about 6% of UK children attend fee-paying schools (Sibieta (2021)), although this proportion rises for 16+ year olds.

¹⁰In what follows, we will use the terms ‘difference in cutoffs’ and ‘difference in marginal performance’ interchangeably

to the case where the decisionmaker wants to simply admit the highest performing applicants and thus ends up admitting more Blues than Reds. The bottom panel, on the other hand, corresponds to the decision where the number of Blue and Red students entering is made equal by fixing the admission cutoff to be lower for the Red group. These represent the two extreme cases where the decisionmaker cares only about efficiency or only about equity.

Figure 1: Potential equity-performance trade-off



Now consider a decisionmaker who values both objectives. We say that a trade-off between equity and performance exists if such a decisionmaker has to sacrifice some performance in order to improve equity and vice versa. This is the case that we formalize and estimate using the data. We do this in two steps. First, in subsection 3.1 we write down the

optimization problem of a decisionmaker who needs to choose how many students to admit from two groups, and who values both performance across and equity between these groups. We show that we can then back out whether the policy maker faces a trade-off between performance and equity by looking at the difference in the de-facto admission cutoffs for the two groups implied by the policy maker’s admission decisions. However, empirically these cutoffs are often unobserved, particularly when decisions are based on a number of factors, some of which, like non-anonymous reference letters, may not be stored in the data and hence will be unobservable when these decisions are reviewed either by researchers or the university itself. Hence, as a second step, in subsection 3.2 we develop a method to overcome this problem and to detect the sign and magnitude of the differences in the implicit admission cutoffs using data on student exam performance and the two-tier admissions structure.

3.1 Diversity-performance Trade-off

Let the two groups of applicants under consideration be denoted by g and h – e.g. g can be males and h females, or h can be state-school educated and g privately educated, etc. Denote by A_g the academic ability (i.e. future academic performance) of a g -type applicant, as inferred by the university.

We assume that the university’s objective function is composed of two terms – one being the average expected performance of the admitted and the other reflecting concerns for equity.

Since the university cares about average performance, they will admit the best students from each group. Recall that in our setting the university practices a two round admission process, first is direct admission, and second the so-called pool. Hence, the admission decision can be summarized via three cutoffs $g_2 < g_3 < g_1$ such that if $A_g > g_1$, then the applicant is admitted directly, if $A_g < g_2$ then s/he is rejected straight away, and if $g_2 < A_g < g_1$, then the candidate is put in the pool. Finally, if $g_3 < A_g < g_1$, then the candidate is eventually admitted from pool. For h -type applicants, denote the analogous quantities by A_h and h_1, h_2, h_3 respectively.

The main insight of our framework is the following: observed difference in cutoffs for direct admission, g_1 and h_1 , will identify the presence and give the size of the equity-performance trade-off faced by the policy maker (if any), and allow us to estimate the implied weight of equity v efficiency motives in the policy maker’s objective.

In what follows we assume that, as its equity objective, the university wants to equalize

the number of admitted students from the two groups, g and h .¹¹ Then our problem can be formally stated as follows. Suppose number of applicants in groups two groups g and h is N_g , N_h respectively; let $\pi_g \equiv \frac{N_g}{N_g + N_h} \equiv 1 - \pi_h$ and total number of places is M . The overall density of A , as defined above, is $\pi_g f_g(\cdot) + \pi_h f_h(\cdot)$, where $f_g(\cdot)$ and $f_h(\cdot)$ are the densities of A_g and A_h in the two groups g and h respectively, with associated CDFs F_g and F_h .

Assuming a quadratic loss function for departures from equality, the university's optimization problem can be formally stated as

$$\max_{g_1, h_1} \left[\begin{array}{l} \pi_g E(A_g \times 1\{A_g > g_1\}) + \pi_h E(A_h \times 1\{A_h > h_1\}) \\ -\beta (\pi_g \Pr(A_g > g_1) - \pi_h \Pr(A_h > h_1))^2 \\ \text{subject to} \\ \pi_g \Pr(A_g > g_1) + \pi_h \Pr(A_h > h_1) = \frac{M}{N} \end{array} \right]$$

where $N = N_g + N_h$ and $\beta \in (0, \infty)$ denotes the value of performance the university would sacrifice per unit of social welfare gained from reduction in inequality.¹² Equivalently, $\frac{\beta}{1+\beta}$ equals the weight of equity relative to efficiency in the university's objective function.

Using the notation introduced above, the above optimization problem can be written as

$$\begin{aligned} \max_{g_1, h_1} & \left[\pi_g \int_{g_1}^{\infty} a f_g(a) da + \pi_h \int_{h_1}^{\infty} a f_h(a) da - \beta (\pi_g (1 - F_g(g_1)) - \pi_h (1 - F_h(h_1)))^2 \right] \\ \text{s.t. } & \pi_g (1 - F_g(g_1)) + \pi_h (1 - F_h(h_1)) = \frac{M}{N} \end{aligned} \quad (1)$$

Claim 1 *If A_g and A_h are continuously distributed and $0 < \beta < \infty$, then the problem (1) has a unique interior maximum.*

Proof. Let $r = 1 - F_g(g_1)$ and $s = 1 - F_h(h_1)$. Since $F_g(\cdot)$, $F_h(\cdot)$ are strictly increasing and continuous, maximizing with respect to g_1, h_1 is equivalent to maximizing w.r.t. r, s . Therefore, (1) is equivalent to

$$\begin{aligned} \max_{r, s} & \left[\pi_g \int_{F_g^{-1}(1-r)}^{\infty} a f_g(a) da + \pi_h \int_{F_h^{-1}(1-s)}^{\infty} a f_h(a) da - \beta (\pi_g r - \pi_h s)^2 \right] \\ \text{s.t. } & \pi_g r + \pi_h s = \frac{M}{N}. \end{aligned}$$

¹¹In the Appendix, we also consider an alternative equity objective, where the university wants to equalize applicant success rates between the two groups. This can be a more reasonable objective if the number of applicants from the two groups is very different, as is the case with females in mathematically intensive subjects in our data. However, our substantive results do not depend on which of the two equity objectives we choose.

¹²We assume $\infty > \beta > 0$ since (i) negative beta would imply desire for inequality and (ii) zero or infinite beta is the case where policy maker does not care about one of the two objectives and hence makes the central question of this paper irrelevant.

Now note that

$$\begin{aligned}
\phi(r) &= \int_{F_g^{-1}(1-r)}^{\infty} a f_g(a) da \\
\implies \phi'(r) &= F_g^{-1}(1-r) f_g(F_g^{-1}(1-r)) \frac{1}{f_g(F_g^{-1}(1-r))} \\
&= F_g^{-1}(1-r) \\
\implies \phi''(r) &= -\frac{1}{f_g(F_g^{-1}(1-r))} < 0,
\end{aligned}$$

so $\phi(r)$ is strictly increasing and concave in r ; similarly, $\int_{F_h^{-1}(1-s)}^{\infty} a f_h(a) da$ is strictly increasing and concave in s . Clearly, $-\beta(\pi_g r - \pi_h s)^2$ is concave in s and r since $\beta > 0$ and $\pi_g, \pi_s > 0$. Since sum of strictly concave functions is strictly concave, it follows that

$$\left[\pi_g \int_{F_g^{-1}(1-r)}^{\infty} a f_g(a) da + \pi_h \int_{F_h^{-1}(1-s)}^{\infty} a f_h(a) da - \beta(\pi_g r - \pi_h s)^2 \right]$$

is strictly concave in (r, s) , while the constraint is linear in (r, s) . Therefore, first-order conditions on the Lagrangian are necessary and sufficient for a maxima. ■

First-order conditions for (1) are:

$$\begin{aligned}
-\pi_g g_1 f_g(g_1) + 2\beta \times (\pi_g(1 - F_g(g_1)) - \pi_h(1 - F_h(h_1))) \times \pi_g f_g(g_1) &= \lambda \pi_g f_g(g_1) \\
-\pi_h h_1 f_h(h_1) - 2\beta \times (\pi_g(1 - F_g(g_1)) - \pi_h(1 - F_h(h_1))) \times \pi_h f_h(h_1) &= \lambda \pi_h f_h(h_1)
\end{aligned}$$

where λ denotes the Lagrange multiplier, implying

$$\begin{aligned}
g_1 - 2\beta \times (\pi_g(1 - F_g(g_1)) - \pi_h(1 - F_h(h_1))) \\
= h_1 + 2\beta \times (\pi_g(1 - F_g(g_1)) - \pi_h(1 - F_h(h_1)))
\end{aligned} \tag{2}$$

Note that $\pi_g(1 - F_g(g_1)) - \pi_h(1 - F_h(h_1))$ captures the difference between the proportions of the two groups among admitted students and hence the deviation from equality. Denoting this by Δ , (2) is equivalent to

$$\beta = \frac{g_1 - h_1}{4\Delta} \tag{3}$$

The numerator of the above expression $g_1 - h_1$, the difference between the cutoffs for direct admission for the two groups, is also the difference between the perceived ability of marginal candidates from the two groups. If the decisionmaker were maximizing just the performance objective ($\beta \rightarrow 0$), they would set $g_1 - h_1 = 0$. Conversely, if they were only maximizing equity ($\beta \rightarrow \infty$), they would set $\Delta = 0$. Note also that the derivation of equation (3) did not require any functional form assumptions on unobservables, e.g. that ability is normally distributed etc. In that sense, it is a nonparametric (identification) result that

expresses the latent structural object β , characterizing the university's preferences, in terms of the reduced form parameters $g_1 - h_1$ and Δ .

Suppose we can observe $g_1 - h_1$ and Δ empirically. Given our assumption that $\infty > \beta > 0$, i.e. that policy maker cares about both objectives, we can split the space of these observations into two cases arising from the first order condition (3):

- $g_1 - h_1 = \Delta = 0$. This implies that decision-maker has maximized both objectives, and hence we can conclude that there is no trade-off between them. This will happen if, for example, the right tail of potential performance distribution is identical for the two groups. In absence of a trade-off, the decisionmaker will maximize both objectives regardless of the value of β , and so we cannot learn the value of β from the data.
- $g_1 - h_1 > 0$ and $\Delta > 0$. This implies that the decision-maker has not managed to maximize either objective, and hence we can conclude that they face a trade-off between them. Equation (3) also provides us with a natural measure of the trade-off in our context: we will define the size of the trade-off as $g_1 - h_1$, marginal performance sacrificed to keep inequality from rising beyond the current level Δ ¹³.

Furthermore, from this we can back out the decision-maker's $\frac{\beta}{1+\beta}$, i.e. the implied relative weight of equity in terms of performance of the university, since β is given by equation (3) above.

To summarize, the above results tell us how to detect the existence of equity-performance trade-off and its size from the admissions decisions, provided we know Δ , the differences between the shares of the two groups among admitted students, and the performance cutoffs g_1 and h_1 for the two groups. Whilst Δ is observed in the data, performance cutoffs are not, since admission depends on a lot of different variables associated with each applicant, some of which are often unobservable in the data. Therefore, in the next subsection, we develop a method to identify the differences in admission cutoffs, by exploiting Cambridge's two-tier admission process.

¹³Alternatively, one could measure the size of trade-off as Δ since it shows how much equality has been sacrificed to keep performance at the current level. However, this alternative definition relies on the specific equity objective we have chosen (in particular, that the policy maker wants to equalize the number of admitted students from the two demographic groups). In contrast, the measure of trade-off we adopt in the text, $g_1 - h_1$, does not have this problem and is the same regardless of the details of the equity objective.

