Economics Working Paper Series

2020/011

School choice, admission, and equity of access: Comparing the relative access to good schools in England

Ian Walker and Matthew Weldon

The Department of Economics
Lancaster University Management School
Lancaster LA1 4YX
UK

LUMS home page: http://www.lancaster.ac.uk/lums/

Lancaster University
Management School

School choice, admission, and equity of access

Comparing the relative access to good schools in England

Ian Walker and Matthew Weldon
Lancaster University
Department of Economics

Acknowledgement: This work has been funded by the Nuffield Foundation. The Nuffield Foundation is an independent charitable trust with a mission to advance social well-being. It funds research that informs social policy, primarily in Education, Welfare, and Justice. It also funds student programmes that provide opportunities for young people to develop skills in quantitative and scientific methods. The Nuffield Foundation is the founder and co-funder of the Nuffield Council on Bioethics and the Ada Lovelace Institute. The Foundation has funded this project, but the views expressed are those of the authors and not necessarily the Foundation. Visit www.nuffieldfoundation.org

Nuffield Foundation

1 Introduction

This report summarises much of our research project, funded by the Nuffield Foundation, that has been concerned with secondary school choice in England – a task that parents of year 6 primary school do every year to seek admission to publicly funded secondary schools for their 10 year old child.

The essence of our analysis is to explore the determinants of parental preferences from observing the preference lists of parents – the list of schools that they would like their child to be considered for admission. The rank-order of schools that parents list is their way of expressing their preferences for schools. Parents might be likely to have idiosyncratic reasons for choosing schools but it is also likely that all parents will attach some weight to the quality of the school. On the other hand, it is likely that parents would prefer not to subject their child to long travel times. The list for each child provides admission authorities with information about which child would like to go to which school and their characteristics relevant to admission criteria. Schools have some control over their admission criteria (although only grammar schools are allowed to prioritise children by ability) and the admission authorities then have the job of matching children to schools and based on *only* the information provided by parents relevant to admission. In principle, any child can list any school but one admission criterion is invariably a geographical one - it is common to use proximity as the criteria that breaks "ties" in the higher eligibility criteria. For example, schools might give priority to SEN children, and to "looked-after children" (LAC are in the care of local authorities. usually through foster parents), and then to the year 6 children who have older siblings in the particular school. Any remaining places are usually awarded according to the proximity order of applicants until the capacity of the school is full. Parents are limited in the number of schools they can list, depending on Local Authority – many are limited to only three, while the most is six. An important issue for our research is that the finite nature of the list discourages parents from taking big risks in the schools they list – by listing schools that one's child is unlikely to be admitted to. In particular, parents when faced with limited choice will be strategic in how they list schools. That is, parents have an incentive to think about what other parents are choosing to ensure that they do not "waste" a choice on a school that will be full of children with greater eligibility.

The project uses the National School Preferences (NSP) data that contains the lists of all parents. This is linked to detailed records on pupil and school characteristics from the National Pupil Database (NPD), that together allow us to investigate school choices and admission decisions, and heterogeneity in these, across types of parents. England is a good laboratory for research into school choice because the money follows the student, there is little admission by ability, and there is a national database that records the education history and attainments of all pupils. We have detailed data on preferences for two cohorts. We know from this data that only 65% of parents list more than one choice; only 27% make as many choices as they can; only 39% put their local school top; and only 55% include their local school as one of their choices.

The main questions we address are: to what extent do different demographic groups, and so different localities, face real inequalities in choice after accounting for heterogeneity in preferences and strategies; and, in what ways do choices respond to and mitigate inequalities of access?

We answer these questions, first using detailed parental ranked preferences data to examine the descriptive evidence of variation in decision-making and choice strategies by demographic group; and second, by statistically modelling parental preferences and schools' admissions decisions, to derive measures of the quality of schools available to parents.

The balance of evidence on the *causal* effects of the quality of school attended on educational outcomes, suggests that getting into a "good" school matters. Indeed, we find that it matters over and above any effect on educational outcomes (see Gorman and Walker, 2020). Data from the Longitudinal Study of Young People in England (LSYPE) includes responses from parents about how satisfied they are with their child's school. Among parents whose children missed out on their preferred school 27% were very satisfied. Among those that were admitted to their preferred school 43% were very satisfied. Satisfaction comes from many sources – not just the educational outcomes. Parents, possibly more than anything, value the fact that their child likes being at school, feels safe at school, and has good friends at school.

There is also widespread interest in the suggestion that allowing parents a degree of choice over which school to send their children to, may raise the quality of all schools through a competitive mechanism. Schools will strive to be better, not least because the money will follow the students who come. While, there is not a lot of evidence that directly addresses this suggestion, nonetheless school choice has been a major theme associated with improving school outcomes …. "the tide that lifts all boats" (Hoxby, 2003). At a minimum this requires that parents care about school quality, and that schools care about attracting students.

In principle, parents have always had a choice of school: until the late 1980's in England, parents could exercise choice by choosing where to live, or choosing to opt out of the state sector and send their children to private schools. School choice policies, such as those introduced from 1988 in England, are really best viewed as being concerned with changing the costs of the choices that are available. When looked at in this light, we can think about school choice systems not in terms of whether or not they facilitate a particular choice, but as a progressive measure that reflects the costs associated with many possible choices. Seen in this way, it is possible to evaluate each individual aspect of a school choice system in terms of the extent to which it does, or does not, reduce the costs of choice, and for whom.

However, measuring school quality is difficult because it is both multidimensional and subjective. Parents value the consumption benefits of a safe and happy environment that is engaging for their children - that provides not just child-care but the satisfaction that their children's developmental needs are being provided for. They also value the investment benefits of the promise of better outcomes – in terms of not just a better job but, more generally, equipping their children to have a more fulfilling life. Perhaps unsurprisingly, a simple operational definition is hard to come by.

We focus here on two main proximate determinants of parental preferences: the distance from home to school, and the quality of outcomes. School quality is typically measured by researchers using test score outcomes – although policymakers, teachers, and parents, also attach importance to such outcomes. Distance is usually measured as a crow-flies distance by school admission administrators seeking to break ties in oversubscribed schools. The importance that parents attach to proximity likely comes from the fact that it proxies for a variety of characteristics – such as being part of a familiar home neighbourhood near to wider family and extended friendship groups, as well as the low cost and high convenience attached to close proximity. Valuing proximity would be consistent with the family already having chosen to live in the neighbourhood. The extent to which parents choose schools close to the residential neighbourhood they have already chosen to live in, is revealing of parents' attitudes towards the social mix of their neighbourhoods[1]. Proximity brings with it convenience – that it allows the child to more easily take advantage of after-school activities, and saves the time, trouble and costs associated with a long school commute. Since proximity is also commonly used as an admission criteria, the ease with which it can be measured is important – and crow-flies distance is likely to be least controversial.

Relying on test scores as a measure of school quality is more problematic. Quality is likely to be driven by the "match" between child and school, which will inevitability be idiosyncratic. But it is very likely that parents may not be good at imagining how their own child might develop in one school relative to another. What is a good match for one child may not be a good match for another. Looking at absolute quality such as average KS4 test achievements does not control for the quality of the inputs. Given that school funding is increasingly formula driven there is, in practice, little variation in financial inputs – apart from the additional funding for low income children, the Pupil Premium. The education economics literature shows large school fixed effects in the determination of school outcomes, suggesting that schools matter to outcomes – through greater effectiveness in their teaching, for example. The same might be said of individual teachers. However, there is little literature on teacher level fixed effects for England, and there is relatively little research on which to base a consensus view over what the characteristics of a good teacher are. Children also learn from other children – there are peer effects within schools. Moreover, prior achievement, measured say by KS2 tests, is also an input into future achievement.

The UK government has implemented value-added measures to measure the outcomes children achieve, adjusting for differences in starting-points. However, even these are imperfect predictors of one's own child's success: they are measured with a lag of several years, and school resources, staffing, and teaching practices may have changed in the interim; they are noisy measures with substantial year-on-year variation for a single school; and they provide a single average measure of performance, rather than a tailored prediction for a child at a given starting point. In addition to these problems, any measure of school performance, raw or adjusted, can only be as good as the pupil assessments that it is based on – assessments may be vulnerable to perverse incentives, teaching-to-the-test, and grade inflation.

In addition, the allocation of children to schools depends not only on who wants to attend which school, but also on school capacities. The *mechanism* by which places are rationed matters too. The school choice mechanism embodies the criteria used to rank child applicants. Varieties of the "Deferred Acceptance" (DA) school choice mechanism are now commonplace because of the proposition that they incentivise parents to report their *true* preferences to the school choice administrators (unlike the "First Preference First" (FPF) mechanism which favoured first preferences). Missing out on one's top choice under FPF implied that one was also likely to miss out on one's second choice. Thus, the consequences

[1] This is to characterise the choice of residence and of school as sequential, whereas common sense suggests parents consider schools as part of their choice of where to live. This problem is discussed in section 5.2 below.

of missing out were likely to be larger than under FPF than under DA. [2] We provide more detail on DA and FPF mechanisms in Section 3.

DA is thought to encourage parents to report their preferences *without consideration of the likelihood of admission*. This theoretical proposition is, however, contingent of parents being able to rank all possible schools. Limits on this would give parents an incentive to choose their listed schools is a "strategic" manner – by which we mean that parents would need to second guess what other parents where choosing to list. In practice, the allowable number of choices that one might list is usually quite prescribed (as few as 3 in many cases and never more than 5, except in major metropolises where 6 is common) and a degree of strategic decision-making might therefore be incentivised.

Records of parents' school preference lists in England, made available by the UK Department for Education, reveal that a majority of parents do not plump for their nearest secondary school, but the most popular state schools can draw applications numbering many multiples of their year seven entry quotas. When we look more closely at the data on families' preferences, and the schools that children ultimately get allocated to, we find surprising levels of variation in both the preferences people express and their chances of admission to their preferred schools. We see quite varying choices and outcomes for different demographic groups - by income, ethnicity, and prior attainment. We also see variation by local school market and by region. Perhaps most strikingly, we uncover what appear to be substantial inequalities in access to chosen schools, for minority ethnic families, when compared to white families.

However, given that admissions oversubscription criteria are tightly circumscribed by regulations intended to protect children from discrimination, is it likely that the observed patterns reflect real inequalities of access to good schools? Or, should alternative explanations draw on differences in preferences and choice strategies to explain the gap? This project uses the national school preferences data, linked to detailed records on pupil and school characteristics from the National Pupil Database (NPD), to investigate heterogeneity in choices and admissions. The questions we address are: to what extent do different demographic groups and different localities face real inequalities in choice, **after** accounting for heterogeneity in preferences and strategies; and in what ways do choices respond to and mitigate inequalities of access?

There is a burgeoning literature on school choice in the English context, but in fact there are only a handful of papers that specifically look at variation in access to schools for different groups in different locations, and there are no other papers, as far as we know, that address strategic choices. Allen et al. (2014) analyse national administrative data to show that parents

[2] Prior to 2006 secondary schools in England were able to use the First Preference First (FPF) admission mechanism – some did, while others used the DA mechanism that is now obligatory. Gorman and Walker (2020) consider the role of the school admission mechanism by comparing areas of the country that used DA with areas that used FPF. Using distance-matching to weight their observations (in the LSYPE data) they find that missing out one's preferred school had much larger detrimental effects on outcomes in an FPF area than it would in a DA area. Missing out under FPF reduced the chance of gaining 5 GCSE A*-C grades by a statistically significant and economically meaningful 11% compared to not missing out, while under DA it was a statistically insignificant 2%. Staying on post-16 was greatly reduced, and there were long term impacts on earnings (at age 25) of 6% under FPF, again statistically significant and economically meaningful, while it was just 1% under DA. There were important adverse effects of missing out on one's first choice under DA on mental health at age 25 but this was only half of the size of the mental health effect under FPF.

do not always choose the highest-performing school that is available to them. Weldon (2018) documents the ethnic gap in chances of admission to one's first choice school. Recent work by the Education Policy Institute (Andrews and Perera 2017; Hunt 2018b, 2019a, 2019c) has analysed the administrative preferences data to shed light on varying access over time and by location. Wilson and Bridge (2019) review international evidence on school choice systems, finding that increased choice usually goes hand in hand with increasing social stratification between schools. Allen and Higham (2018) assert that new free schools are socially selective, but they do not account for the possibility that the sorting they observe is caused by choices, rather than constraints. The previous work to examine the variation in access to school quality in England includes Allen et al. (2014), Hunt (2018b, 2019a, 2019c), and Burgess et al. (2019). This report provides evidence that between-school stratification along social and ethnic lines is driven primarily by a combination of heterogeneity in preferences and residential segregation, rather than by the way in which selection by schools has worked.