3.2 Identifying differences in cutoffs

Assume that, for admitted students, performance in the subsequent university exams is generated by $Y_g = A_g + \varepsilon_g$ for g -type and $Y_h = A_h + \varepsilon_h$ for h -type respectively, where ε_g and ε_h are stochastic noise terms that affect exam scores over and above A_g and A_h , ability levels observed at admission. We will assume that admission officers' inference is correct on average, i.e. $E(\varepsilon_g|A_g) = 0$ and $E(\varepsilon_h|A_h) = 0$. This may be justified as follows. Suppose the admission officers observe the set of characteristics X_g for each group g , and from this infer the ability $A_g = E(Y_g|X_g)$. Then $\varepsilon_g = Y_g - E(Y_g|X_g)$, implying by definition that $E(\varepsilon_g|X_g) = 0$, and therefore $E(\varepsilon_g|A_g) = 0$, since A_g is solely a function of X_g . In particular, this implies that with $F_g(\cdot)$ denoting the marginal CDF of A_g , any set C with $F_g(\cdot)$ -positive probability, we have that:

$$E(\varepsilon_g|A_g \in C) = \int_{a \in C} \underbrace{E(\varepsilon_g|A_g = a)}_{=0} dF_g(a) da = 0. \quad (4)$$

This implication will be used below.

The econometrician observes the distributions of exam scores of all entrants and, in particular, those of pooled g -type admits: $Y_g|g_3 < A_g < g_1 = A_g + \varepsilon_g|g_3 < A_g < g_1$, and of directly admitted h -type admits: $Y_h|A_h \geq h_1 = A_h + \varepsilon_h|A_h \geq h_1$. Therefore, average exam score of pooled g -type admits equals

$$\begin{aligned} E[Y_g|g_3 < A_g < g_1] &= E[A_g + \varepsilon_g|g_3 < A_g < g_1] \\ &= E[A_g|g_3 < A_g < g_1] + \underbrace{E[\varepsilon_g|g_3 < A_g < g_1]}_{=0 \text{ by (4)}} \\ &= E[A_g|g_3 < A_g < g_1] < g_1, \end{aligned} \quad (5)$$

while the average exam score of directly admitted h -type admits equals

$$\begin{aligned} E[Y_h|A_h \geq h_1] &= E[A_h + \varepsilon_h|A_h \geq h_1] \\ &= E[A_h|A_h \geq h_1] + \underbrace{E[\varepsilon_h|A_h \geq h_1]}_{=0 \text{ by (4)}} \\ &= E[A_h|A_h \geq h_1] \geq h_1. \end{aligned} \quad (6)$$

It follows from (5) and (6) that

$$g_1 - h_1 > E[Y_g|g_3 < A_g < g_1] - E[Y_h|A_h \geq h_1]. \quad (7)$$

The RHS, which is estimable from our data, thus provides a **lower bound** on the difference in the cutoffs for direct admissions $g_1 - h_1$. In particular, if the average exam score for pooled g -type admits is (weakly) higher than that of directly admitted h -type admits, i.e.

$$E[A_g + \varepsilon_g | g_3 < A_g < g_1] \geq E[A_h + \varepsilon_h | A_h \geq h_1], \quad (8)$$

then $g_1 > h_1$. In fact, if $\Pr(g_3 < A_g < g_1) > 0$, and $\Pr(A_h > h_1) > 0$ – corresponding to the likely scenario that perceived ability is continuously distributed – then even *equality* of mean exam scores, i.e. $E[A_g + \varepsilon_g | g_3 < A_g < g_1] = E[A_h + \varepsilon_h | A_h \geq h_1]$ will also imply $g_1 > h_1$.

These ideas are graphically illustrated in Figure 2, where the ability distribution of the two hypothetical groups are plotted in the two panels, and the rightmost vertical lines in the top and bottom graph represent g_1 and h_1 , respectively. The shaded areas in the top and the bottom panels correspond to $E[A_g + \varepsilon_g | g_3 < A_g < g_1]$ and $E[A_h + \varepsilon_h | A_h \geq h_1]$ respectively. If $g_1 = h_1$, then every individual in the top shaded area (for group g) must score lower than every individual in the bottom shaded area (for group h), up to random noise. Therefore, the average in the top must be less than the average at the bottom. A contradiction implies that the rightmost vertical line in the top figure must be at a higher level of ability than the rightmost vertical line in the bottom figure.

Note that reversal of the inequality (8), i.e. $E[A_g + \varepsilon_g | g_3 < A_g < g_1] < E[A_h + \varepsilon_h | A_h \geq h_1]$, is consistent with both $g_1 > h_1$ and $g_1 < h_1$, and is therefore inconclusive.

Note also that if we find that (8) holds empirically, we do not and cannot distinguish between the possibility that the observed difference in de-facto cutoffs arises either (i) because admission officers have systematically biased beliefs against g -types (hence violating our earlier assumption that $E(\varepsilon_g | A_g) = 0$) but use the same actual cutoff of expected performance for everyone, or (ii) because they have correct beliefs on average but use a systematically higher cutoff for g , for example to promote equity.

Risk and Stochastic Dominance: Suppose admission officers are risk averse and base the admission decision on $B_g \equiv E(U(Y_g) | X_g)$ for a concave, increasing $U(\cdot)$, instead of $E(Y_g | X_g)$. Now if the distribution of $Y_g | g_3 < B_g < g_1$ first-order stochastically dominates (FOSD) that of $Y_h | B_h \geq h_1$ (as found in the empirics and reported below), then $E[U(Y_g) | g_3 < B_g < g_1] \geq E[U(Y_h) | B_h \geq h_1]$ for *all* increasing $U(\cdot)$, and therefore,

$$g_1 - h_1 > E[U(Y_g) | g_3 < B_g < g_1] - E[U(Y_h) | B_h \geq h_1] \geq 0. \quad (9)$$

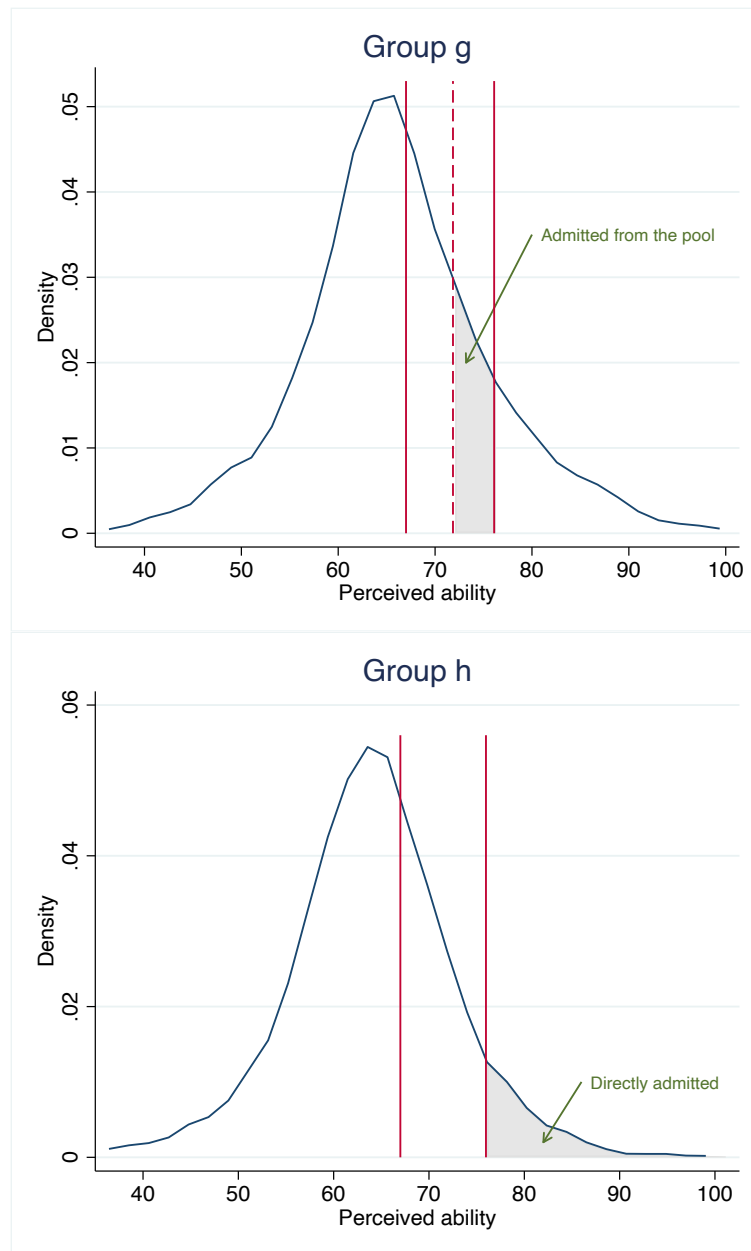
Thus under FOSD, the conclusion of higher admission standards for g -types, i.e. the positive sign of $g_1 - h_1$ is robust to risk-aversion considerations.

Upper Bound: A completely analogous argument leading to (7) applies mutatis mutandis to the upper bound, viz.

$$g_1 - h_1 < E[A_g + \varepsilon_g | A_g > g_1] - E[A_h + \varepsilon_h | h_3 < A_h < h_1] \quad (10)$$

The RHS of (10) is the difference in average performance between directly admitted g -types and the h -types admitted from the pool, and it provides an estimate of the upper bound of the difference in admission cutoffs. In the empirical application, we will use (7) and (10) to bound the numerator of (3) which, together with the point-identified Δ , will then yield bounds for β , the relative weight on equity.

Figure 2: Identifying differences in cutoffs



3.3 Summary of empirical implications

We now summarize the empirical implications of our formal framework.

In section 3.1, we have shown that if we observe Δ , the gap in the shares of the two groups among admitted students, and $g_1 - h_1$, the difference between cutoffs for direct admissions or, equivalently, performance of marginal candidates from the two groups, we can establish whether there is a trade-off between performance and equity. If there is trade-off, we can then establish its direction, and estimate its magnitude at the margin, as well as the magnitude of β , the amount of performance the decisionmaker is willing to sacrifice in order to improve equity at the margin, or equivalently $\frac{\beta}{1+\beta}$ which represents the relative weight of equity vis-a-vis performance in the university's objective function.

While Δ is observed in the data, $g_1 - h_1$ is not. Nevertheless, in section 3.2 we have shown how to identify lower and upper bounds of $g_1 - h_1$ (given by (8) and (10) respectively). From now on, we will denote the lower bound as $E(Y_g|pool) - E(Y_h|direct)$ and upper bound as $E(Y_g|direct) - E(Y_h|pool)$.

As discussed in section 3.1, a key to identifying whether there is a trade-off between performance and equity is knowing whether $g_1 - h_1$ is above zero. Since we do not observe $g_1 - h_1$, but rather its lower and upper bound, we have some freedom as to how we decide when we identify a trade-off. We will use the most conservative rule as follows:

- **Trade-off:** We will say that we identify a trade-off only when *both* estimated lower and upper bounds of $g_1 - h_1$ are significantly above zero. Then, $E(Y_g|pool) - E(Y_h|direct)$ and $E(Y_g|direct) - E(Y_h|pool)$ will give us the lower and the upper bound for the magnitudes of the trade-off on the margin. This, in turn, will allow us to calculate the lower and upper bounds on $\frac{\beta}{1+\beta}$ using equation (??), as well as their confidence intervals which follow from the calculations standard in the partial identification literature, cf. Stoye 2021. Again, to be conservative, we will focus on the lower bound of our estimates of $\frac{\beta}{1+\beta}$.
- **No evidence of trade-off:** When the lower bound of the cutoff difference is significantly below zero or includes zero in its confidence interval and the upper bound is significantly above zero, we will say that there is no conclusive evidence of trade-off.

Note that the above cases cover all possible combinations of lower and upper bound estimates, except the one where $E(Y_g|direct) - E(Y_h|pool) > 0$ and $E(Y_g|direct) - E(Y_h|pool) \leq 0$. This case means that the candidates taken from the pool generally outperform directly

admitted ones, implying that the two-tier admission system is not working as intended and hence the basis of our identification strategy is not valid. We show in section 4.3 that in our data, this violation never occurs, i.e. both in aggregate and within each groups of interest, pooled candidates underperform directly admitted ones, as intended by the two-tier admission system.

4 Data

In this study, we utilize administrative micro-data on individual students from the Cambridge Admissions Office. The data contain student characteristics (gender, school type, school leaving grades, nationality etc.), their placement in Cambridge (admission tier (pool or direct), subject, college) and their exam performance after entering Cambridge.

Our sample consists of 5,888 students who entered Cambridge in 2013-2016 for the following subjects: Economics, Engineering, Mathematics, Biological Natural Sciences, Law, Medicine and Physical Natural Sciences¹⁴.

Among these, Economics, Engineering, Mathematics and Physical Natural Sciences have a higher mathematical contents and require more advanced mathematical preparation than Law, Medicine and Biological Natural sciences. In Cambridge, the former group of subjects have a minimum mathematics requirement that the applicants must fulfil before coming to Cambridge, while the latter group do not (University of Cambridge 2022a)¹⁵.

Henceforth, we will refer to the former group as mathematically intensive (MI) subjects. The two groups of subjects are similarly competitive as measured by the offer rates: they are 22% for non-MI subjects and 24% for MI subjects¹⁶.

In studying equity-performance trade-off, we focus on two student characteristics that

¹⁴This is all students accepted in these subjects after data cleaning and exclusion of those who dropped out, changed subject, postponed their studies or did not sit exams for some other reason.

¹⁵In more detail, Economics, Engineering and Maths require applicants to have a minimum level of mathematics, e.g. Mathematics A-level in the UK school system, for their application to be considered. In contrast, Medicine requires one of ‘biology, physics or mathematics’, and Law does not have any such requirement. Historically, Cambridge groups Biological and Physical Natural Sciences under the heading of ‘Natural Sciences’, which requires ‘any two of Biology, Chemistry, Physics, Mathematics and Further Mathematics ... also subject to the requirements for Year 1 options’. This refers to the fact that the students decide whether to do Physical or Biological Sciences as soon as they start their degree, and in order to enrol into Physical Sciences courses they are required to have taken Mathematics prior to coming to Cambridge. There is no such requirement for Biological Sciences courses (University of Cambridge 2022b).

¹⁶This excludes Biological and Physical Natural Sciences due to data limitations.

are of key concern when it comes to equity: gender and socioeconomic status. Gender is observed in our dataset. As proxy for socioeconomic status we will use whether the student attended a UK state-funded school. Although crude, this is one of the key variables used by the UK government as their broadening participation measure, as we discussed in section 2, so is very relevant in our context.

For student performance, we use marks in their end-of-year exams, the same measure used by the University of Cambridge (see section 2 for more details). The exams are held at the end of the academic year, one for every course the student took that year. Typically, we do not have access to individual course exam marks.¹⁷ Instead we use the average percentage of marks obtained by the student across all of their exams taken that year, standardized by subject, as our outcome of interest. Since exams take place once a year over three years (the length of most undergraduate degrees at Cambridge), we have three performance observations per student, one in each year of their degree¹⁸. Our initial analysis focuses on first year performance and is then extended to subsequent years to examine the longer term validity of our main conclusions.

4.1 Gender

The gender composition of our sample is summarized in table 1. Among the intakes we analyze, 36% are female, but with marked difference across subjects. In Law, Medicine and Biological Natural Sciences, the number of females is close to or over 50%. However, in mathematically intensive subjects it is significantly lower. This pattern of female underrepresentation in mathematically intensive subjects is in line with what has been widely documented before in different settings (e.g. Wang and Degol 2017). At the same time, the gender gap in probability of offer conditional on application is negligible – 23.9% and 23.2% respectively for males and females, with some heterogeneity across subjects^{19, 20}.

It is also apparent from table 1 that in terms of average performance, males outperform females in the first-year exams in all subjects. The average first year gap is 3.5 percentage

¹⁷The one exception is Economics third year exams, which we use in our analysis in section 5.3.

¹⁸In our sample, there are some exceptions to the three year rule: Engineering degree is four years long, Mathematics has an option to proceed to a fourth ‘bonus’ year and Medical degree is six years, including clinical study. For these, we use the first three years of exam performance, to allow comparability.

¹⁹To compute offer rates, we use the data on all 58,890 applicants to Cambridge in 2013-2016 in the relevant subjects; whereas, as described above, the rest of our analysis is based on the subsample of 5,888 admitted to Cambridge in the same years.

²⁰The by-subject offer rate breakdown excludes Biological and Physical Natural Sciences due to data limitations.

Table 1: Gender composition and performance

	N admitted	Percent female	Offer rates		Female score - Male score	
			Females	Males	Year 1	Year 3
	(1)	(2)	(3)	(4)	(5)	(6)
<i>All subjects</i>	5888	36	23	24	-3.5	-0.9
<i>MI subjects</i>						
Economics	547	34	16	16	-2.7	-1.2
Engineering	957	24	23	17	-4.2	-2.2
Mathematics	818	15	38	39	-6.8	-5.7
Physical Sciences	1222	27			-3.6	-1.6
<i>Non-MI subjects</i>						
Biological Sciences	804	59			-1.7	0.1
Law	645	56	25	25	-1.2	-0.9
Medicine	895	45	17	23	-1.4	1.0

Note. *Percent female:* Percentage of female students in total admitted. *Offer rates:* Percentage receiving offer for admission, by gender. *Female score - Male score:* Percentage point difference between the average exam score achieved by females and the average exam score achieved by males. All differences are significant at 1% level except for Economics Year 3 (significant at 5% level) and Biological Natural Sciences Year 3 (insignificant at conventional levels).

points (the average first year exam score is 64.9%). The differences are significantly larger in mathematically intensive subjects than in the rest. Among the former group of subjects, those with a lower proportion of females have a bigger gap between average male and female performance. By the time the students are in their third year, gender differences in performance shrink across all subjects, by an average of 0.9 percentage points (relative to the average third year exam score of 66.5%), though they are still statistically significant except in Biological Natural Sciences and in Medicine, where third year females outperform their male counterparts. It is worth noting that the overall average performance also improves from first to third year in almost all subjects.

4.2 School type

We split our sample into those students who went to UK state-funded schools (46%) and all the rest (54%), the latter group including both students from UK fee-paying schools and non-UK students (32% and 21% of the total, respectively). In line with the UK government

approach, attending a state-funded UK school is taken to represent a lower socioeconomic status. In terms of offer rates, students from state-funded schools do better than the rest, with a probability of 27% of getting an offer compared to 21% for the rest of the students.

In first year exams, students who went to state-funded schools, underperform the rest on average by 1.7 percentage points, with the gap somewhat larger for the mathematically intensive subjects compared to the rest. By the third year, the average difference shrinks to 0.6 percentage points, converging between the two groups of subjects. These gaps are statistically significant throughout.

4.3 Pool

Recall that the key to our empirical approach is the two-tier admission process, where some students are admitted directly (first tier) and others are admitted after being put in the pool (second tier). About 1,215 students, or 20% of our sample were admitted from the pool²¹.

Performance wise, table 2 shows that in the first year, students admitted directly significantly outperform those admitted from the pool. This is true for both, mathematically intensive and other subjects. These differences shrink by the time the students reach their third year, but still remain statistically significant overall and for each group of subjects²². The fact that candidates taken from the pool are weaker is in line with the rationale of the two-tier admission system and confirms that, on average, admissions officers predict performance correctly when allocating the students to these tiers. This is important because it confirms two the central assumption of our empirical model: first that the pool contains candidates who are systematically weaker than directly admitted ones, and second, that the admissions officers predict performance correctly on average.

²¹In 4/5 of these cases, the student is taken by a different college from the one they had applied to, in line with the main rationale of the pool. The rest of the time, they are taken by the college they had applied to. This happens when a college puts a candidate in the pool but then takes them back, for example, if they could not find someone else they would rather take from the pool.

²²The MI subjects start out with a bigger gap in year 1, and it shrinks more over time. By year 3, in both groups of subjects directly admitted candidates outperform pooled candidates by 1 percentage point, statistically significant at 1% level.

Table 2: Two tiers of admission

	Taken from the pool (%)	Exam score differences	
		Direct - Pool	
		Year 1	Year 3
	(1)	(2)	(3)
<i>All subjects</i>	21	1.9***	0.9***
<i>MI subjects</i>	18	2.1***	0.9***
Economics	21	1.2**	0.3
Engineering	19	1.8**	0.9*
Mathematics	15	4.3***	2.0**
Physical Natural Sciences	19	1.6***	0.8
<i>Non-MI subjects</i>	24	1.2***	0.8***
Law	24	1.0***	1.0***
Medicine	22	1.6***	1.4***
Biological Natural Sciences	25	0.9*	0.2

Note. *Direct - Pooled:* Percentage point difference between the average exam score achieved by all directly admitted students and the average exam score achieved by all students taken from the pool, in Year 1 and Year 3 as denoted by column headings. ***, **, and * indicate that the score difference is different from zero at 1, 5 and 10% significance levels, respectively.

5 Results

5.1 Identifying and measuring the trade-off

Recall that our focus is on two key characteristics of applicants: (A) gender and (B) socioeconomic status as proxied by whether they attended a UK state-funded school. For each one of these, there are two categories: for gender, we label them h =female and g =male, and for school type, applicants from UK state-funded schools are labelled h and applicants from other schools g .

For each characteristic, gender and school type, table 3 shows the estimates that allow us to identify whether there is a trade-off between performance and equity, and, if present, its direction (which group is better performing), its lower and upper bounds and the implied

bounds on the weight of the decisionmaker places on equity, $\frac{\beta}{1+\beta}$.

Table 3: Identifying trade-offs

	Δ	Difference in cutoffs		$\min \frac{\beta}{1+\beta}$	$\max \frac{\beta}{1+\beta}$
		Lower bound	Upper bound		
	(1)	(2)	(3)	(4)	(5)
Gender (h = female, g = male)					
non-MI subjects	-0.02 (0.02)	0.02 (0.08)	0.39*** (0.06)		
MI subjects	0.54*** (0.01)	0.25*** (0.06)	0.59*** (0.07)	0.1*** (0.02)	0.2*** (0.02)
School type					
h = UK state, g = all other	0.08*** (0.02)	-0.04 (0.04)	0.37*** (0.05)		

Notes: Standard errors in parenthesis. For columns (4) and (5) they are calculated using bootstrapping as is standard in the partial identification literature, cf. Stoye 2021.

Δ is the difference in the shares of g -types and h -types among students admitted directly (first round).

*, **, and *** significantly different from zero at 10%, 5% and 1% level respectively.

Recall from section 3.3, that we identify a trade-off between performance and equity using the most conservative rule available to us, viz. that there is a trade-off if and only if both the lower and upper bound of the cutoff differences are significantly greater than zero. Hence, we can summarize the conclusions from table 3 as follows:

A. Gender

- **Non-mathematical subjects.** We find no conclusive evidence of a trade-off between gender equality and performance for these subjects. Although the upper bound of cutoff differences is positive and significant (column 3), the lower bound is not significantly different from zero (column 2). At the same time, Δ is not significantly different from zero (column 1), and hence full gender equality is achieved, consistent with the no trade-off outcome in our model (see section 3.1). We reach this conclusion despite the fact that, on average, females slightly underperform males in these subjects, as shown in table 1. This underscores the fact that whether the trade-off exists depends on the performance of the marginal,

rather than average candidates, and conclusions based on simple averages can be misleading. An alternative explanation for the contrast between estimated average and marginal differences is that our lower bound estimate of the latter is below the true value of marginal differences. Regardless, the point remains that our method has detected no evidence to support that the admission cutoff is different across genders or the existence of a performance-equity trade-off in these subjects.