At first glance the patterns, revealed by the recent availability of data on parental preferences and by school allocations data, point to a picture of admissions inequality, where schools appear to preferentially select white, non-poor pupils and exclude those eligible for Free School Meals, and minority ethnicities. However, cream-skimming to the extent suggested by the data is implausible, as a majority of schools do not even have sufficient admission autonomy to discover the income, ability, or ethnicity of pupils, let alone act on that information. Indeed, further inspection shows that the patterns instead point to the role of parental preferences and the strategic element of their decision-making.

This is the first work to attempt the task of disentangling choices from admissions success. This report summarises how we achieve this by explicitly modelling the probability of admission to schools, and how we use this model to decompose the variation in the quality of allocated schools into components due to geography, admissions, and preferences.

- What weight do parents place on the factors that they trade-off against each other when evaluating schools: school performance, proximity, and admission chance?
- How much variation is there in the weights that parents use? In particular, do these weights, that determine preferences, vary across types of parent?
- To what extent does the design of the system affect the quality of choices that parents experience? By *design* we mean: school locations and capacities, the relevance and availability of information, and the nature of oversubscription priority rules.
- Are there simple interventions that can improve choices by parents?

Before we spell out how we attempted to answer these questions, we preview our findings. In brief, these are:

- On average, we estimate that parents place a considerable weight on school performance (our proxy for quality). By observing their choices we suggest that parents are prepared to allow their child to travel an additional 0.9 km (when the mean distance is around 2.5 km) to achieve a 10 percentage point better quality school. This is a considerable burden that households seem to be willing to pay.
- Holding other variables equal, minority ethnic groups are more willing to travel for incremental improvements in school performance than white parents. White parents are willing to send their children 11% further for a 10 percentage point improvement

in a school performance measure (proportion achieving 5+ good GCSE's), whereas minority ethnic parents are willing to travel 21% further for the same improvement.

- Minority ethnic families are, on average, 17% less likely to achieve their first-choice school (and more so for Black, than Asian or Other), and this pattern persists when looking only at London. Overall, Londoners are less likely to get their first-choice.
- After accounting for ethnicity, parents of children with attainment in the top tercile of KS2 (end of primary school) tests, are willing to travel 50% further for a 10 percentage point improvement in test scores, than the families of children in the bottom tercile.
- We explore the simplest and cheapest possible intervention – extending the length of lists that parents are able to specify. Our data suggests that many Local Authorities restrict the ability to list sufficient schools – something that causes parents to be too conservative in their choices. While this enables LAs to say that it has a high proportion of children attending their first choice – but if their first choice is an unduly safe one then this is a hollow achievement.
- But further reforms could improve the choices that parents make – through better information provided using tools that are familiar to us as consumers in the context of choosing hotels, flights, and movies.
- In work in progress, we show that reforms to admission criteria offer the possibility of manipulating the allocation of children to improve the chances of disadvantaged children attending more effective schools and benefitting from high ability peers. Additional social mobility might be obtained at minimal cost.
- In the longer term, we believe that the research here could be extended to include not only the role of peer effects but other policies could also be considered – those that change the nature of the choices available through improving schools, changing the nature of others, and closing/expanding/relocating existing schools. A major challenge to this long-term agenda would be to incorporate location decisions by parents.

The rest of the report is structured as follows. An important innovation here is the use of data on parental preferences, so Section 2 is dedicated to explaining the data that we use. The salient basic details of the way in which school choice works in England in given n Section 3. A major aspect of our work is to be able to distinguish between the roles of preferences that drive demand and school capacities and the process by which schools ration their capacity when they are oversubscribed. Section 4 explains our approach to modelling preferences and describes how we model admission. The probability of admission plays into preferences because the school choice system incentivises strategic behaviour in parents and this is outlined in Section 5 and the basic statistical findings are presented. Section 6 then uses the statistical estimates to make inferences about how well different groups of the population are served by the school choice system. It also shows how it well works, and how it might work better, across areas of the county. Section 7 speculates on how the school choice system could be improved. Finally, Section 8 explains how the research is being extended to show how social mobility might be served by changing admission criteria and how the work could be extended further to analyse a wider variety of policy issues.

2. Setting and data

To gain insights into choices and admissions, the project capitalises on the newly available parental preferences data for England. The full dataset consists of records for all children who were in year 6 in the academic year 2013/14, and were applying to English state-maintained secondary schools for entry in September 2014. For each child we have a record listing up to six schools that were ranked by the child's parents as part of the Local Authority's co-ordinated allocations process. We also have the child's Key Stage 2 (KS2) results and year 6 census record, and the linked year 7 census record for 2014/15, were also available. The school identifiers allow linking to school performance tables (containing school-level GCSE performance measures), the school-level census data (for demographic information), historical Ofsted data, and to Edubase, the public database of schools' information.

The preferences dataset identifies the schools listed by each parent, and also identifies the school that was offered to the parent at the conclusion of the admissions process. From this we can determine the family's chosen school, whether or not the child gained admission to that school, and the rank of the school that the child *did* gain admission to. The linked datasets provide the child's location, prior attainment, ethnicity, Pupil Premium status and other characteristics, as well as identifying the secondary school the child actually enrolled in, if the child stayed in the state sector.

The cohort contains rank-order lists for around half a million children who are in their final year of primary school (aged 10-11), linked to their home postcode, ethnicity, gender, Pupil Premium status, and primary school test scores from the National Pupil Database (NPD). From the home postcode we can calculate their location and proximity to all schools [3]. The NPD also provides detailed information on schools, including location, test scores, demographic composition, governance and religious denomination. We also obtain current and historical government inspections (Ofsted) scores from the online schools database, Edubase.

Individuals are eligible for inclusion in the sample if there exists a rank-order list record for the child, and additionally they are either included in the year 6 census[4] (final year of primary school), or the year 7 census (first year of secondary school), or both, and their home postcode, ethnicity and KS2 attainment are not missing. This means that those children transitioning from a private primary school, or who eventually transition to a private secondary school, are potentially included in the sample as long as their parents participate in the state-school admissions cycle. Our report employs a consistent demographic cross-classification to allow for variation in parental income, academic attainment, and ethnicity.

Much of the existing work on school choice focusses upon variation by socio-economic status or incomes, and finds significant variation, so it is important to allow for this dimension in our classification by including Pupil Premium status as a demographic indicator. The Pupil Premium (PP) is an initiative that awards additional funds to schools for each child on roll who is eligible. Pupils are eligible for PP if they have been eligible for Free School Meals at any time in the last six years, and pupils are eligible for Free School Meals if their parents have low incomes and are in receipt of certain benefits. PP is therefore a proxy for low incomes.

[3] School proximity is measured as the straight-line distance between the centroid of the home postcode and the centroid of the school postcode, in kilometres.

[4] Children in private schools are not included in the National Pupil Database census.

We also control for variation in attainment at primary school by including tercile of Key Stage 2 (KS2) SAT scores in the classification. KS2 tests are taken by children at the end of the final year of primary school (at the age of 11) and give an aggregated score for ability in English and Mathematics. It is important to note that, as the admissions cycle takes place at the beginning of the final year of primary school, children will not have taken the test when their parents are choosing schools, and schools will not observe KS2 attainment. However, KS2 terciles are used as a correlate of unobservable family and child characteristics that affect school choice and admissions. It is possible that KS2 terciles are endogenous with respect to admissions outcomes – for example, it may be the case that the parents who live close to a good school are also likely to urge their children to do well. However, this is not likely to substantively impact our results as KS2 tests are low stakes, that have no bearing on admission priority at schools, even though they may cause parents and children some stress. Since they have no bearing on admissions, there are no incentives for parents or children to alter effort in response to admission decisions.

However, KS2 may be correlated with unobservable variation in parental characteristics and behaviour. In particular, KS2 outcomes may be affected by spillovers from parental investments in gaining entry to grammar schools, such as helping children with homework, maintaining discipline and procuring private tutoring. KS2 tercile therefore incorporates information about the parents', as well as the children's, unobservable type with respect to human capital investments.

Ethnicity is measured in the NPD based on parental reports. There are 18 ethnic groups in the original data but for most of the analyses here a two-way white/minority classification has been used. While this two-way classification hides all distinctions between non-white ethnic groups, it ensures that the aggregated sample size of the minority group is large enough when cross-classified by income and attainment, and also ensures sufficiently large sample sizes outside London and the major cities. A four-way classification (white British; black; south Asian; other) has been used in some of the graphical descriptive statistics. Weldon (2018) presents descriptive analyses and also a discrete choice model estimated using the four-way classification, with the conclusion that the important distinction with regards to school choice and admissions is between white families and all other ethnic groups. Ethnicity is related to other constructs such as nationality, religion, language and the length of time spent living in the UK. It is possible that phenomena ascribed to ethnicity in the paper should more properly be ascribed to one of these other constructs. However, the other constructs are not recorded in the data used in this study.

There are therefore 12 demographic groups defined by the cross-classification of ethnicity (white/minority), Pupil-Premium status (eligible/not eligible) and KS2 attainment terciles (high/middle/low). For each child we observe their demographic group, home location at the postcode level, rank-order list, the offered school and the school that the child was subsequently enrolled in.

3. The English admission system

In England, children transition from primary school to secondary school at the end of the school year in which they reach 11. To apply for a place at a state secondary school in England, parents submit a ranking of their preferred schools. They can list between 3 and 6 schools depending on the Local Authority (LA). Parents must apply within their own LA, but may include any school within or outside their own LA on their list. Local authorities allocate places according to a matching mechanism co-ordinated by each LA and through collaborations between neighbouring LAs.

A matching mechanism is a (usually computerised) procedure taking as inputs parents' stated preferences (submitted as a rank-ordered list of schools) and schools' capacities and over-subscription policies (submitted by schools as a rank-ordered list of eligible pupils), and outputting a school allocation for each child. The rationale for matching mechanisms is to provide a fair allocation, acceptable to all parties, while avoiding the time and inconvenience for both parents and schools, that would arise if allocations were decentralised[5]. The two main types of matching mechanism that have been used in England are:

First-Preferences-First (FPF) In the first round, children are allocated to their first choice. If there are not enough places, the school's priority rules are used to decide which first-preferers gain a place. If there are any places remaining, the procedure is repeated for those ranking the school second, third etc. This method was common up to 2006.

Deferred Acceptance (DA)[6] In the first round, pupils are *provisionally* allocated to their first choice, using the schools' priority rules where schools are oversubscribed. In the second round, pupils left without a place in the previous round are provisionally allocated to their second choice. However, in DA (and not in FPF), a pupil who has higher priority at her second choice than someone who was provisionally allocated in the first round, can knock that person out of their place. Therefore, having a first preference for a school does not give a child priority at the school. In the third round, those without a place (including those knocked out of their provisional place in the previous round) are provisionally allocated to their next choice, and so on. This was mandatory from 2007 but was also used in many LA's prior to this.

Although efficient (as they maximise the number of people obtaining their first preferences) FPF mechanisms were criticised on the grounds that they forced parents to consider the chances of admission when stating their preferences, and were vulnerable to gaming by sophisticated parents (Abdulkadiroğlu and Sönmez, 2003). Since FPF mechanisms were banned in 2007 (Department for Education and Skills 2007), use of the DA algorithm (Gale and Shapley, 1962) has become ubiquitous (Coldron et al. 2008; Pathak and Sönmez, 2013).

For oversubscribed schools, allocation of places is prioritised based on a set of criteria which depends on the school, typically including: whether the child is under the care of the local

[5] To give an example of the difficulties with decentralised matching, suppose a parent applies separately to two schools. He receives an offer in March from his second favourite school, and accepts. His favourite school has this child on a waiting list until July, and then finally offers this child a place. He rejects the offer from his second favourite school, which now has an empty seat, which it manages to fill in August. At this point, 3 weeks before the first day of term, some other school has an empty seat and has to make another offer, and so on.

[6] There are two variants of this mechanism (pupil-proposing and school-proposing). The pupil-proposing mechanism is described here. The two mechanisms are similar, but do not always produce equivalent outcomes, and it is not known whether all English LA's use the pupil-proposing variant.

authority (``looked-after'' children, LAC); whether the child has an older sibling at the school; and, finally, the distance from the school. In some schools, religious worship or baptism forms an additional criterion that is inserted somewhere between the proximity tiebreaker and the LAC criterion. In a small proportion of state schools, a priority is attached to those displaying an *aptitude* at a particular subject or a range of subjects. Finally, there are a small number of grammar schools which fill all of their places based on performance in tests or an exam, known as the 11+ exam.