- **Mathematically intensive subjects.** There is a trade-off between gender equity and performance in these subjects, as both the lower and the upper bounds of the differences in cutoffs are significantly above zero (columns 2 and 3). Neither of decisionmaker’s individual objectives are maximized. On the one hand, equity is not achieved, as the difference between the share of males and females among the students is large – 54 percentage points (column 1). On the other hand, marginal performances are not equalized: marginal males outperform marginal females by somewhere between 0.25 of a standard deviation (lower bound in column 2), and 0.59 (upper bound in column 3). For a benchmark standard normal distribution for normalized test-scores, at the 80th percentile (the overall admission success rate is about 20%), this represents about 10% of the total mass of admitted girls. In other words, between 10-20% of the admitted girls got in due to the lower cutoff. These estimates give us the lower and the upper bound of gender equality - performance trade-off faced by the university in mathematically intensive subjects. The decision maker’s implied equity weight $\frac{\beta}{1+\beta}$ is between 10% and 21% (and hence the weight placed on performance is 79-90% respectively; columns 4 and 5).

B. School type: We find no conclusive evidence of trade-off between socioeconomic equity and performance, using school type to proxy for socioeconomic status. Although the upper bound of cutoff differences is positive and significant (column 3), the lower bound is not significantly different from zero (column 2). Note also that equity is not achieved, since Δ is 0.08 and is significantly different from zero, i.e. the percent admitted from UK state-funded schools is 8 percentage points lower than percent admitted from other backgrounds²³.

²³Recall from section 2 that the University of Cambridge faces a government target of ensuring that 62-64% of UK students are from state-funded schools. In our data this fraction is 64%, so the upper limit of the target is exactly met.

Note that in all cases, the upper bound of the cut-off differences is positive which, combined with the sign of the lower bound estimates, implies that the pool system is working as intended, with pooled candidates weaker on average than directly admitted ones. Hence, the key assumption of our cut-off identification model is satisfied (see section 3.2).

To sum up, these results show a trade-off between gender equality and performance in mathematically intensive subjects, but not in other equally competitive ones. They also show that there is no conclusive evidence of a trade-off when it comes to school type.

To examine the robustness of these conclusions, we now move beyond simple mean comparisons. In particular, since identification of the trade-off crucially depends on our estimates of the lower bound of cut-off differences, in the subsequent subsections we estimate this lower bound more carefully, controlling for a number of potential confounders. We also extend the analysis beyond first year exams to see whether our results continue to hold for performance in later years. We first consider the case of gender, and then school type.

5.1.1 Gender

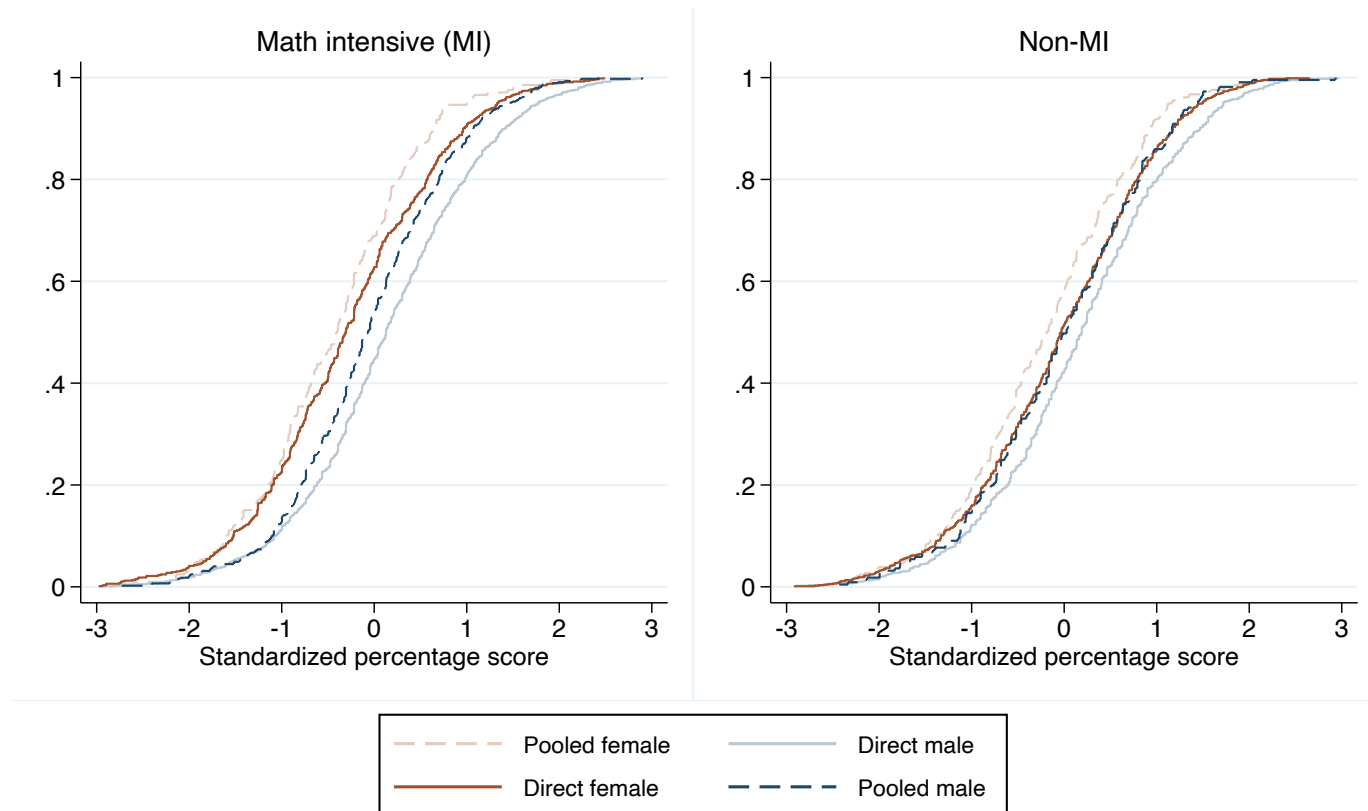
In figures 3, we plot the cumulative distribution function of first year exam percentage scores for four subgroups of admitted students: pooled male, directly admitted male, pooled female and directly admitted female. The panel on the left is mathematically intensive subjects, the panel on the right is the rest.

In mathematically intensive subjects, we see clear evidence that the distribution of scores for directly admitted males first-order stochastically dominates the rest, followed by pooled males, directly admitted females and, finally pooled females. The fact that pooled males have stochastically higher exam scores than directly admitted female entrants *throughout the distribution* suggests that $g_1 - h_1 > 0$, i.e. the de-facto admission cut-off is higher for males, seen in the light of equations (7) and (10) above. In contrast, the distribution of exam scores for pooled females is first-order stochastically dominated by the distributions of both pooled and directly admitted male scores. This confirms our finding that the university faces a gender equality performance trade-off in these subjects. Figure 4 shows that the same pattern also holds in each individual subjects in the mathematically intensive group (Economics, Engineering, Mathematics and Physical Sciences).

In contrast, in non-mathematically intensive subjects (figure 3, right graph), the performance of pooled males and directly admitted females is similar, in line with our earlier finding that there is no gender equality-performance trade-off for these subjects. Again, in figure 5 we see the same pattern in the individual subjects that comprise this group (Biological

Sciences, Law and Medicine).

Figure 3: First-year exam scores by pool status and gender, by subject group.



Note. The graph shows the cumulative distribution function of first year exam percentage scores for different subgroups of students. The functions are plotted separately for MI and non-MI subjects

The graphical evidence above is then corroborated by regression analysis on the combined sample of pooled male and directly admitted female entrants to check robustness of our estimates of the lower bound of equity-performance trade-off. We do this in column 1 of table 4, Panel A, where we regress first year exam performance (standardized across subjects) on a dummy indicating pooled males. The positive, statistically significant coefficient 0.11 on the dummy variable implies that pooled males score an average of 0.11 standard deviations higher than directly admitted female applicants.

There are three key challenges in interpreting this estimate as an indication of de-facto lower admission cutoffs for females. Firstly, males could apply to subjects that are more selective in the first round of admissions than females. We address this by including subject fixed effects (column 2 onwards). Secondly, decisions to admit and pool candidates are made

by individual colleges. This opens up a possibility that the propensity to put students into the pool varies across colleges in a way that is correlated with the overall performance. Indeed, confirming this, figure 6 shows that colleges that put more candidates into the pool tend to be better performing, on average, and so may be receiving higher quality applicants. To address this we include application college fixed effects (column 3), and confirm that our result remains the same when we compare directly admitted females and pooled males who initially applied to the same college.

Finally, the effect could be driven by variations in the quality of teaching provision across colleges, e.g. if colleges that have better teaching and hence better overall performance are more likely to take students from the pool. Offer college fixed effects would control for this; however, since there is a number of large colleges which very rarely take from the pool, including these fixed effects effectively reduces our sample by nearly 1/3, and renders the estimates insignificant (not reported). At the same time, figure 7 demonstrates that this endogeneity concern is not empirically relevant: if anything, college ‘quality’ is inversely correlated with the propensity to take applicants from the pool, and hence cannot explain the positive coefficients we obtain.

Henceforth, we will use the specification with subject fixed effects and application college fixed effects as our preferred one.

Columns (4) and (5) confirm the previous finding that these results are being driven entirely by MI subjects. In these subjects, the first year gap is 0.24 of standard deviation, significant at 1% level. This is the lower bound of the difference in admission cutoffs and confirms our earlier result that using even the most conservative measure, the university faces a trade-off a gender equality-performance trade-off in these subjects. The coefficient gives the lower bound estimate of this trade-off and is the same as we found in mean comparison tests in table 3.