If a pupil cannot be allocated to any of their listed preferred schools, they are assigned to a school with spare capacity and this is not necessarily the nearest. In spite of this, many parents list fewer than the maximum allowed number of schools.

For the purposes of this study, schools have been categorised into six broad groups, according to their admissions policy and ethos:

Community schools The first group comprises those schools owned and controlled by LAs, comprising 20% of secondary school places in 2014. These schools generally have simple admissions criteria, prioritising siblings, 'looked after' children and those living within a designated zone, with straight-line distance used as a tie-breaker. Until recently the largest group, since 2010 many formerly community schools have been converted into state-funded autonomous schools called 'academies'.

Non-faith academies[7] This second group is now the largest group, enrolling 57% of state-funded secondary school children. These schools have some autonomy to set their own admissions criteria although within the strict guidelines set by the government. Many academies have similar admissions criteria to community schools. Some academies include aptitude in a particular subject or range of subjects in their admissions criteria. Some schools operate 'fair banding' criteria, where a quota of children is admitted from each attainment (KS2) quantile. This category includes Free schools, which are a type of academy set up by parents or other interest groups.

Roman Catholic schools 11% of secondary school children are enrolled in Roman Catholic schools, making this the largest faith school denomination in secondary schooling. RC schools usually select up to 100% of their intake on religious grounds. At many schools proof of baptism is sufficient, although at the more popular schools proof of regular church attendance may be required. Some schools reserve a proportion of places for children of other faiths/no faith.[8]

Church of England schools The second largest providers of denominational secondary schooling are Church of England (C of E) schools (7% of places). A majority of C of E schools also require proof of religious worship for some or all places. The admissions criteria of C of E

[7] Many faith schools and grammar schools are also designated as academies, but for the purposes of this study this category only includes those academies which are not faith schools or grammar schools.

[8] Faith schools are commonly thought of as high-quality schools. They are often oversubscribed. Yet there is almost no quantitative literature, for England, that addresses the effectiveness of faith schools. McKendrick and Walker (2020) uses LSYPE data to estimate the effect of attending a faith school. The raw data shows that faith schools do have better educational outcomes. However, they find that the alleged advantages of faith schools are not robust in the data to more detailed analysis. The only robust finding is that children that profess faith and attend a faith school do tend to be more likely to retain their faith into adulthood.

schools are more heterogeneous, and support for religious selection less unanimous in the Anglican sector than in the Roman Catholic sector.

Other faith schools There are a very small number (less than 1%) of schools with religious denomination other than Anglican or RC. These include schools catering to Jewish, Muslim, Buddhist and minority Christian sects. These schools often operate faith-based admissions criteria for some or all places.

Grammar schools There are just 162 academically-selective grammar schools (enrolling less than 5% of children) -- the remainder of a much larger system of academic selection that existed prior to the 1970's. To obtain a place at a grammar school, children are required to sit an academic exam. In some Local Authorities all 11-year old children sit a common exam, whereas in others only children who wish to apply to a school sit the school's exam.

Figure 1 summarises the excess capacity for the major school types listed above. There *is* excess capacity **overall** – driven by Academy and Community schools. However, this aggregate picture masks the uneven geographical supply of capacity, with large surpluses in some areas, and a dearth of places in others. Catholic and CofE schools have small overall over-capacity, while Grammar Schools face excess demand and every place is filled so that admissions equal capacity.

Figure 1: Capacity, demand and admissions in 2014 at each school type, as a proportion of total capacity in year 7 in England.

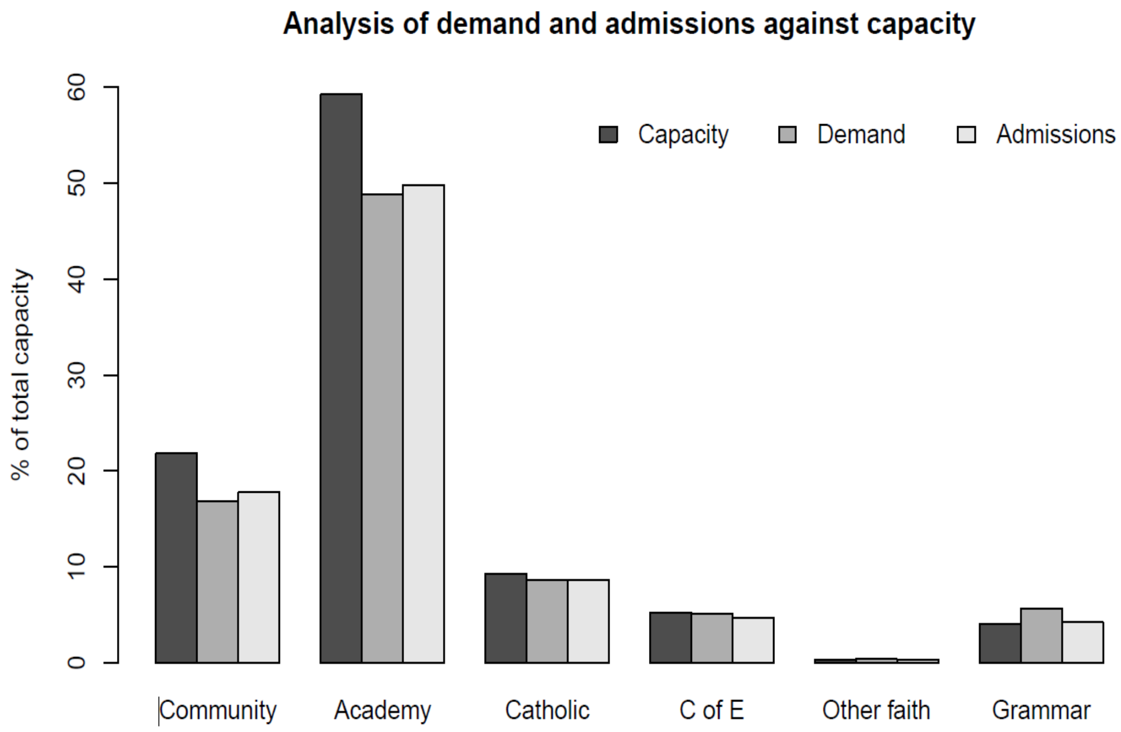

Strict guidelines regulate the admission criteria that schools are permitted to use. Some unlawful criteria include: interviews or other face-to-face contact; rank-order (eg. first-preferences-first); and parental financial contributions or volunteering. In addition, new schools have more stringent restrictions on religious and academic selection than existing schools. All own-admissions schools may opt to receive a list of applications (without information on the preference order) and rank them before sending the ranking back to the LA for the matching to be computed. Commonly, local authorities provide services to assist the ranking of students (for example, the calculation of home-school distances) as traded services to own-admission schools. For community schools the LA computes the ranking.

In contrast to some other countries that operate centralised school admissions, lotteries are almost unheard of as a priority tie-breaking mechanism[9] in England. The final tie-breaker for the majority of community, academy and religious schools is almost invariably the straight-line distance between the school and the pupil home address. The distance tiebreaker creates the conditions for house prices to be affected by the proximity of good schools (Gibbons and Machin 2006, 2008; Gibbons, Machin, and Silva 2013; Machin and Salvanes 2016).

Secondary school admissions are coordinated at the Local Authority (LA) level in England. In practice, this means that LA's must provide a centralised clearing house, with a common admission cycle to which all schools must subscribe. There is no option for state-funded schools to admit pupils into the entry year outside the co-ordinated admission mechanism. However, in-year admissions for years other than the entry year are decentralised, and handled by each school separately. Since the introduction of school choice in 1988, successive Admission Codes have stipulated progressively tighter controls on the operation of admission systems. In 2007, the Admission Code outlawed mechanisms that allow schools to prioritise pupils based on their order of ranking (called 'first-preferences-first' mechanisms). Although efficient in the sense that they maximise the number of families who access their first-choice schools, such mechanisms are thought to provide parents with incentives to game the system by misreporting their preferences, which could risk benefitting more sophisticated, well-informed, parents at the expense of others, and skew statistics on the admission system by inflating the apparent number of families accessing their first-choice school. The admission legislation has meant that, in terms of the clearing house mechanism itself, and permitted admission oversubscription rules, LA's are similar to each other. However, in one aspect of the admission mechanisms there is still substantial variation - the maximum permitted size of submitted preference lists.

Figure 2 shows that LA's allow parents to rank between three and six schools. The figure covers all English LAs and the first two rows are London LA's. London LA's all allow up to six schools to be ranked and listed in submissions. Elsewhere, there is variation in the number of choices offered even across LA's with a similar level of urban density. Moreover, parents do not always use all of their available choices, and there is wide variation in the distribution of choices used as shown in Figure 2. Many of the LAs that allow six choices have very few parents that use more than three; but many find that all 6 choices are often used. Among LA's with a maximum of three slots, for example in County Durham or Cornwall, a majority of parents use only one preference slot, whereas in Lancashire the majority of parents do use all of their three available preference slots. Some allow three and many such LAs find that many parents submit only one or two schools suggesting that three is enough (for example, Gloucestershire). But other LAs allow three and many parents use all three (for example, Lancashire) – suggesting that if parents were allowed to list more than three then they would use the extra choices made available.

[9] One exception is Brighton and Hove LA, which has, since 2007, operated a system based upon catchment areas within which lottery tie-breakers operate.

Figure 2: The proportion of parents that submit rankings of between one and six schools by LA

4. Analytical framework

Parents are asked to rank schools in order of their preferences – which depend on quality and proximity, but we also allow for preference for a school to depend on the probability of acceptance – that is we allow for risk aversion.

Child application details are confronted with the admission criteria of schools that they apply for. This establishes their eligibilities and then children are allocated to the highest ranked school that they are eligible to attend. We do not know the details of the actual admission criteria used by each school. But we know that schools will have an ordered set of priorities such as SEN, LAC, and siblings; and that some schools may also have geographical priority areas: some will prioritise faith in some way; and some may offer a priority for "aptitude" (but not "ability") say for music, sports, or even maths. Remaining places are typically allocated by crow-flies proximity from home to school. We do not know the details, but our analysis starts by inspecting the data so as to infer the point at which priority applicants have been exhausted and remaining places are allocated for proximity. We do know the capacity of all schools. So, when a school is full remaining applicants are then considered for their second placed school. Simply by knowing the point at which allocations of school places begin to be awarded by proximity and the point at which this ends (i.e. the school is then full) we are able to consider the effects of a number of possible reforms. For example, we could add a further priority criterion below whatever the existing priorities are. One obvious example would be to prioritise FSM children above those that live nearby. We could also replace the proximity criterion by a lottery. We could choose to leave Grammar Schools outside of any reform package, or we could include them by replacing admission by ability (which we assume can be proxied by KS2 since we do not know the actual grammar entrance test scores).

Having allocated all children to all schools we know which school each child attends and therefore we know the peer mix at each school. To predict the KS4 outcomes we use an estimated education production function that includes own KS2, peer KS2 effects and an addictive school fixed-effect.

4.1 Parental preferences for schools

Our modelling of preferences is based on the presumption that parents like good quality schools (denoted Q) but dislike their child travelling a long distance to attend (denoted D). We assume that parents are "agents" for their children in such a way that child preferences are reflected in those of their parents – so we do not distinguish between the preferences of parents and children. We also allow preference variation across observable characteristics of the children, **Z** (which includes ethnicity, FSM status, and our measure of prior ability - the child's KS2 score at age 10 prior to secondary schooling).

We also allow for preferences to vary across the unobservable characteristics of parents, denoted by ξ_i. Note that we do not know ξ_i – we think of this as something that is unobservable but explains the variation that we cannot explain using the variables that we do observe. Astrophysicists, faced with the problem of understanding the universe, name ξ_i "dark matter" – a name that acknowledges the mysterious nature of the universe and yet suggests the hope that it may ultimately be observed. In contrast, economists often refer to ξ_i as the "error term" – a name that suggests that we have made a mistake and holds out little hope that this will be ultimately reversed. Economists have been famously, and perhaps correctly, described as "dismal", while astrophysicists are famously expensive in their pursuit of knowledge.

Hereafter, we will refer to ξ_i, more agnostically, as "unobserved heterogeneity" and we refer to the associated variation in behaviour as "idiosyncratic".

Our shorthand for the well-being of parents which, following standard economics nomenclature, we refer to as "utility", is given by

$$U_{is} = U(Q_s, D_{is}; \mathbf{Z}_{is}, \xi_i) \quad (1)$$

where U_{is} is the utility of the parent of child i when the child attends school s. This is assumed to depend (positively) on Q_s (the "quality" of school s), which we think as a school level outcome. This can be measured by KS4 performance at the school such as the proportion of children at *s* attaining 5+ GCSE "passes", or some other summary metric.[10] It could be a school contextual value added measure. In addition, U_{is} depends (negatively) on the distance that child i would need to travel to attend school s. We measure this as the crow-fly distance from the centroid of i's postcode to the centroid of the school's postcode, in km. We think of this as a sufficient statistic for the time and trouble incurred in having i attend s. In principle, but this could be refined using actual travel time (and cost) and actual distances (from Google's API or other source). Parental (and child) background variables are also included in U_{is} and this is summarised by the vector (i.e. list) of variables denoted by $\mathbf{Z}_{is}$. But conditional on $\mathbf{Z}_{is}$ there is, on average, a systematic relationship between U_{is} and its other observable determinants - Q_s and D_{is}.

Thus, our model regards well-being, or utility, as a component that varies systematically with its observable determinants, but there is an unobservable component, denoted ξ_i, that is assumed to be randomly distributed, according to a bell-shaped distribution, across parents. This is sometimes referred to as a "Random Utility Model". So, although all parents feel differently about the school options available to them we can incorporate these differences by assuming that they follow a specific statistical distribution. Here we assume ξ_i to be randomly Normally distributed, and independently of the other variables in the model.

We aim to construct measures of school choice and of choice set "amenities" (the characteristics that affect parental preferences: e.g. distance and quality, among other things that might be attractive, or not, in the school neighbourhood) that are more informative than the more usual "% achieving first choice" summary statistic. We also wish to decompose the variation in school matches that arise because of differences in parental preferences and how much arises because of differences in school capacity constraints. To do this, we need to define the structural models of both parents' preferences and of constraints that are used in these measures.

Utility is a function of observable "amenities" provided by a school (such as test scores, and proximity to the family home), and there are likely to be other amenities that matter to the parents, but are unobservable to researchers. We make no assumptions about the relationship between utility and parents' expectations of the short- or long-run outcomes

[10] The data represents a cohort of parents who did not have a history of performance measures available to them. The contextualised value-added measure used before 2010 had been abandoned, and the Attainment 8 and Progress 8 measures had not yet been introduced. The performance measure in most widespread use in 2013 was the proportion of children at each school achieving the threshold of five GCSEs at grades A* to C. So, we adopt this measure as our measure of quality because of its ubiquity and salience to parents. However, it is a raw attainment measure, rather than a measure of quality per se, and as such it is highly correlated with the demographic characteristics and prior attainment of the school's intake.Nonetheless, none of the qualitative results hang on the precise definition.

from education; we simply assume that each parent implicitly assigns utility scores to alternatives, and chooses the set of *L* schools that would generate the highest expected utility score, out of all other possible combinations of *L* schools.

We need estimates of the parameters of these systematic determinants, Q_s and D_{is}, to be able to construct the level of well-being, or utility, associated with any match of i to s, so we need to adopt a specific shape, or *functional form*, for equation (1). This function incorporates the assumptions required to aggregate the determinants of U_{is} that reflects the way individuals weight together D and Q. The specific equation that we adopt for utility is a simple additive one – utility depends linearly on Q and linearly on log D. We choose a logarithmic transformation of D, rather than D itself, to reflect our presumption that an extra kilometre makes a smaller difference to the intensity of preference for proximity the further away one is. This log transformation of D effectively assumes that U is linear in *proportional* (i.e percentage) increases in D. That is, we make (1) more specific in (2) by assuming that it is a specific “structure” that is linear in Q and log-linear in D, and that there are number of other observable variables, in the vector **Z**, that also affect well-being, so that

$$U_{is} = U(Q_s, D_{is}; \mathbf{Z}_{is}, \xi_i) = Q_s + \rho.\log D_{is} + \mathbf{Z}_{is}\chi' + \xi_i \qquad (2)$$

Thus, Q_s is some measure of quality of school s (which might be a value added measure, or simply the proportion of children in s that obtains, say, 5, KS4 passes, or something more complicated), D_{is} is distance from i’s home (postcode) to school s (postcode), and **Z** is a list of school and individual level characteristics. The variable ξ_i represents the unobservable determinant of U. ρ is the unknown parameter that drives the trade-off between proximity and quality that parents are willing to make – we anticipate that ρ<0 because we expect individuals to be prepared to go further to get better Q but D itself has a negative effect on U for given Q. Note that we assume that the coefficient on Q_{is} is 1 - this is innocuous because we are only concerned about the *relative* effects of Q and log D. Our aim is to estimate the value of ρ based on what we observe in the data on parental preferences.

The interpretation of ρ is the *percentage* increase in D that a family would allow their child to travel to school to get a positive unit higher level of Q. We therefore think of ρ as the “willingness to travel” to obtain higher Q . In practice we allow the ρ parameter to differ across ethnicity etc. by estimating our equation (2) for U(.) separately for each of our 12 groups of children (categorised by own ethnicity, FSM status, and range of KS2 score). This is a straightforward way of allowing the parameters of the preference model to vary across groups of parents who differ in their observable characteristics, **Z**.

Figure 3 is a graphical description of what U_{is} looks like. Figure 3 is drawn in Q vs log D space. That is the vertical axis is Q and the horizontal is log D. One can rearrange equation (2) to have Q on the left hand side so that it reads $Q_s = \rho.\log D_{is} + \mathbf{Z}_{is}\boldsymbol{\chi}' + \xi_i - U_{is}$, and when written in this way it shows that Q rises (linearly) by ρ per unit change in log D (i.e. doubling D) for a given level of U. Or, put another way, a parent would be equally happy that the child be at a school that is ρ% further away if its Q were one unit higher.

The solid blue lines show combinations of Q_s and log D_{is} , for an assumed level of ρ, that provides the same level of well-being to i. So any two schools on a given line of slope ρ, say s_1 and s_2, are such that $U_{i1} = U_{i2}$ – that is, the parent is *indifferent* between s_1 and s_2 since they lie on the same blue line. That is, s_1 and s_2 have the same value of U because 1’s lower Q , relative to 2, is just low enough to offset its closer proximity. And any school on a higher line, such as s_3, , such as s_1 or s_2 , will have higher utility. Thus, s_3 is preferred to s_2 because it has higher Q and is no further from home, while s_3 is preferred to s_1 because, although its further from home, it is much higher quality. Thus, a parent faced with these choices of school would rank s_3 higher than s_1 and would place schools s_1 and s_2 as equal second.

But preferences can differ across parents. For example, the dashed blue lines shows the preferences of another parent who has a much higher value of ρ, say ρ'> ρ, then she would demand a much larger increase in Q to be attracted to a more distant school than was the case with the parent who had a slope of ρ. We allow for such heterogeneity in the analysis.

Figure 3 *Graphical description of preferences*

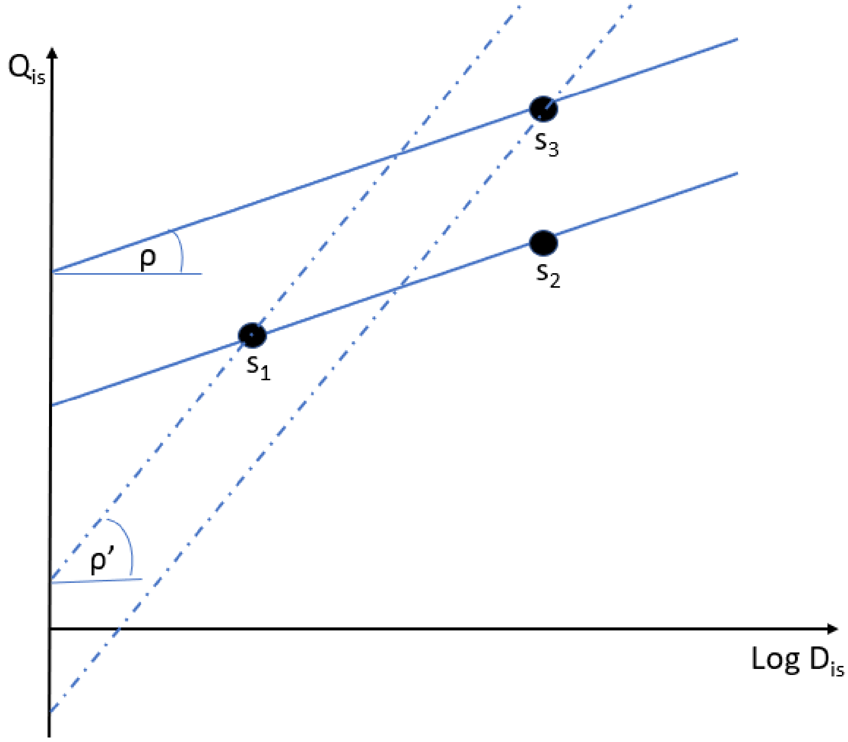

4.2 A model with strategic choices

If all schools had surplus capacity, this model would be sufficient since strategic behaviour is unnecessary if capacity constraints were non-binding: each parent would simply choose the alternative that maximises U_{is}. If parents were allowed to rank all possible schools then they would be rational to do so. If the list length is limited then parents would then realise that the probability of admission would depend on the rankings of other parents. This then generates incentives for parents to rank schools "strategically" – i.e. take into account the behaviour of other parents. It is unclear how they would do this, but a sensible strategy would be to choose, not the *L* best schools from the universe of schools, but to choose *L* schools that would generate the highest **expected** utility or well-being.

The logic of the DA admission mechanism is that parental rankings will express the parents' true preferences – that is, parents rank **only** on the basis of the well-being that they would enjoy if their child *were attending* each school. If you can list as many schools as you like then you can be sure to get into the best school in your list that you meet the criteria to be admitted. In particular, parents will not need to (or wish to) take account of the probability of being *able to* attend each school. Since parents get the match to the school that they rank highest provided they can be admitted, they have no cause to even consider the chances of getting in.

The "mechanism design" theory underlying school choice that lies behind this presumption is based on the assumption that parents are able to rank ALL schools. In practice, Local Authorities limit parents to creating a listing of only a small number of choices of schools, say L of them. If L is small (and, in our data, it is often as small as 3, and never more than 6) then parents are then incentivised to factor into their calculation the chances that they will be **accepted** at the schools they would like to attend. That is, a parent may think she needs to play safe – to consider, in a "strategic" way, the rankings of other parents, which affects the chances of getting into schools that she likes, when choosing her own ranking.

Here we pragmatically allow for this by incorporating an additive direct effect of the probability of i's admission to s, denoted P_{is}, into our expression for U_{is}, so that (2) becomes

$$U_{is} = U(Q_s, D_{is}; \mathbf{Z}_{is}, \xi_i) = Q_s + \beta.\log P_{is} + \rho.\log D_{is} + \mathbf{Z}_{is}\boldsymbol{\chi}' + \xi_i, \qquad (3)$$

where $\beta<0$ represents the parent's attitude to the risk associated with choosing s, P_{is}. We use the log of P_{is} because when $P_{is}=1$ the model becomes the same as (2), and if it were 0 (i.e. no chance of acceptance) then that school would be given infinitely small weight in the parent's preferences and it would be as if it did not exist for this parent. Note that P_{is} depends on how many children live nearer or have higher priority than i at s for some other reason, such as having a sibling at s. On the other hand, if i were to have a sibling at school s_0, say, then $P_{i0}=1$ because this sibling provides a guaranteed entry ticket to school s_0. But, even so, the parent might still want to consider different schools to the child's sibling, $s \neq s_0$ since a much better school might outweigh the convenience if sending the younger sibling to the same school as the older one.

We estimate the parameters of this model using, so called, *discrete choice* estimation, and we allow each parameter to vary across each of our 12 groups defined by ethnicity (white British vs the rest;), income (on FSM vs not), and ability (bottom , middle, and top thirds of the KS2 distribution).

4.3 Probability of admission

P_{is} in (3) is not something we can observe – it has to be calculated.[11] In order to quantify the expected quality of choice sets, and to account for school capacities and admission probabilities, we need to estimate the chances that a child will be successful in obtaining a place at a school, if she applies. P_{is} will be 1 is *i* for all possible *s* if she in LAC or SEN. It will equal 1 for a particular *s* if she has an older sibling in that *s*. Otherwise, *i* will have some probability, less than 1, of gaining admission to each *s* depending on the capacity of *s* and *i*'s proximity to *s.* The bigger is *s* and the closer that *i* lives to *s* the better the chances of gaining admission; but the more SEN and LAC children and others with older siblings in the school that would also like to attend the smaller will be the chances.