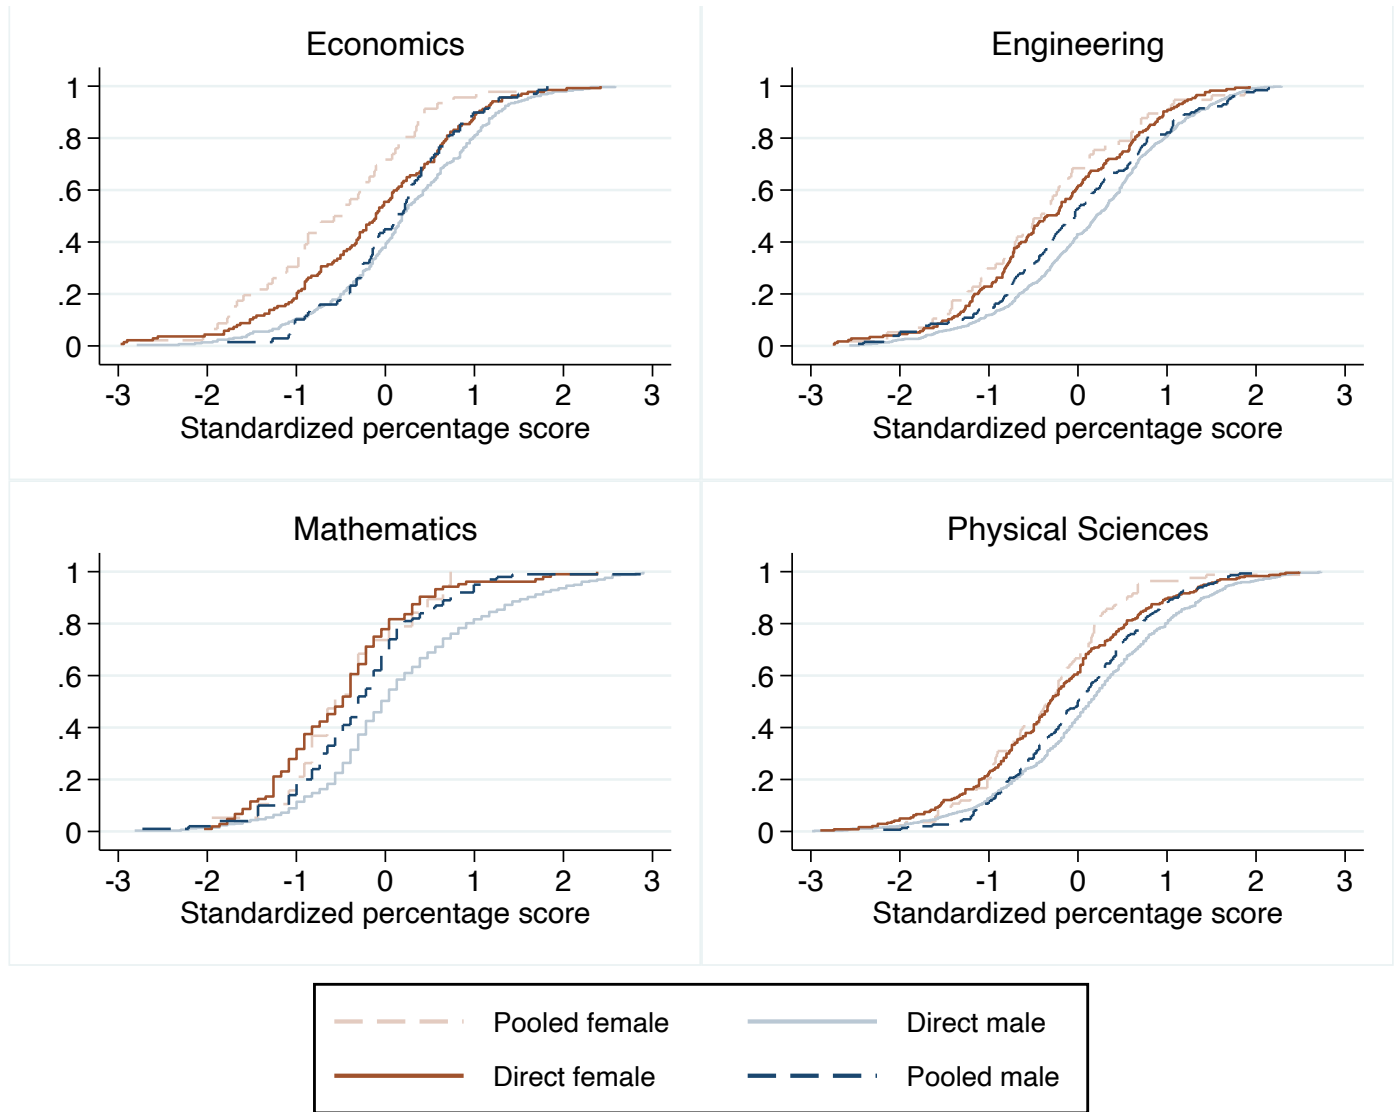
It may, however, be that first year performance of students is not the best measure of their overall performance at university. We therefore, re-estimate our regressions using the students performance in final exams in the remaining two years of their university degree (table 4, Panel B). Focussing on MI subjects only, first, we see that the statistically significant performance gap persists in second and third years. Hence, using this longer run measure of performance, we continue to find existence of a equity-performance trade-off. In magnitude, the gap shrinks, from 0.24 standard deviations in the first year to 0.15 standard deviations in the third year, which means that the implied lower bound of the equity-efficiency trade-off

is smaller when we look at performance in later years.²⁴ This is in line with our earlier observation that the gender gap in average performance falls between the first and the final year.

To summarize, our regression analysis controlling for a number of factors has confirmed our preliminary gender results reported in Table 1. There is no difference in lower bound of admission cutoffs for different sexes in Law, Medicine and Biological sciences, confirming the conclusion that there is no equity-performance trade-off there. However, in mathematically intensive subjects (Mathematics, Engineering, Physical Sciences, and Economics), there is a significant difference in cutoffs, even using the lower bound estimates, thus implying that there is a gender equality-performance trade-off in these subjects. The lower bound of that trade-off signifies that at least 0.15-0.24 of a standard deviation in performance is given up on the margin to maintain the current proportion of female students at 23% in these disciplines.

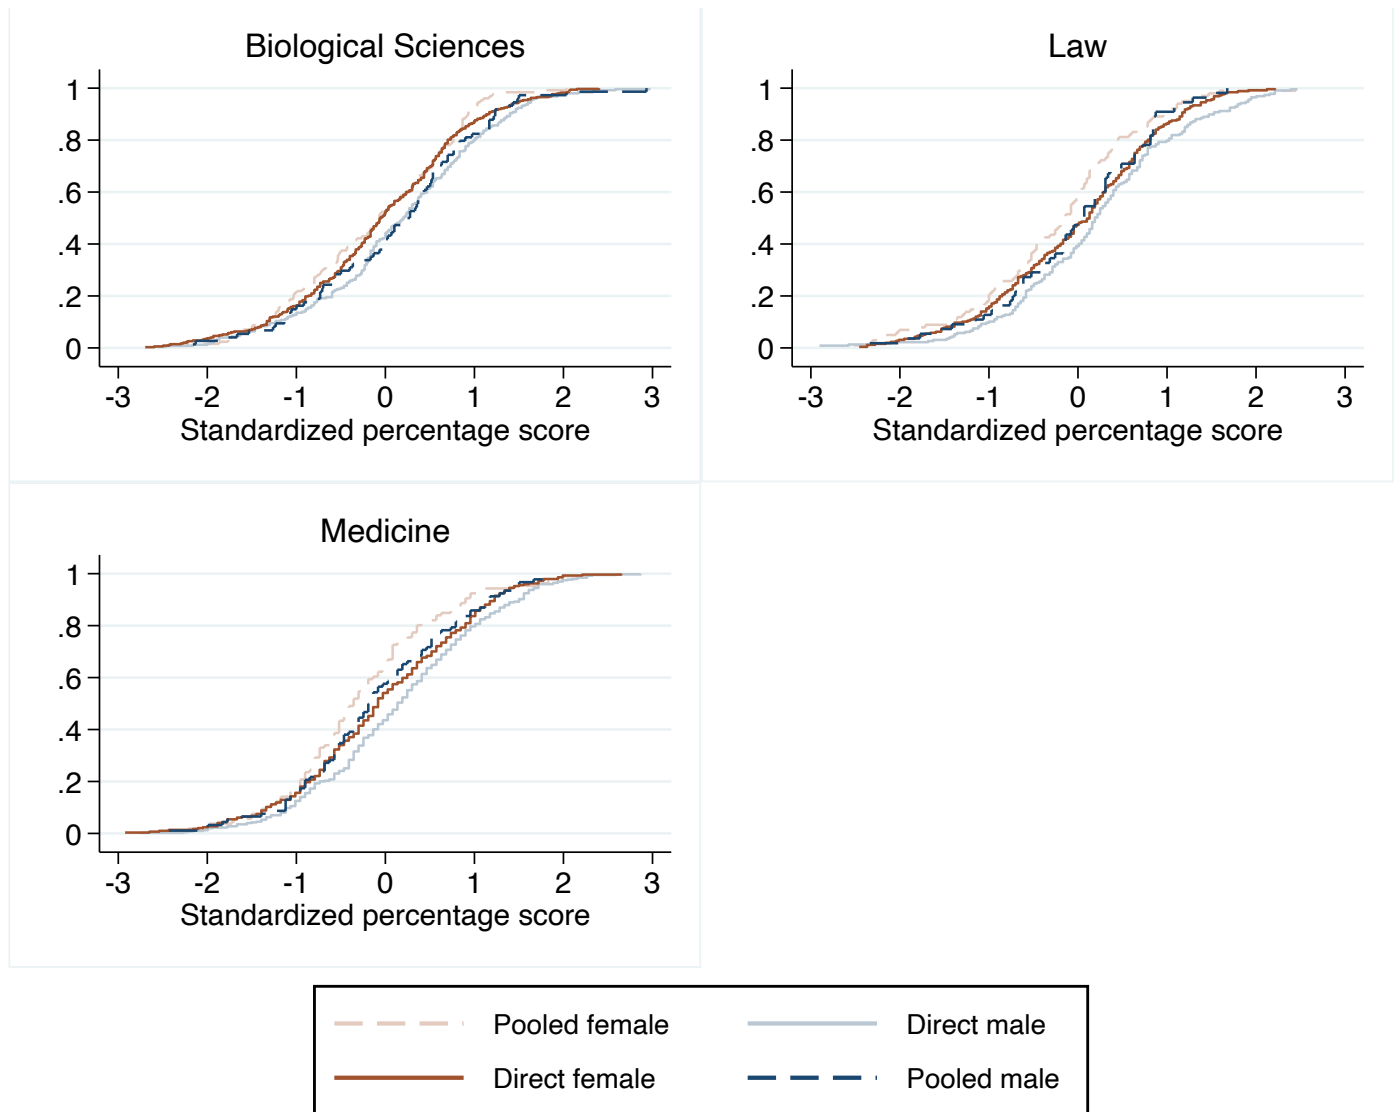
²⁴We also performed the same exercise for non-MI subjects and found that the estimates remain statistically insignificant in later years.

Figure 4: First-year exam scores by pool status and gender, MI subjects.



Note. The graph shows the cumulative distribution function of first year exam percentage scores for different subgroups of students in MI subjects, by subject.

Figure 5: First-year exam scores by pool status and gender, non-MI subjects.



Note. The graph shows the cumulative distribution function of first year exam percentage scores for different subgroups of students in non-MI subjects, by subject.