There are a handful of other papers that explicitly estimate joint models of strategic choice and the probability of admission (Ajayi 2013; Calsamiglia, Fu, and Guell 2014; Fack, Grenet,

[11] Note that P_{is} depends on D_{is} because D_{is} determines the number of children who live closer and so have higher priority than pupil *i* at school s. In fact, P_{is} will be (approximately) inversely proportional to the **square** of D_{is} because the number of children living closer is likely to be proportional to the *area* of the circle with radius D_{is}, which depends on D_{is}^2.

and He 2015; Ajayi and Sidibé 2016; Agarwal and Somaini 2018; Luflade 2018). Of these papers, Fack et al. (2015) avoid modelling strategic choice, but instead posit methods for recovering the distribution of preferences that are robust to strategy. Calsamiglia et al. (2014) and Agarwal & Somaini (2018) model decision-makers divided into *strategic* and *sincere* types, and estimate the proportion of each type in the population. Both papers find that the population contains both strategic and sincere decision-makers, but do not characterise this variation in terms of demographic variables. This report takes an approach similar to Ajayi (2013) and Luflade (2018) in modelling admission probabilities directly from what we observe about *I*'s characteristics.

We take this probability approach throughout the report. Even if the priority mechanism were completely transparent and all data were available, admission would not be deterministic because; although, in that case, the schools orderings of students would be known, the proximity threshold (cut-off) at which schools cease to admit pupils would still not be known. This priority cut-off cannot be known in advance of the allocation, either by parents or to researchers, because it depends on the interaction of the demand for the school (i.e. the aggregate preferences of parents) and the school's capacity, and the capacities of other local schools due to 'overflow'. Additionally, parents often do not have complete information about their priority For example, if the priority is based on religious adherence and must be evaluated by school staff. There may also be uncertainty about a child's position in the priority ranking, since this depends on overall demand for the school for those with greater priority – for example, the number of children in the same cohort that have older siblings at the school. Uncertainty about the cut-off therefore induces uncertainty about the chances of admission. There are therefore multiple sources of uncertainty for parents that mean that it is more appropriate to model admission as a subjective probability.

4.4 Measurement of choice set quality

Our main aim is to use our "structural" model of school demand and supply to evaluate the quality of school amenities (i.e. quality and proximity) available to parents of different demographic groups, in different locations. Previous work to examine the variation in access to school quality in England includes Allen et al. (2014), Hunt (2018b, 2019a, 2019c), and Burgess et al. (2019).

Our approach contributes to this literature by explicitly considering the level of school quality that families of different demographic groups, in different localities, can expect to be able to access, taking into account not only geographical proximity, but also admission criteria, school capacity, and differences in individual preferences.

In order to decompose the relative contributions of market, admission constraints, and choice behaviour, we quantify the quality of choices available to each child in terms of three possible definitions:

A. The well-being associated with the local market, defined as the utility of the best school in each family's choice set, where 'best' is defined using estimated preferences, and ignoring any capacity constraints or admissions rules. This may not coincide with the school the family actually states to be first on their rank-order list, because it may be a "strategic" choice.

B. The best well-being that a family can expect to achieve, given constrained school capacities and admission rules, defined as the school performance and proximity that a family can expect to achieve if it submits its best possible rank-order list – where the best possible rank order list takes into account admission probabilities, and the constraint on the number of schools that can be listed..

C. The utility of the lists that parents actually submit. This is the actual utility of the school that children were allocated to.

Definition A picks one school from each family's choice set based on estimated willingness to travel for school quality[12] and yields a level of well-being W_A. This is the most common way of thinking about the quality of choice and is the way that has most often been studied in the past. But it is an appropriate definition only if there are no binding capacity constraints. Definition B reflects what parents can expect to experience, given both the local distribution of schools, *and* the constraints in gaining admission to those school and yields a level of well-being W_B Welfare according to definition C, call this W_C, can be measured directly from observable data. It is the realised utility from the lists that parents submit resulting in the schools they actually attend.

These definitions decompose the contributions of market, admission constraints, and behaviour in the following way:

- The difference between W_A and W_B for a given family is the welfare foregone due to admission constraints – if all schools were undersubscribed then they would be equal.
- The difference between W_B and W_C is the welfare forgone due to parents' idiosyncratic preferences, which may trade-off observable amenities for amenities that are *unobservable* in the data. In addition, the difference between W_B and W_C will, subsume any mistakes due to imperfect information about performance or the chances of admission.

Variation in W_A across families in different localities represents the part of the variation in quality of choice that is attributable to the uneven geographical distribution of high-performing schools. When policy-makers speak of 'cold spots' in school quality they often operationalise the idea in this way. However, our framework takes account of, for example, that there may well be a high-performing school within a reasonable travel time of the vast majority of children, but often those schools are over-subscribed, and therefore only a small minority of children may be able to access them.

The framework also allows us to evaluate how much of the variation, across the 12 different demographic groups, that achieve a place at a high-performing sch is due to individual decision-making: either parents not valuing performance (as measured) or not being aware of measures of performance, or not using the admissions system in the optimal way.

Definitions A and B do not model the admission system and capacity constraints. Definition B is our main contribution to this literature. W_B provides a realistic measurement of variation in the choice sets available to families of different demographic groups, allowing for the possibility that two families living in the same street may have very different choice sets, due to having differing admission priorities at local schools.

[12] The discrete choice model is estimated separately for each of the 12 demographic groups defined by cross-classifications of ethnicity, pupil premium and KS2 attainment. We predict choices using type-specific estimates of willingness to travel.

Given an admission system, comprising an school choice mechanism including details such as the maximum number of preferences that parents can state, a set of school capacities, and each school's admission priority rules, a parent maximises the expected quality of the allocation[13] by submitting a preference list that optimally trades off the quality of listed schools against the probability of admission to each of those schools. This would require comparing the expected utility of submitting each preference list; for an average-sized school market this would be a daunting task for a parent. A parent living in London, with a maximum of six preferences to list, and having 15 nearby schools to choose from, would have to consider approximately 3.6 million different ways of filling in the preference list[14].

[13] We do not assume here that parents are maximising expected utility (the standard definition of a rational decision-maker in economics and decision theory) or that maximising the expected quality, measured by one observable quantity, is equivalent to maximising expected utility, or that parents' subjective probabilities of admission coincide with our estimates. Instead, we use the maximisation of a single measure of quality as a benchmark against which to assess the variation in choice sets, and parents' responses to this variation.

[14] The number of ways of choosing six ordered items from 15 is 15!/(15-6)! = 3.6m . This assumes that full lists are always better than partial ones. If the parent also has to consider partial lists, the number is even greater. Fortunately, Chade and Smith (2006) show that the decision-maker can select an optimal lottery sequentially by adding, at each stage, the school that provides the largest marginal improvement in expected utility. We use this *marginal improvement algorithm* to calculate W_B for each child.

5. Parents' preferences and "strategic" choice behaviour

5.1 Descriptive evidence on choice behaviour

Figure 4 shows that demographic groups differ systematically in their use of the preference lists provided by local authorities. Minority ethnic families use a considerably larger number of preference choices, on average, averaged across all England (left) and for just London (right). For each ethnicity, the parents of children scoring in the top tercile at KS2 use more choices than those with scoring lower children. For each ethnicity type and for each tercile (i.e third) of the KS2 distribution parents of children eligible for Pupil Premium use fewer choices on average, although this variation is smaller than the variation by ethnicity and KS2.

Figure 4: Mean list length used by groups, classified by binary ethnicity, Pupil-Premium status, and KS2 attainment terciles

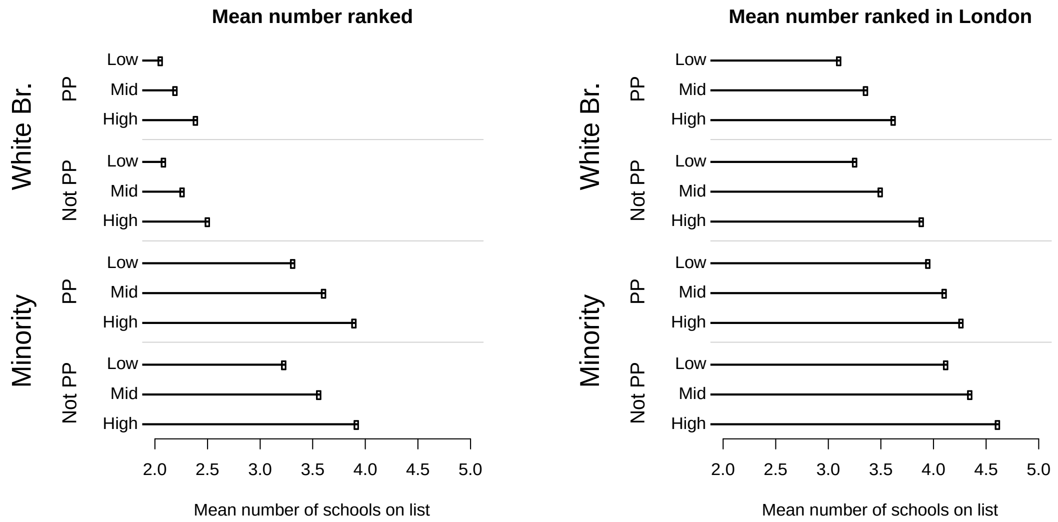

Figure 5: Mean quality (5+ A-C) of listed schools by ethnicity, Pupil-Premium status, and KS2 tercile.*

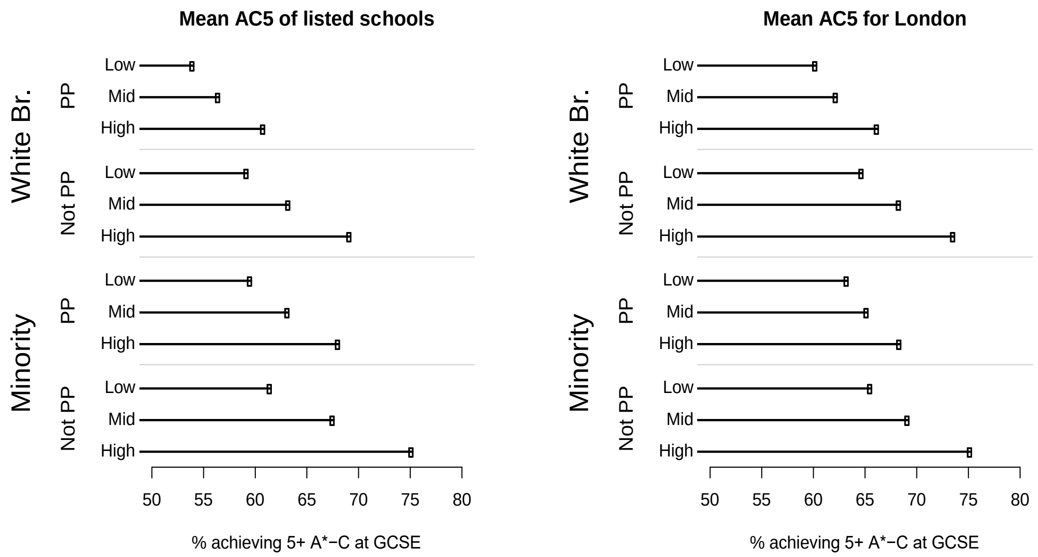

According to Figure 5, the average quality of listed schools reflects the same pattern. If minority ethnic parents are ranking higher-quality, presumably more popular, schools, this is part of the explanation of the difference in proportions achieving their first-choice school.

5.2 Discrete choice evidence

The aim of discrete choice analysis is to model parental decision-making, between mutually exclusive alternative schools, when faced with a trade-off between proximity (for example, avoiding large home-school commutes) and school performance.

Tabulating the proportions of parents who choose different schools does not account for the uneven spatial distribution of pupils and schools. But, by using a statistical model, we can account for these spatial irregularities and uncover patterns in the choices that parents make, given the options that they face.

The main focus of the current analysis is to produce estimates of parents' Willingness to Travel (WTT) for improvements in academic performance, and compare these estimates across demographic groups and across locations. Thus, this way we can directly quantify any geographic or demographic variation in parental engagement with school choice.

The estimated willingnesses to travel will be used in the evaluations of choice sets to predict the 'best' school in each family's choice set given its estimated preferences, and the 'best' school they can expect to achieve after taking into account the chances of admission.

Previous work to model parental school preferences is limited, because of the lack of availability of choice data. In the UK, Burgess et al. (2015) used the Millenium Cohort Study to conduct a discrete choice analysis of primary school choices. In the US, Hastings et al. (2005), and Abdulkadiroğlu et al. (2015) have also used choice data to study the determinants of preferences. The present study uses a similar methodology to these previous studies. However, the data at our disposal is of a scale and quality that has not been used before in the UK or the US, and consequently much richer models may be estimated.