Figure 6: College performance and contributions to the pool

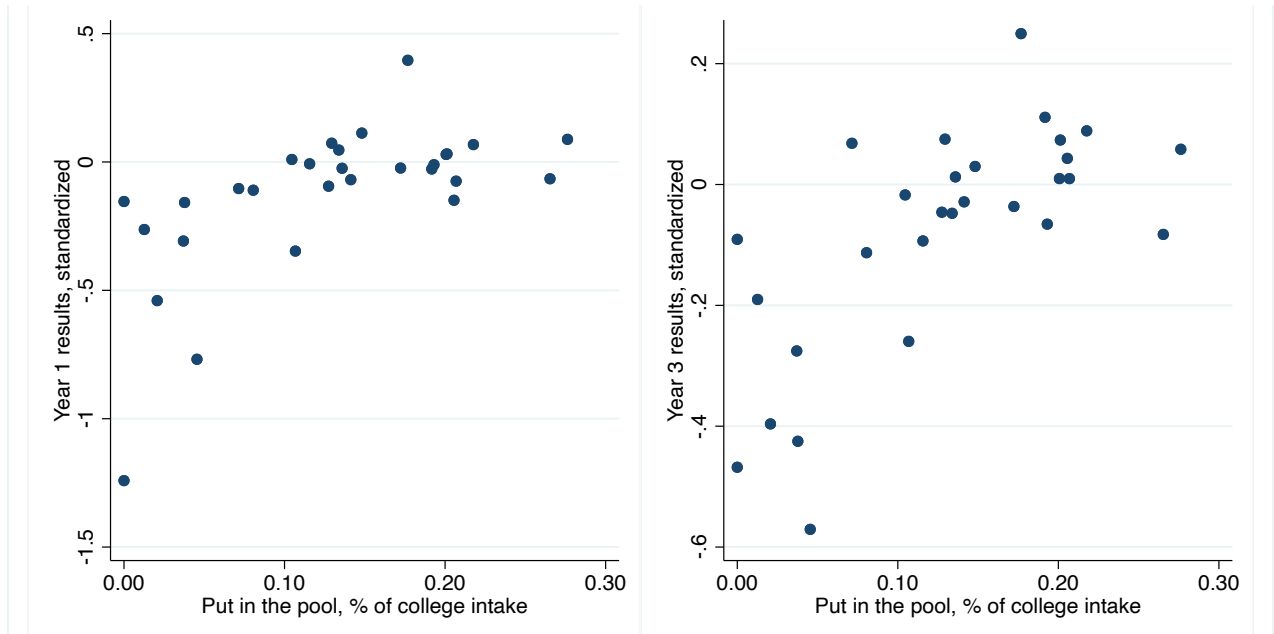


Figure 7: College performance and withdrawals from the pool

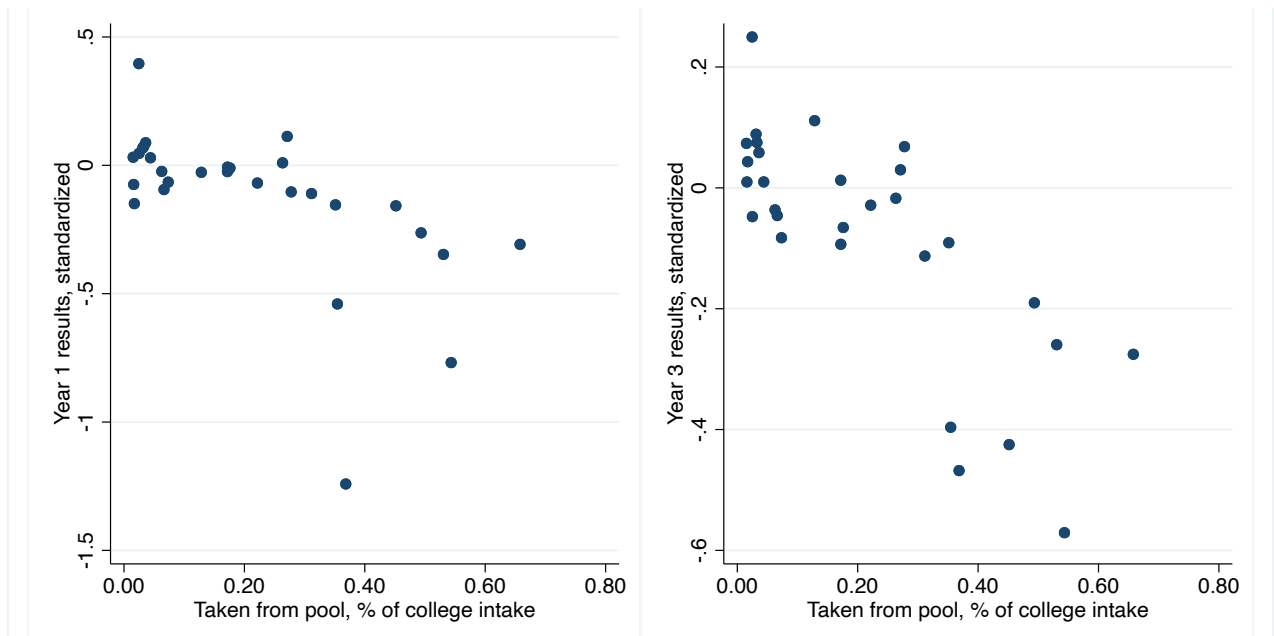


Table 4: The gap between pooled males and directly admitted females

Panel A. First year

	All subjects			MI	Non-MI
	(1)	(2)	(3)	(4)	(5)
Pooled male	0.11** (0.05)	0.17*** (0.06)	0.15*** (0.05)	0.24*** (0.06)	0.02 (0.08)
Observations	2235	2235	2235	1106	1129
Subject FE		✓	✓	✓	✓
Application college FE			✓	✓	✓

Panel B. Later years: Mathematically intensive subjects

	Second year	Third year
	(1)	(2)
Pooled male	0.16*** (0.06)	0.15*** (0.06)
Observations	1106	1106
Subject FE	✓	✓
Application college FE	✓	✓

Note. Each column reports the results from a different OLS regression. *Pooled male:* a dummy variable that equals one for pooled males and zero for directly admitted females. *MI:* Subjects with high mathematical contents: Economics, Engineering, Mathematics and Physical Natural Sciences. *Sample:* Pooled males and directly admitted females. The sets of subjects used to obtain the corresponding results are indicated in the column header in Panel A, and is MI in Panel B. *Dependent variable:* the standardized score obtained in exams, in first year in Panel A, in second and third year in Panel B, indicated by column headers. Standard errors clustered at the application college level (columns 1 and 2 in Panel A) and robust standard errors (the rest of the columns in Panel A and all columns in Panel B) are reported in parentheses. ***, **, and * indicate that the coefficients are different from zero at 1, 5 and 10% significance levels, respectively.

5.1.2 School type

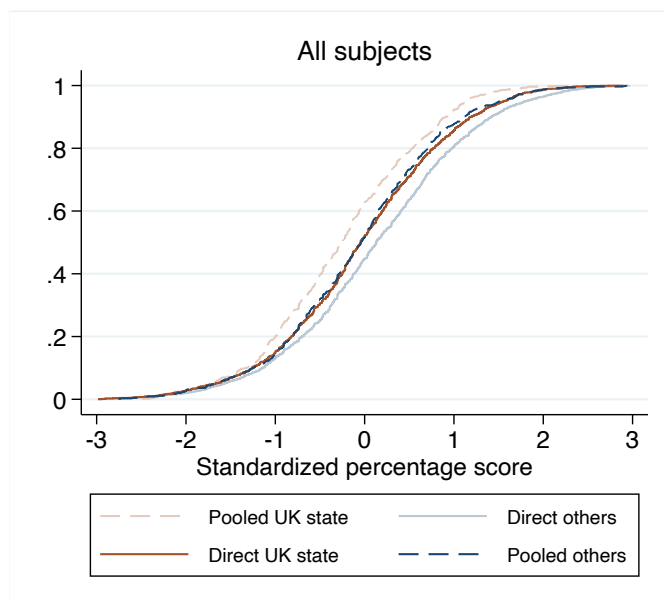
Recall that our category of interest here is whether a candidate went to a state-funded school in the UK (h) or not (g) which, in UK policy debates, is used as a key indicator of socioeconomic background (see section 2).

Plotting the CDFs of first year exam scores for pooled and directly admitted subgroups (figure 8), shows that the directly admitted dominate the pooled, within and across school-types. To further see whether there is a trade-off at admission, we now test for a positive difference between performance of directly admitted candidates from state-funded schools and pooled candidates from other schools.

Table 5 shows that the mean differences are insignificant in all specifications with first year performance (columns 1-3), confirming our estimates in table 3. The differences are negative for second and third year performance (columns 5 and 6). Recall that negative differences do not provide evidence of a trade-off (sections 3.2 and 3.3).

To recap, the fact that there is no positive difference in performance between super-marginal candidates from state-funded schools and marginal candidates from other schools implies that $\min g_1 - h_1$, the lower bound of the equity-performance trade-off, is not significantly different from zero. Hence, based on this conservative measure, we confirm there is no conclusive evidence of a trade-off between equity and performance when it comes to school type of candidates.

Figure 8: First-year exam scores by pool status and school type.



Note. The graph shows the cumulative distribution function of first year exam percentage scores for different subgroups of students, by school type, for all subjects combined.

Table 5: The gap between directly admitted students from UK state-funded schools and pooled students from other schools.

	Year 1			Year 2	Year 3
	(1)	(2)	(3)	(4)	(5)
Pooled others	-0.044 (0.044)	-0.048 (0.046)	-0.067 (0.045)	-0.101** (0.045)	-0.077* (0.045)
Observations	2795	2795	2795	2795	2795
Subject FE		✓	✓	✓	✓
Application college FE				✓	✓

Note. Each column reports the results from a different OLS regression. Sample: directly admitted students from state-funded schools and pooled students from other schools. Dependent variable: indicated in the column header, where *Year 1*, *2*, and *3* stand for standardized scores obtained in first-, second-, and third-year exams, respectively. *Pooled others*: a dummy variable that equals one for pooled candidates from other schools and zero for directly admitted candidates from state-funded schools. Standard errors clustered at the college of application level (Columns (1) – (2)) and robust standard errors (Columns (3) – (6)) are reported in parentheses. ***, **, and * indicate that the coefficients are different from zero at 1, 5 and 10% significance levels, respectively.

5.2 Robustness

We perform several further robustness checks on our results from mathematically intensive subjects, where we find a significant gender gap performance of marginal candidates, even with the lower bound estimates. We report outcomes of these robustness checks in table 6.

- Women-only C=colleges: Cambridge has three colleges that admit only women; these colleges were established later than most others, and tend to have lower financial resources. To check that our gender results are not driven by these, in addition to including application college fixed effects, we drop the women-only colleges from the sample (6, column 1).
- Other background controls: Another possible concern is false attribution, e.g. the gender gap stems from school type if pooled females come mainly from state schools whereas pooled males come mainly from private schools. To check this, we include dummies for school type (UK state funded, other UK, and non-UK) and the student's place of residence (EU, the UK, and other) as additional controls (table 6, column 2)²⁵.
- Year effects: We include fixed effects for year of application to ensure that the scores are comparable across years (table 6, column 3).

Table 6 shows that in MI subjects the gender gap is robust to all three issues discussed above.

Lastly, we consider an alternative way to identify the marginal candidates by focussing on a small number of pooled students who were eventually admitted by the same college that had placed them into the pool. These students could be considered to be the marginal admits, since they had applied to and were not admitted by the pooling college in the first round, but were eventually admitted by them. Column 4 of table 6 shows that, conditional on subject and application college fixed effects, marginal males score 0.21 standard deviations higher in first-year exams than marginal females. This result lends further support to our finding that the admission cut-off for females is lower than that for males. The number of such students is small (242), hence we did not break this up by subject type, but since the overall difference is large, we conjecture that it would be higher still if we considered MI subjects alone.

²⁵The student's place of residence does not always coincide with the school location.

Table 6: Gender: robustness checks, Year 1 performance

	Main	Robustness			
		Mixed	Controls	Year FE	Marginal
	(1)	(2)	(3)	(4)	(5)
Pooled male	0.24*** (0.06)	0.24*** (0.06)	0.24*** (0.06)	0.25*** (0.06)	
Marginal male					0.23* (0.13)
Observations	1106	1052	1106	1106	242
Subject FE	✓	✓	✓	✓	✓
Application college FE	✓	✓	✓	✓	✓
Application year FE				✓	

Note. Each column reports the results from a different OLS regression. Dependent variable: the standardized score obtained in first-year exams. Sample: in all columns except (5), the sample is pooled males and directly admitted females in MI subjects. In column (5), the sample consists of students in all subjects taken back from the pool by the college they applied to. *Marginal male*: a dummy variable that equals one for marginal males and zero for marginal females. *Main*: main specification (table ??, panel A, column 5). *Mixed*: only mixed gender colleges. *Controls*: additional controls for school type and residence (UK, European Union, other). *Application year FE*: with fixed effects for year of application to Cambridge. Robust standard errors are reported in parentheses. ***, **, and * indicate that the coefficients are different from zero at 1, 5 and 10% significance levels, respectively.

5.3 Understanding the gender gap

In this section, we try to further understand the nature of the gender gap in exam results and the implied trade-off between gender equality and performance in mathematically intensive subjects. As we saw in Figure 4, the performance distribution of females admitted in the first tier is first order stochastically dominated not only by the performance of males admitted in the first tier, but also by that of males admitted in the second tier (pool). The underlying factors we investigate can be grouped into two categories: those that are present at university and those in place prior to university admission. These factors may lead to different policy implications, which we briefly discuss this at the end of the section.

We first consider two factors that may influence the gender gap while at university: the

first is the possibility of gender peer effects, and the second the opportunity to select courses in one's degree. We now discuss them in turn, with the results summarized in table 7.

Peer Effects: The gap between pooled males and directly admitted females could be driven by peer effects. In our sample, only 36% of students are females, and the proportion female is systematically lower in subjects where we find the gender gap in performance to be larger. It is conceivable that females are affected negatively by environments with fewer female classmates. To test this, we exploit the fact that education in Cambridge is organized in colleges which create small communities. Hence, we investigate whether the gap in first-year performance is lower in subject-year-college combinations with a higher share of females. We do not find any evidence supporting this hypothesis (column 1 in table 7).