5.2.1 The model

A discrete choice model shares the same basic structure as a regression model. There is a dependent variable, in this case a binary variable indicating whether a pupil chose a particular school, and a set of independent (predictor) variables whose influence upon the dependent variable is being modelled[15]. The convention in discrete choice modelling is to conceptualise the regression model as predicting a latent (unobserved) variable - called utility. Each family is assumed to have a separate utility value for each school, which is known to the family but unobserved by the researcher. The family simply chooses the school which gives it the greatest amount of utility. A family's utility for a school is a function of the explanatory variables – observed variables relating to the school's characteristics, the family's/child's characteristics, and interactions of the two – plus a random idiosyncratic element that captures the effect of unobserved heterogeneity between parents (akin to the error term in a regression). Such models are therefore sometimes known as random utility models in the literature.

[15] Since the dependent variable is a binary variable, and not a continuous variable, an extra stage in the modelling is needed - to transform the continuous linear predictions of the regression model into statements about the probability of the dependent variable taking a value of one, which corresponds to the family choosing the school, and zero otherwise. A purely linear model would not be constrained to predict the probability of choosing a school to be between 0 and 1. However a model that specifies the log of the odds ratio as linear in the explanatory variables is know as a Logit model and the corresponding version of this, use here, that considers many possible choices is the multinomial logit model.

5.2.2 *The Convenience vs Performance trade-off in school choice*

The estimated coefficients of a multinomial logit discrete choice model can be interpreted as log odds ratios, relating the odds of a pupil choosing a given school when a particular variable in the model is at a particular value, to the odds of choosing a school after a unit change in that variable. However, this interpretation is not intuitively easy to use. Additionally, as the number of demographic groups for whom the coefficients are separately estimated increases, the interpretation of interactions between variables and pupil group identifiers becomes increasingly complicated.

We require a straightforward representation that relates the estimated choice model to parental decision-making, so that the relative strength of demand for academic performance is revealed. Typically, in discrete choice contexts where money is involved, coefficients can be converted into monetary values, known as "willingness to pay". In the context of school choice no money changes hands, so we use distance, rather than money, as a yardstick against which to measure intensity of preference for test scores. The idea is that parents may face a trade-off between seeking a school with high academic performance, and the convenience of settling for a closer school. Note that, for any given family, no trade-off between distance and performance may be necessary – the best school may also be the closest – but the model allows the estimation of trade-offs that parents would make, if they had to. A parent's Willingness to Travel (WTT) can be interpreted as an approximation to the additional distance that a parent would be "willing" (willingness here meant in an economic sense as encompassing both willingness and ability) to travel for a 10 point improvement in test scores (say, from 70% achieving 5+ A*-C at GCSE to 80%).

5.2.3 *Non-randomness of home location*

In interpreting the results of the choice model, and especially in interpreting estimated willingness to travel, the non-random assignment of pupils to home locations, arising from their choice of home location being motivated, in part, by the proximity to good schools, must be considered. The locations of family homes are not completely independent of the quality of the local school market, since families choose where to live partly on the basis of the quality of the local school market. We do not model this location choice process. It is possible that part of a family's overall Willingness to Travel for academic quality has been subsumed within previous residential moves, for example by compromising between the child school commute and parental work commutes in choosing where to live. This is, of course, mediated by house prices, which is likely to depend, in part, on demand for nearby schools.

If there are systematic variations in the ways that different demographic groups choose where to live, whether because of differences in ability to pay or for any other reason, the estimated Willingness to Travel may reflect these differences, rather than underlying attitudes to schooling. For example, a demographic group may appear to be less willing to travel whereas, in fact, they are more likely to choose to live in an area that is already near to preferred schools. What this means is that we can only interpret Willingness to Travel relative to the residential distribution of demographic groups. This is a limitation of all existing discrete choice estimation of parents' preferences for schools (see, for example Hastings, Kane, and Staiger 2005; Deming et al. 2014; Burgess et al. 2015). By ignoring the potential endogeneity of location we rely on an assumption of "selection-on-observables", whereby parents choose where to live based on observable amenities whose effects we can allow for by including them in the explanatory variables in our discrete choice model.

This is a limitation of all existing discrete choice estimation of parents' preferences for schools (see, for example Hastings, Kane, and Staiger 2005; Deming et al. 2014; Burgess et al. 2015). We impose an assumption of selection-on-observables, whereby parents choose where to live based on *observable* amenities which we can adjust for by including them in the discrete choice model. Thus, our estimates should be interpreted as the determinants of school choice *conditional on* location.

Unlike previous work, our estimation method yields not just willingness to "pay" for quality – the trade-off that parents appear willing to make in order to access better quality schooling. But it also provides estimates of the probability that a parent will "strategise" – ie make choices based on the expected utility of the listed choices, rather than selecting purely on the basis of the most attractive L schools, where L is the number of choices that parents are allowed to list. This estimated probability of being an expected-utility maximiser rather than a simple utility maximiser allows us to compute the welfare implications of the distribution of school quality.

5.3 Discrete choice results

In order to estimate ρ we need data on the distances to all schools, the qualities of schools, their **Z**'s, each parent's ranking of schools, and the probability that i will get admitted to s based on the admission criteria (distance, and whether i has an older sibling in s) and on P_{is}.

The estimation method involves choosing a value for ρ (and one for β and for **γ**) computing the rankings of schools predicted by that value for each and every parent, and then compare the resulting rankings with the actual rankings of parents. Then, by varying the "guess" for ρ, we can find the best "fitting" value of ρ (and, similarly, for β and **γ**), which is those that best explain the observed rankings across parents. Of course, no value of ρ (together with corresponding β and **γ**) will completely match every parents' ranking because there is unobservable variation across i due to the distribution of ξ_i - which represents the idiosyncratic elements of preferences that we cannot measure. This estimation process might sound laborious but it is not computationally very intensive even for such a huge dataset in the case of the multinomial logit model.

We estimate the parameters separately for each group, where the groups are divided according to FSM=0 or not, ethnicity being white British or not, and KS2 being in the top, middle, or bottom third of the national distribution of KS2. That is, the data is divided into 12 cells according to these characteristics. Each cell is sufficiently large that we can obtain really quite statistically precise estimates of the parameters. In Figure 6 we present the estimates of ρ, for each cell, together with the corresponding 95% confidence intervals that describe how precise each estimate is. Since the dataset is very large, these confidence intervals are quite narrow so that even small differences tend to be statistically significant.

The estimates of ρ shown in Figure 6 suggest that, for given Q, No-FSM, non-white, high KS2 children are prepared to travel much further (almost 1.4km each way), compared to white British, low KS2, FSM children (who are only prepared to travel an addition 0.4km for the same increase in Q). The estimates are surprisingly consistent across groups: for each of the four combinations of ethnicity and FSM, the higher is the KS2 of the child the further he is willing to travel to get an extra unit of Q; and for each KS2 cell non--FSM children are willing to travel further than the FSM, and the minority children are willing to travel further than the white British. The differences in the estimates of ρ shown in Figure 6 are sufficiently precisely

Figure 6 Estimated ρ (willingness to travel for greater quality) parameters for each group

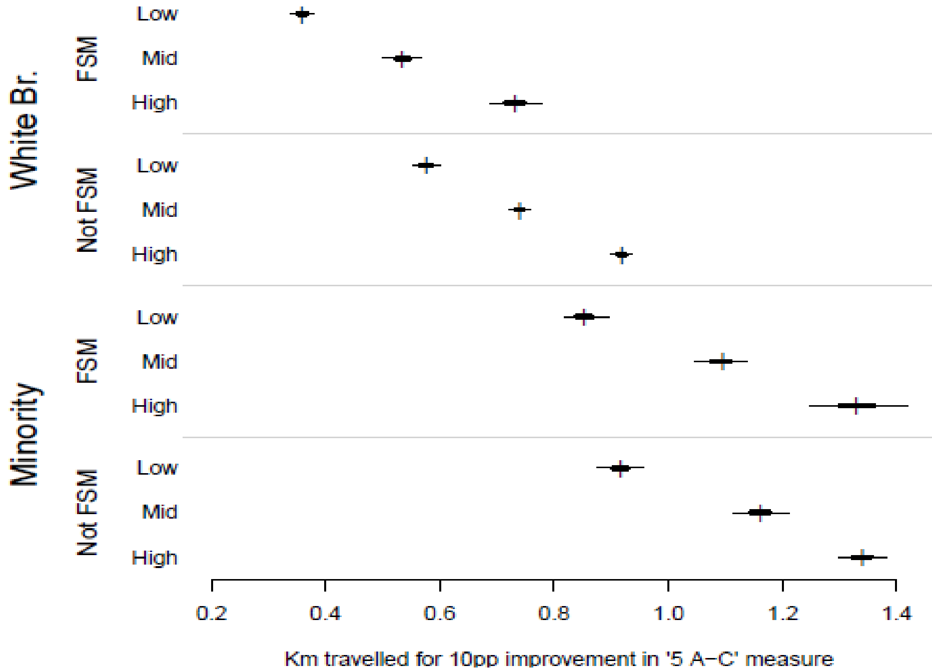

Notes: The 12 demographic groups are cross-classified by ethnicity (white British, Minority); Pupil Premium eligibility; and terciles of KS2 (end of primary school test) attainment (high, medium, low). Willingness to Travel is evaluated at 2.5 km. For each estimate, the vertical line is the point estimate, while the thick bars show the interquartile range, and the thin bars show the 95% confidence intervals of the estimate.

Figure 7 Estimated β (risk aversion) parameter for each group

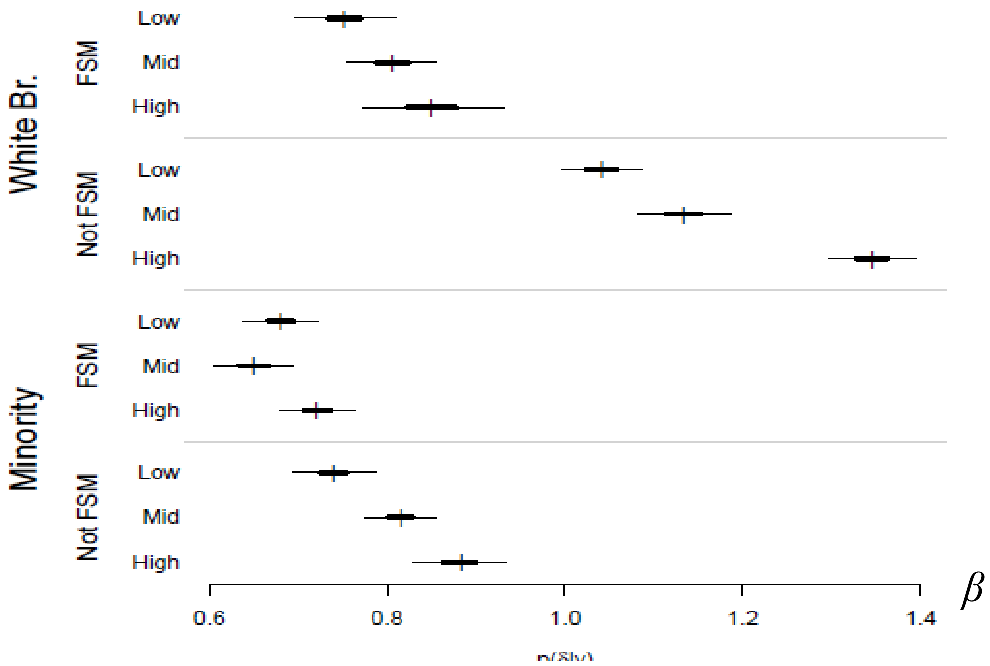

Notes: The 12 demographic groups are cross-classified by ethnicity (white British, Minority); Pupil Premium eligibility; and terciles of KS2 (end of primary school test) attainment (high, medium, low). Willingness to Travel is evaluated at 2.5 km. For each estimate, the vertical line is the point estimate, while the thick bars show the interquartile range, and the thin bars show the 95% confidence intervals of the estimate.

estimated that they are statistically significantly different from each other. For example, the differences across KS2 groups for each if the 4 cells of FSM/white are significantly different. And, for given KS2 group, the differences between each of the four FSM/white cells are significantly different from each other, at the 95% level of confidence, except for the top KS3 comparisons between No-FSM and FSM non-white groups.

Dividing the data in this way into 12 cells seems to be a useful way of capturing the differences in preferences according to some of the principal observable variables in the data.

We also visualise, in Figure 7, the estimates of the strategizing parameter, β, that captures the risk aversion of the 12 different types of parents. These β estimates are slightly less precise than our ρ estimates in Figure 3, but similar patterns emerge. Risk aversion is measured along the horizontal axis in this case. Thus, Non-FSM / High-KS2 / white British parents are more cautious then otherwise similar non-white parents.

In interpreting the results of the choice model, and especially in interpreting the estimated "willingnesses to travel", we take no view of the source of the variation in preferences. There are a variety of reasons why these may vary across individual families within groups and across groups. A number of factors come to mind. The return on investment in education through school quality might be higher for minority children – foresighted parents may wish to compensate for the disadvantage that they feel that their children might experience in the labour market and respond to this by placing a greater weight on school quality in their preferences. Even if they do not, minority parents might face tighter credit market constraints than other parents because of their lower average wealth – this makes it more difficult, at any given income, to raise a mortgage large enough to finance a location more favourable to school quality. In response they locate further away and encourage their children to travel further. The locations of family homes are not completely independent of the quality of the local school market, since families choose where to live partly on the basis of the quality of the local school market. We do not model this location choice process. It is possible that part of a family's overall willingness to travel for academic quality has been subsumed within previous residential moves, for example by compromising between the school commute and work commutes in choosing where to live. This is, of course, mediated by house prices, which may differ according to demand for, and supply of capacity in, nearby schools.

If there are systematic variations in the ways that different demographic groups choose where to live, whether because of differences in ability to pay or for any other reason, the estimated Willingness to Travel may reflect these differences, rather than underlying attitudes to schooling or to the well-being of children. For example, a demographic group may appear to be less willing to travel, whereas in fact they are more likely to choose to live in an area that is already near to preferred schools. What this means is that we can only interpret willingness to travel relative to the residential distribution of demographic groups. With a richer dataset one might be able to capture more of the heterogeneity in preferences by dividing the data into more finely differentiated cells. For the moment we think of our estimates as relevant to the parents in our data and those might vary by age and/or by cohort so that they might not apply to the decision-making made on behalf of subsequent cohorts of children.

6 School admissions and equity of access

In this section we explore the relationships between geography, demographics, and access to high-performing schools. Subsection 6.1 begins by reviewing evidence from the raw administrative records, and reveals striking variation in the probability of accessing the first-choice school by ethnic group. We then present results in subsection 6.2 from the choice-set decomposition, showing that much of this apparent inequality of access is explained by different demographic groups making more ambitious first choices, although some inequality of access remains.

6.1 Descriptive evidence on admissions

Using the newly available national preferences data, Weldon (2018) documented that minority ethnic families are much less likely to be admitted to their first-choice school than white British families. This pattern does not disappear when socio-economic status (measured by Free-school-meals eligibility) and prior attainment are accounted for, and does not appear to be an artefact of the geographical distributions of each group.

Figure 8 shows that within each ethnic group, Londoners (red) are much less likely to be admitted to their first-choice school than England as a whole (grey) . Minority ethnic families are, on average, 17% less likely to achieve their first-choice school than white British, and the gap is larger for black than Asian or Other, and this pattern persists when looking only at London. Figure 9 shows the breakdown of this

Figure 9 shows how the first-choice admission proportion varies across groups - outside London, and within London. The proportion, for all groups, is only slightly lower in London than in the rest of the UK. However, across the country, the proportion is lower for minorities than whites. And there is a consistent gradient with respect to the KS2 terciles (labelled low, medium, high) – the most able have a *lower* proportion, and the least able *a higher* proportion, than the middle group, irrespective of ethnic group and Pupil Premium group. The pupil premium effect itself is very small. Our overall impression from Figure 9 is that there is little, if any, evidence of cherry-picking by schools.

To explore ethnicity differences in more depth, Figure 10 shows the proportion admitted to their first-choice school broken down by school category, ethnicity and distance quintile. The gap between ethnic minority groups and the white British majority in the chances of admission widens with distance from the school, as one would expect. The extent to which this happens is likely to be lower for richer families who are better able to locate near their first-choice school. Of particular interest is the slope for community schools, as we know that these schools have no scope for admissions rules that might disadvantage minority ethnic families conditional on distance. That the pattern in this type of school is shared with Academies suggests that the differences are largely caused by different choice preferences and strategies, rather than cream-skimming by particular types of school.

Figure 8: Proportion admitted to first choice school by ethnicity in England and London

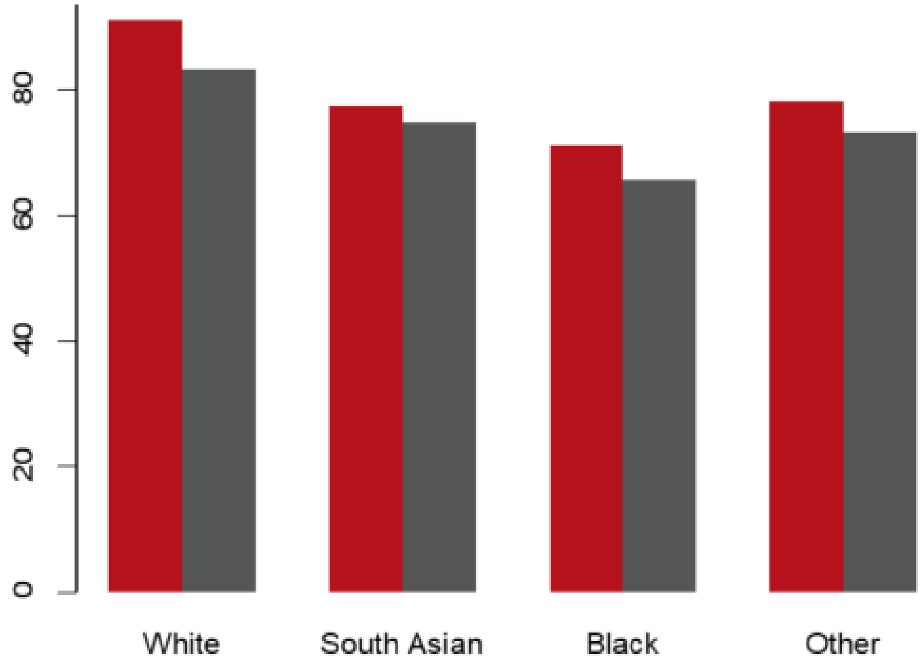

Figure 9: Proportion admitted to first-choice school by binary ethnicity, Pupil-Premium status, and KS2 attainment tercile.

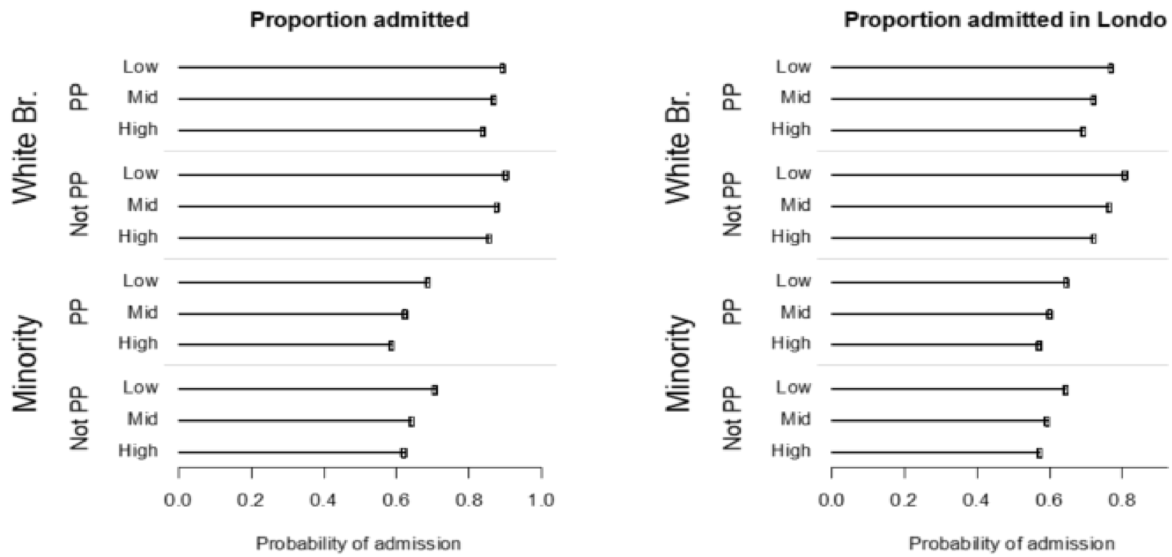

Figure 10: Proportion admitted to first-choice school by school type, ethnicity, and distance(quintiles)