Compulsory and optional courses. Every degree we analyze has some compulsory courses that every one must take and also a range of optional courses. It seems useful to ask whether females do better in one category than another. Although in some degrees the program structure is too complex to test for this, in Economics there is a clear split into optional and compulsory courses in the students' final year. In column 2 of table 7, we use the data from the final year of Economics exams to ask whether the gender gap in the final year performance, found earlier in table 4, is the same across compulsory and optional courses.

Our results show that it is not: pooled males still outperform directly admitted females in the compulsory courses (0.18 standard deviation, significant at 10% level). However, there are no performance differences in optional courses. This is consistent with the idea that once they are allowed to choose, females sort into the courses where they perform better. Since in most degrees there is a lot more choice in later compared to earlier years, this may partly explain our finding that the gap between pooled males and directly admitted females shrinks over time (in table 4). This means that at least to some extent, females are able to reduce the gender gap when they are allowed a choice of courses. This also suggests that the observed gender gap cannot be explained by something endemic to the university environment that disadvantages females across the board in their studies.

Still, this leaves open the question of whether the ability to select allows females to avoid certain courses that are being taught or examined in a way that disadvantages them, or that it enables them to choose courses that are a better match for their pre-admission characteristics. Although we cannot answer this directly, in the next subsection we look at whether pre-admission characteristics, such as school qualifications, can explain some of the gender gap.

Pre-Admission Qualifications: In table 8, we investigate whether the observed gender difference are related to pre-admission qualifications. To do this, we will focus on roughly 80% of our sample who sat British-system school exams²⁶. In this system, two sets of exams marks go onto the students application form to the university. The first are the General exams, known as GCSEs in the UK and IGCSE internationally. They are nationally graded, compulsory board exams for all school pupils in the UK, whether or not they would like to go on to higher education. They are taken when students are 16 years old, so roughly 16 months before a student applies to Cambridge. In this exam, a pupil receives grades A*, A, B,...,U, where A* is the highest and U is the lowest. It is common for universities to use the number of A*s as a summary measure for performance in these exams. Using this measure, we see that the directly admitted females in fact dominate pooled males (table 8, column 1).

The second qualification are the Advanced exams, known as A-levels, and they are taken only by those who plan to continue in higher education. These exams come into two parts because they are spread over two years. The first and earlier part, known as Advanced Subsidiary (AS) levels is the relevant one to our analysis because these exams place before the students apply to university²⁷. These exams are typically taken four-five months before a student applies to Cambridge. Hence, AS-levels are the most recent set of school results available to university admission officers at the time of the candidates' application to Cambridge, and tend to carry significant weight in the admission process²⁸.

Recall that the main gender differences we find are in those subjects where mathematics plays a central role. Hence, it is of particular interest to ask whether we see similar differences in mathematics AS exams taken before the application to Cambridge. We report our results in table 8 panel A, columns 2 and 3.

A key characteristics of the A-level system is that school pupils have some flexibility in the number of AS modules or courses (and hence exams) they choose to enrol in for each subject they study at A-levels. The candidates taking more modules are covering more material and coping with a higher workload. The average estimated number of mathematics modules taken by the candidates our sample is 6, but it varies with most candidates taking between 3 and 9 modules²⁹. In column 2 of table 8 our dependent variable is average standardized

²⁶These are used in the vast majority of schools the UK (except Scotland) and also a number of other countries, primarily former territories of the British Empire, e.g. Singapore, Hong Kong, and India.

²⁷The second part, known as A2, is taken in the summer before starting university, by which time university have already made their offers.

²⁸The A-level system underwent a substantial reform in 2017, which, inter alia, involved the AS-exams. The data used in this paper are pre-reform, and so the description here refers only to the pre-reform arrangements.

²⁹Although we cannot directly observe the number of modules taken by each candidate, our data contain

score per math module, and in column 3 it is the number of math modules taken.

We see that pooled males have a slightly higher average AS mathematics score than directly admitted females, significant at 10% level, but the magnitude of the difference is small (table 8, column 2). They also take significantly more maths modules, approximately 1/2 module more, which compared to the average of 6, is a sizeable difference and is significant at 1% level (table 8, column 3). We do not observe significant differences between our groups of interest when we look at AS-level performance in other subjects, such as social and natural sciences (not reported). The difference in Maths AS is important: as shown in Panel B of the same table, conditional on being admitted, it is positively correlated with subsequent performance in university exams, at 1% level in both first and third year results.

To summarize, there is evidence that females outperform males in the compulsory general exams taken at 16+. At the same time, when it comes to advanced exams taken at 18+, the directly admitted females are found to have taken less maths modules and perform marginally worse on average than the pooled males. These results are consistent with the conclusion that males come to university with stronger preparation in maths, and this is then contributing to the gender performance gap in mathematically intensive subjects in university exams.

Although the gender gap in our data may be partly explained by prior mathematical preparation, there are other possible explanations as well. In particular, our result that the gender gap is only found in mathematically intensive subjects but not in equally competitive disciplines like Law, Medicine and Biology is consistent with the idea that girls, and bright girls in particular, are reluctant to or discouraged from applying to quantitative disciplines. Indeed, in our dataset, females constitute only 25% of applicants in mathematically intensive subjects, whilst accounting for 55% of applicants in the other subjects³⁰.

the maximum points that a candidate can earn in each subject they are taking aggregated across all modules they take. Using the fact that a typical module has 100 points available, we create a proxy for the number of modules taken by each candidate in each subject.

³⁰This excludes Biological and Physical sciences due to data limitations.

Table 7: Understanding gender gap at university

	Peer effects (1)	Compulsory v optional courses (2)
Pooled male	0.26*** (0.12)	0.04 (0.11)
Share female	0.05 (0.27)	
Pooled male \times Share female	-0.04 (0.37)	
Compulsory		-0.02 (0.07)
Pooled male \times Compulsory		0.18* (0.10)
Observations	1106	1028
Subject FE	✓	
Application college FE	✓	✓

Note. Dependent variable: Column 1 the overall standardized score obtained in first-year exams, overall. Unit of observation: student. Column 2 the standardize score obtained in third-year exams for each course taken within the Economics degree. Unit of observation: student-course pair. Sample: Column 1 pooled males and directly admitted females in MI subjects. Column 2 pooled males and directly admitted females in Economics. 2013-2016 cohorts. Standard errors: Robust in column 1; clustered at student level in column 2. *Pooled male*: a dummy variable that equals one for pooled males and zero for directly admitted females. *Share female*: for each student, we calculate the share of females in his/her subject-year-college combination (including that student). *Compulsory* a dummy that equals one if the course is compulsory, 0 otherwise. ***, **, and * indicate that the coefficients are different from zero at 1, 5 and 10% significance levels, respectively.

Table 8: Gender gap: characteristics at admission

<i>Panel A. Gender gap in characteristics at admission</i>			
	General exams	Advanced maths exams	
	A*s	Average score	Number of modules
	(1)	(2)	(3)
Pooled male	−0.37*** (0.07)	0.04* (0.02)	0.47*** (0.15)
Observations	810	850	850
Subject FE	✓	✓	✓
Application college FE	✓	✓	✓

<i>Panel B. University performance and characteristics at admission</i>		
	First year results	Third year results
	(1)	(2)
Advanced maths score	0.13*** (0.02)	0.06*** (0.02)
Number of A* in General exams	0.11*** (0.02)	0.15*** (0.02)
Observations	2504	2504
Subject FE	✓	✓
Application college FE	✓	✓

Note. The multicolumn headers give groups of dependent variables. Robust standard errors are reported in parentheses. Panel A notes: *Pooled male*: a dummy variable that equals one for pooled males and zero for directly admitted females. *Sample*: (1) Pooled males and directly admitted females enrolled in MI subjects and who have (I)GCSE scores. MI are subjects with high mathematical contents: Economics, Engineering, Mathematics and Physical Natural Sciences. (2) Pooled males and directly admitted females enrolled in MI subjects, who take at least three Advanced subsidiary (AS) modules. *Dependent variables*: (1) standardized A* count obtained in General ((I)GCSE) exams; (2) standardized exam score obtained in Advanced subsidiary (AS) maths modules, average across modules; (3) number of AS mathematics modules taken by the candidate. Panel B notes: *Sample*: All students in MI subjects for whom AS and GCSE exam information is available. *Dependent variable*: the standardized score obtained in exams, column 1 in first year and column 2 in third year. *Regressors*: AS maths score – total maths score obtained across all AS maths modules; Number of A* in General exams – number of A* achieved in (I)GCSE exams. ***, **, and * indicate that the coefficients are different from zero at 1, 5 and 10% significance levels, respectively.

6 Conclusion

Admission practices at elite universities face the dual objectives of maintaining high academic standards while admitting sufficiently many students from under-represented demographic groups. To make informed policy decisions it is important to understand whether there exist trade-offs between these objectives, and if so, what is their realized magnitude.

In this paper, we develop a simple intuitive model of optimal admissions for a university that values both equity and post-entry academic performance, and receives applications from more students than it can admit. We show that in this model, the university’s implied preference for equality vis-a-vis expected future performance is a closed-form function of the difference in admission-rates and the difference in performance between marginal admits from different population subgroups at the optimum. This result is nonparametric in that it does not require assuming any specific functional-form for distribution of unobservables, e.g. that underlying ability is normally distributed.

A challenge in implementing the above idea empirically is that it is usually not possible to determine who the marginal entrants are, particularly when admission decisions are based on many indicators, some of which are unobserved by researchers, leading to well-known infra-marginality problems.

We develop a novel method of addressing this issue by exploiting the fact that many students enter elite universities, including Cambridge from where our data come, through a waitlist or second round clearing. We show that using inter-group performance difference between students admitted directly versus through second round of clearing, we can construct lower and upper bounds on the unobserved difference in the admission cutoffs, which in turn equal the performance of the marginal entrants. These bounds (a) do not require information on rejected candidates and (b) remain valid even if some applicant characteristics viewed by admission officers are unobserved by the researcher. This in turn allows us to identify whether the university faces an equity-performance trade-off, its direction and its magnitude at the margin. Furthermore, this allows us to back out the implied relative weights on equity and performance in the university’s objective function. Given that many, if not most, institutions in the world use waitlists and/or clearing to fill all their positions, this approach is likely to have wider applicability beyond the specific context studied here.

Finally, we apply our method to data from the University of Cambridge, one of the most selective and prestigious higher education institutions in the world. We concentrate on two key applicant characteristics that policymakers typically focus on: gender and socioeconomic background as proxied by the school attended by the applicant.

We find strong evidence that in mathematically intensive subjects where female enrolment is relatively low, there exists a significant trade-off between gender equality and performance. According to our most conservative estimates, the university is giving up at least 0.15-0.24 standard deviations in terms of exam performance to prevent the gender ratio from deteriorating further in these subjects. For a standard normal benchmark, this suggests that between 6-15% of admitted girls got in due to a lower cut-off. Our measured performance gap is resilient to a large variety of robustness checks, persists throughout the length of the degree, especially in the compulsory core papers, and indicates a genuine underlying regularity. When interpreted via our formal model of equity-performance trade-off, we find that the weight on the equity objective has to be between 10-20% and that on academic performance the remaining 80-90% respectively, to rationalize the gap we observe.

In contrast, we detect no gender-equity versus performance trade-off in non-mathematical but equally competitive subjects of Law, Medicine and Biological Sciences, where the gender ratio among admits is close to 50-50.

We also investigate the equity-performance trade-off when it comes to socioeconomic background, as captured by the key variable on which the UK government has long based its admission guidelines, viz. the type of school attended by the candidate. We do not find strong evidence of trade-off between equity and performance in this case.

These findings, which show that the presence and magnitude of the trade-off is context-dependent and stand somewhat in contrast to the current policy debate in the UK which focuses almost exclusively on the merits of increasing intake from state-funded schools, with relatively little attention to gender performance gap and recruitment into STEM fields. In contrast, our results show that in the latter area the university faces much larger trade-offs, and so it should command more, not less, policy discussion.

To better understand the nature and causes of the gender performance gap, we run a few further checks. Broadly speaking, the performance gap at university (a) may be due to pre-existing differences in student preparation and/or (b) may emerge as a result of different impact of the university environment on the two genders, and these two alternative explanations have different policy implications. Note that this duality is inevitable for any *outcome-based* test, where it is difficult to disentangle these two sources of ex-post differences. Indeed, our evidence, which is admittedly more suggestive than conclusive, points to a gap existing already before entry into the university. Although females outperform males in general school exams, when it comes to mathematics preparation at a more advanced level, critical for mathematically intensive degrees at university, they lag somewhat behind

their male counterparts. Based on this, we conjecture that in the short-run, boosting the mathematical preparation for female students admitted to STEM fields and Economics, say through a pre-course preparatory course, can potentially improve their subsequent performance and reduce the existing gender gap. Early childhood intervention encouraging bright girls to pursue mathematical tracks (cf. Heckman and Krueger 2005, Heckman 2006, Ellis et al 2016, Wang and Degol 2017) would be a natural longer-term goal for increasing STEM opportunities and participation for women. Indeed, our results are consistent with the suggestion that girls are discouraged from applying to mathematically intensive subjects, with the stronger females going instead into disciplines like Law and Medicine (see Tonin and Wahba 2014 and Crawford et al 2018 for the specific case of Economics).

On the methodological end, our model of an admissions process that values both equity and efficiency, and implies closed-form expressions for their implicit relative weights in terms of nonparametrically estimable reduced-form parameters, gives our analysis a structural flavour. Secondly, we outline a novel method of set-identifying expected performance differences between marginal admits from different groups – an ingredient of our relative weights calculation – by using candidates admitted via waitlist, a common feature of college-admission around the world. These methods serve to mitigate two common complications related to empirical assessment of selection procedures in universities and other organizations, viz. the unobservability of all the factors affecting selection decisions, and the lack of counterfactual outcomes for rejected applicants. When taken to data from Cambridge, our empirical results provide quantitative evidence on the existence of a trade-off between gender diversity and academic performance in some but not all subjects. Establishing whether such trade-offs exist and measuring their size is important for having an informed discussion about policy when the societal objective is a combination of both diversity and efficiency.

Appendix

Here, we solve the policy maker's problem from section ?? assuming an alternative equity objective, to equalize applicant success rates (rather than the number of admitted students) between the two groups g and h . This may arguably be a more reasonable objective to pursue if the number of applicants from the two groups is very different – as is the case with females in mathematically intensive subjects.

The university's optimization problem can now be formally stated as

$$\max_{g_1, h_1} \left[\begin{array}{l} \pi_g E(A_g \times 1\{A_g > g_1\}) + \pi_h E(A_h \times 1\{A_h > h_1\}) \\ -\beta (\Pr(A_g > g_1) - \Pr(A_h > h_1))^2 \\ \text{subject to} \\ \pi_g \Pr(A_g > g_1) + \pi_h \Pr(A_h > h_1) = \frac{M}{N} \end{array} \right] \quad (11)$$

where, as before, β denotes the social welfare value of performance sacrificed per unit of social welfare gained due to a marginal reduction in inequality.

Now, (11) is equivalent to

$$\begin{aligned} \max_{g_1, h_1} & \left[\pi_g \int_{g_1}^{\infty} a f_g(a) da + \pi_h \int_{h_1}^{\infty} a f_h(a) da - \beta (F_g(g_1) - F_h(h_1))^2 \right] \\ \text{s.t. } & \pi_g (1 - F_g(g_1)) + \pi_h (1 - F_h(h_1)) = \frac{M}{N}. \end{aligned} \quad (12)$$

First-order conditions: The first-order conditions for (12) are

$$\begin{aligned} -\pi_g g_1 f_g(g_1) - 2\beta \times (F_g(g_1) - F_h(h_1)) \times f_g(g_1) &= \lambda \pi_g f_g(g_1) \\ -\pi_h h_1 f_h(h_1) + 2\beta \times (F_g(g_1) - F_h(h_1)) \times f_h(h_1) &= \lambda \pi_h f_h(h_1) \end{aligned}$$

where λ denotes the Lagrange multiplier, implying

$$\frac{g_1 + 2\frac{\beta}{\pi_g} (F_g(g_1) - F_h(h_1))}{h_1 - 2\frac{\beta}{\pi_h} (F_g(g_1) - F_h(h_1))} = 1,$$

which via simplification yields

$$\beta = \frac{(g_1 - h_1) \times \pi_g (1 - \pi_g)}{\underbrace{2(1 - F_g(g_1) - (1 - F_h(h_1)))}_{\text{Difference in admission success-rates}}}.$$

Bounds on $(g_1 - h_1)$ obtained in (7) and (10) above translate into bounds for β , since π 's and success-rates are directly estimable from the data. Expression (??) also means that for positive β , we must have that the difference in cutoffs and the difference in success rates must be of the same sign. That is, g -types would have a higher admission success rate if

and only if the marginal (direct) admit of g -type has a higher expected performance. The point-estimates on the bounds for the implied weight on equity, $\frac{\beta}{1+\beta}$, follow from (8) and (10) and confidence intervals follow from calculations standard in the partial identification literature, cf. Stoye 2021.

References

- [1] Altonji, J.G. and Blank, R.M., 1999. Race and gender in the labor market. Handbook of labor economics, 3, pp.3143-3259.
- [2] Arcidiacono, P., Lovenheim, M. and Zhu, M., 2015. Affirmative action in undergraduate education. Annual Review of Economics, 7(1), pp.487-518.
- [3] Arcidiacono, P. and Lovenheim, M., 2016. Affirmative action and the quality-fit trade-off. Journal of Economic Literature, 54(1), pp.3-51.
- [4] Arcidiacono, P., Aucejo, E.M. and Hotz, V.J., 2016. University differences in the graduation of minorities in STEM fields: Evidence from California. American Economic Review, 106(3), pp.525-62.
- [5] Arcidiacono, P., Kinsler, J. and Ransom, T., 2022. What the Students for Fair Admissions Cases Reveal about Racial Preferences (No. 15240). Institute of Labor Economics (IZA).
- [6] Arnold, D., Dobbie, W. and Yang, C.S., 2018. Racial bias in bail decisions. The Quarterly Journal of Economics, 133(4), pp.1885-1932.
- [7] Arrow, Kenneth J. 1973. The Theory of Discrimination., In Orley Ashenfelter and Albert Rees, Discrimination in Labor Markets. Princeton, N.J.: Princeton University Press, 1973. 3–33.
- [8] Autor, D.H. and Scarborough, D., 2008. Does job testing harm minority workers? Evidence from retail establishments. The Quarterly Journal of Economics, 123(1), pp.219-277.
- [9] Ayres, I. and Brooks, R., 2004. Does affirmative action reduce the number of black lawyers. Stan. L. Rev., 57, p.1807.
- [10] Becker, G.S., 1957. The economics of discrimination. University of Chicago Press.

- [11] Benson, A., Li, D. and Shue, K., 2019. Promotions and the peter principle. *The Quarterly Journal of Economics*, 134(4), pp.2085-2134.
- [12] Bhattacharya, D., Kanaya, S. and Stevens, M., 2017. Are university admissions academically fair?. *Review of Economics and Statistics*, 99(3), pp.449-464.
- [13] Boliver, V., 2013. How fair is access to more prestigious UK universities?. *The British Journal of Sociology*, 64(2), pp.344-364.
- [14] Crawford, C., Davies, N. M. and Smith, S. 2018. Why do so few women study economics? Evidence from England. Royal Economic Society working paper.
- [15] Cummings, Jack, 2015. Girls ‘lack confidence’ for Oxbridge. *The Tab*. <https://thetab.com/2015/09/25/girls-lack-the-confidence-for-oxbridge-54920>
- [16] Ellis, J., Fosdick, B.K. and Rasmussen, C., 2016. Women 1.5 times more likely to leave STEM pipeline after calculus compared to men: Lack of mathematical confidence a potential culprit. *PloS one*, 11(7), p.e0157447
- [17] Fryer Jr, R.G. and Loury, G.C., 2005. Affirmative action and its mythology. *Journal of Economic Perspectives*, 19(3), pp.147-162.
- [18] Heckman, J.J., 2006. Skill formation and the economics of investing in disadvantaged children. *Science*, 312(5782), pp.1900-1902.
- [19] Heckman, J.J. and Krueger, A.B., 2005. *Inequality in America: What role for human capital policies?*. MIT Press Books.
- [20] Hoffman, M., Kahn L. B., and D. Lee, 2018. Discretion in Hiring. *Quarterly Journal of Economics*, 133(2), pp.765-800.
- [21] Keith, S.N., Bell, R.M., Swanson, A.G. and Williams, A.P., 1985. Effects of affirmative action in medical schools: a study of the class of 1975. *New England Journal of Medicine*, 313(24), pp.1519-1525.
- [22] Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J. and Mullainathan, S., 2018a. Human decisions and machine predictions. *The Quarterly Journal of Economics*, 133(1), pp.237-293.
- [23] Kleinberg, J., Ludwig, J., Mullainathan, S., and A. Rambachan 2018b. Algorithmic Fairness. *AEA Papers and Proceedings*, 108, pp.22-27.

- [24] Laband, D.N. and Piette, M.J., 1994. Favoritism versus search for good papers: Empirical evidence regarding the behavior of journal editors. *Journal of Political Economy*, 102(1), pp.194-203.
- [25] Li, D., 2017. Expertise versus Bias in Evaluation: Evidence from the NIH. *American Economic Journal: Applied Economics*, 9(2), pp.60-92.
- [26] Li, D., Raymond, L.R. and Bergman, P., 2020. Hiring as exploration (No. w27736). National Bureau of Economic Research.
- [27] Massey, C. and Thaler, R.H., 2013. The loser's curse: Decision making and market efficiency in the National Football League draft. *Management Science*, 59(7), pp.1479-1495.
- [28] Sander, R.H., 2004. A systemic analysis of affirmative action in American law schools. *Stan. L. Rev.*, 57, p.367.
- [29] Sibieta, Luke 2021. The growing gap between state school and private school spending. Institute for Fiscal Studies Observation, 08 October 2021 <https://ifs.org.uk/publications/15672>
- [30] Stoye, J. 2021. A Simple, Short, but Never-Empty Confidence Interval for Partially Identified Parameters. <https://arxiv.org/abs/2010.10484>
- [31] Tonin, M. and J. Wahba 2014. The Sources of the Gender Gap in Economics Enrolment. IZA Discussion Paper No. 8414
- [32] University of Cambridge 2019. Access Agreement with the Office for Fair Access (OFFA) 2018-19.
- [33] University of Cambridge 2020. Access and Participation plan.
- [34] University of Cambridge 2022a. Course list. <https://www.undergraduate.study.cam.ac.uk/courses/>
- [35] University of Cambridge 2022b. Natural Sciences. Course outline. <https://www.undergraduate.study.cam.ac.uk/courses/natural-sciences#course-outline>
- [36] Wang, Ming-Te, and Jessica L. Degol 2017. Gender gap in science, technology, engineering, and mathematics (STEM): Current knowledge, implications for practice, policy, and future directions. *Educational psychology review* 29, no. 1: 119-140.

- [37] Zimdars, A., 2010. Fairness and undergraduate admission: a qualitative exploration of admissions choices at the University of Oxford. *Oxford Review of Education*, 36(3), pp.307-323.
- [38] Zinovyeva, N. and Bagues, M., 2015. The role of connections in academic promotions. *American Economic Journal: Applied Economics*, 7(2), pp.264-92.