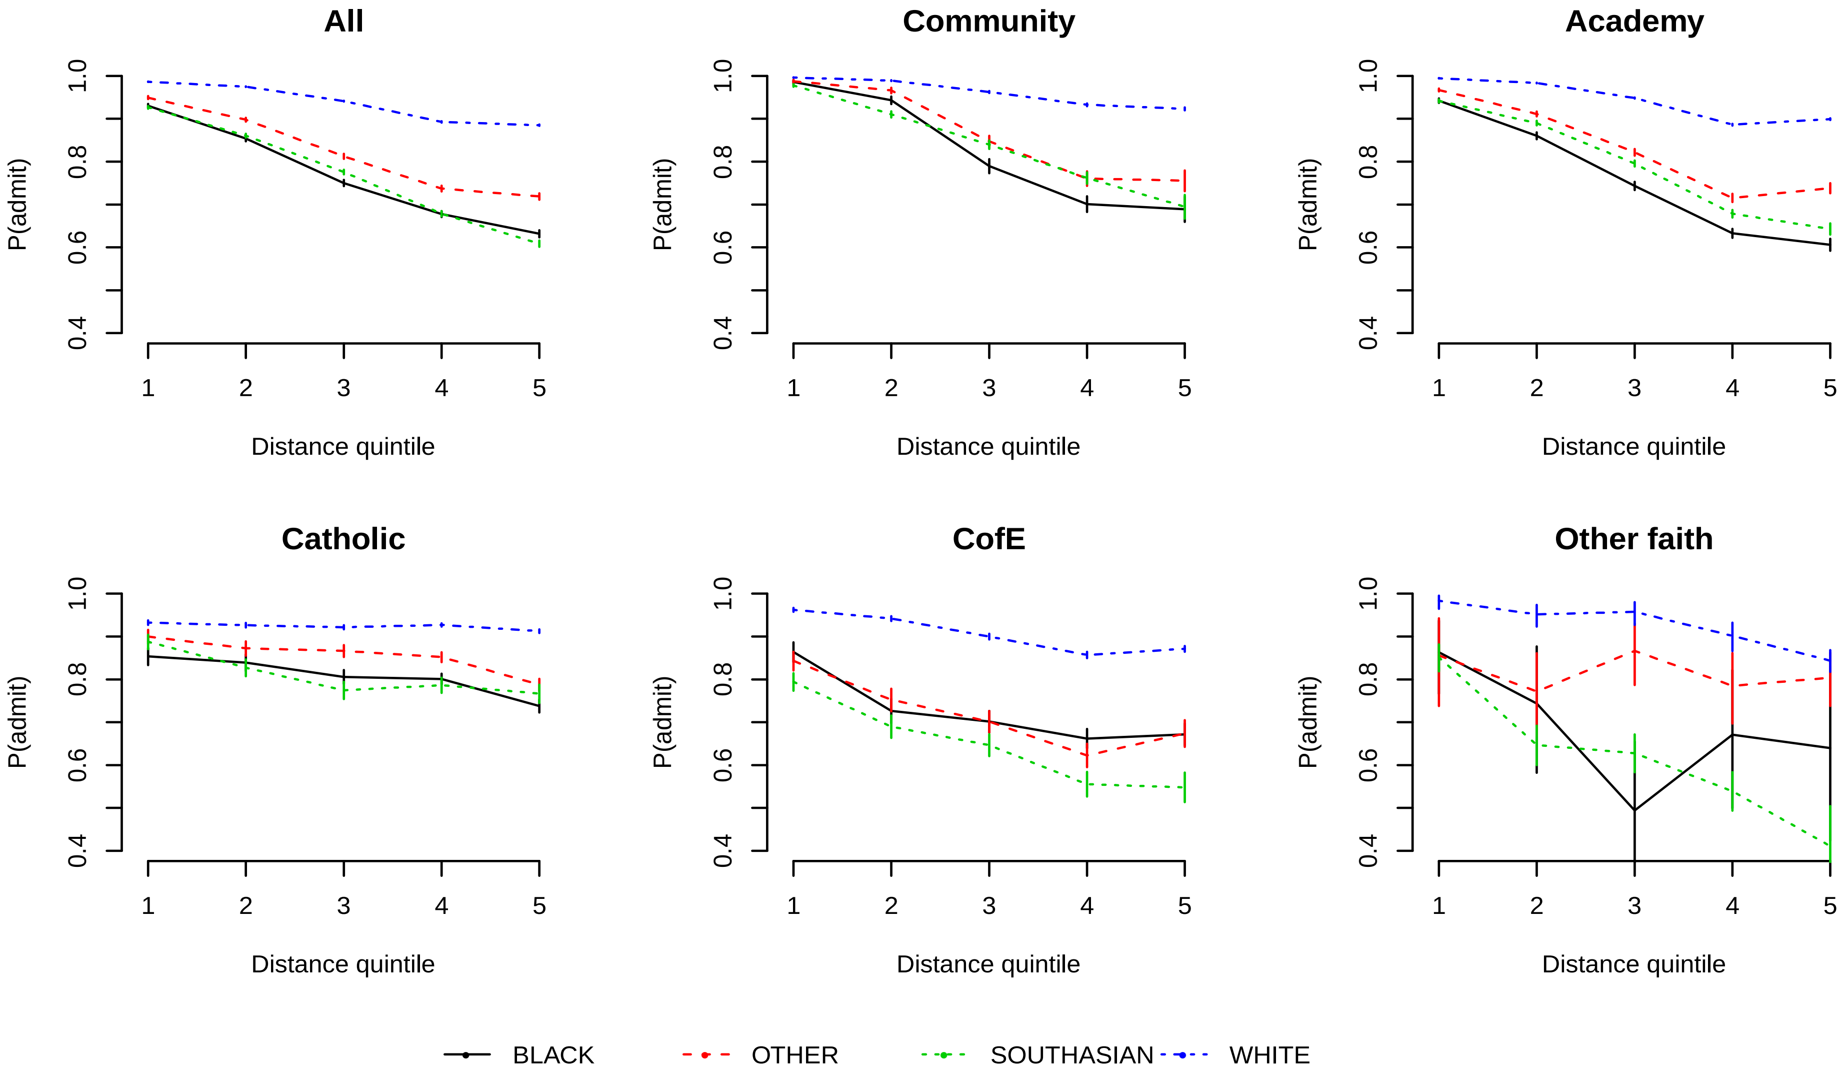

6.2 Results of the welfare analysis

In this section we present results from applying the three definitions of choice set quality to the Autumn 2014 cohort of secondary school entrants. The aim of this analysis is to evaluate whether variation in choice behaviour and chances of admission for different demographic groups results in systematic differences in the quality of outcomes from the admission process.

We compute utility under our three definitions: definition A is the best school that could be achieved given the list that the family submits; B is what is achievable given the list submitted which takes account of the constraints associated with the limited list length ; and C is what is actually achieved.

Welfare according to definition C, call this W_C, can be measured directly from observable data – we simply need to know Q and log D for that school and we can compute utility accordingly. Definition W_A is model-based, it picks one school from each family's choice set based on estimated willingness to travel for school quality[16]. This is the way in which the quality of choice has been measured in the past. Definition W_B reflects what parents can expect to experience, given both the local distribution of schools, *and* the constraints in gaining admission to those schools. Our aim is to decompose the difference in welfare between A and C into the difference that is due to constraints, ie A-B and the difference that is due to preferences, B-C. The definitions of A, B and C allow us to decompose the contributions of admission constraints, and preferences in the following way:

- The difference between W_A and W_B for a given family is the welfare foregone due to admission constraints – if all schools were undersubscribed then these would be equal and there would be no difference. We think of this as the loss in welfare associated with insufficient capacity.
- The difference between W_B and W_C is the welfare forgone due to parents' idiosyncratic preferences, which may trade-off observable for *unobservable* amenities and will, in addition, include any sub-optimal choices due to imperfect information about performance admission chances. We think of this as the role of preferences.

Figure 11 shows how the welfare difference between the ideal school and the actual school attended decomposes for groups defined by ethnicity. The gap is largest for black families at a 1.7 (equivalent to having to travel almost an additional 6-fold difference in distance – close to 15 kms more)! This is followed by other ethnicities, then by south-asian. Finally, the gap for white families is much smaller at 0.7 (equivalent to a 2-fold increase in distance – or 5km).

The breakdown of these large welfare losses, based our estimates, is that for the minority ethnic group the loss due to constraints is approximately ½ which is equivalent to a 50% rise in distance – or 1.3km. The rest of the loss is due to preferences. In the case of white families only about 1/3rd of the overall welfare loss is due to constraints – equivalent to about 0.4km. It is clear that the school choice system that is necessitated by the scarcity of capacity in good schools, imposes a much larger adverse impact on non-white groups compared to white groups.

[16] The discrete choice model is estimated separately for each of the 12 demographic groups defined by cross-classifications of ethnicity, pupil premium and KS2 attainment. We predict choices using the estimates of willingness to travel specific to each type.

Figure 11: Welfare gaps between 'Best' and 'Achievable' (A - B) and between 'Achievable' and 'Achieved' (B - C) by ethnic group.

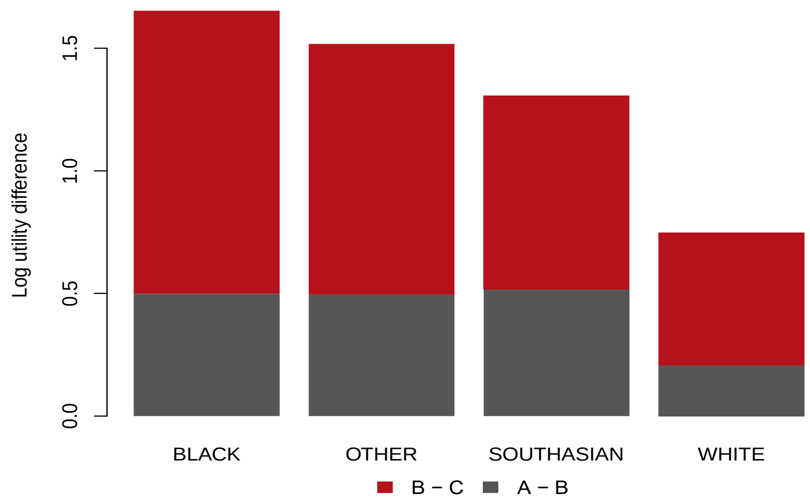

Notes: Utilities are scaled such that log utility differences correspond to log ratios of distance equivalents. For example, a log utility difference of 0.69 is equivalent to the utility loss from travelling twice as far. The difference (B – C) incorporates utility loss due to sub-optimal choice, but also incorporates preferences for unobservable amenities. The difference (A – B) incorporates the potential gains from redistributing existing capacity to eliminate capacity constraints.

Figure 12 slices the results by area type. Families living in major urban areas are much worse off, than those in non-major urban areas, who in turn experience slightly bigger loses than rural area families. Figure 13 slice the results by family income by comparing those who are Pupil Premium eligible and those who are not. On average the loss is close to a unit change in welfare (equivalent to almost a 3-fold rise in travel) – somewhat larger for the PP families compared to non-PP. In both cases the loss due to constraints is approximately 0.3 – equivalent to around a 50% rise in travel distance.

Finally, the heat map in Figure 14 contrasts the distributions of expected school quality across LA's under our three alternative definitions. Bluer indicates higher proportions achieving 5+ GCSEs, redder means lower proportions – than the average (of 63%). Definition C is based on the observed schools and is the one that is typically used to think about the geographical differences in education quality. East Anglia, the SW, parts of the East Midlands and Cumbria are the cold spots, and hotspots are largely limited to the Home Counties. Definition B is what *could* be achieved, given the capacity constraints. Many overall improvements are possible by allocating children to schools using our estimated preferences for them. Shropshire, Lancashire, and Kent become hotter. Definition A shows what might be achieved if only capacity constraints were relaxed. We see dark blue emerging for much more of the SE including Kent, Lincolnshire and improvements across Cheshire, Lancashire and Yorkshire. Our estimates suggest that constraints matter a lot, but differentially across areas.

The school choice system is necessitated by the scarcity of capacity in good schools. The constraints on capacity impose a much larger adverse impact on ethnic minority groups, those living in major urban areas, and families in receipt of pupil premium. These constraints matter a lot, but the effect is felt differentially across different areas of the country. Moving from C to B in Figure 10 suggests that preferences matter, the move from B to A shows that capacity matters too.

Figure 12: Welfare gaps between 'Best' and 'Achievable' (A - B) and between 'Achievable' and 'Achieved' (B - C) by major-urban/urban/rural types

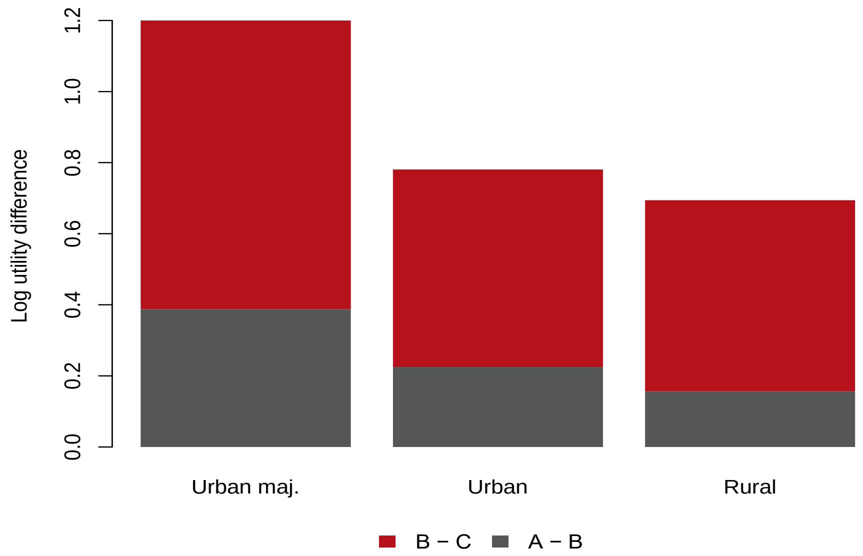

Notes: Utilities are scaled such that log utility differences correspond to log ratios of distance equivalents. For example, a log utility difference of 0.69 is equivalent to the utility loss from travelling twice as far. The difference (B – C) incorporates utility loss due to sub-optimal choice, but also incorporates preferences for unobservable amenities. The difference (A – B) incorporates the potential gains from redistributing existing capacity to eliminate capacity constraints.

Figure 13: Welfare gaps between 'Best' and 'Achievable' (A - B) and between Achievable' and 'Achieved' (B - C) by Pupil Premium

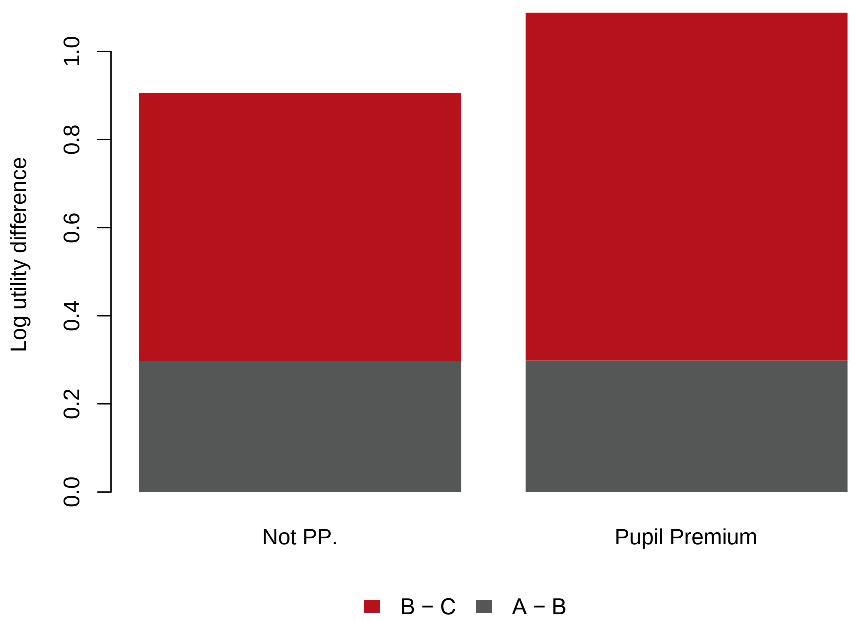

Notes: Utilities are scaled such that log utility differences correspond to log ratios of distance equivalents. For example, a log utility difference of 0.69 is equivalent to the utility loss from travelling twice as far. The difference (B – C) incorporates utility loss due to sub-optimal choice, but also incorporates preferences for unobservable amenities. The difference (A – B) incorporates the potential gains from redistributing existing capacity to eliminate capacity constraints.

Figure 14: Expected school qualities by Local authority (definitions A, B, C)

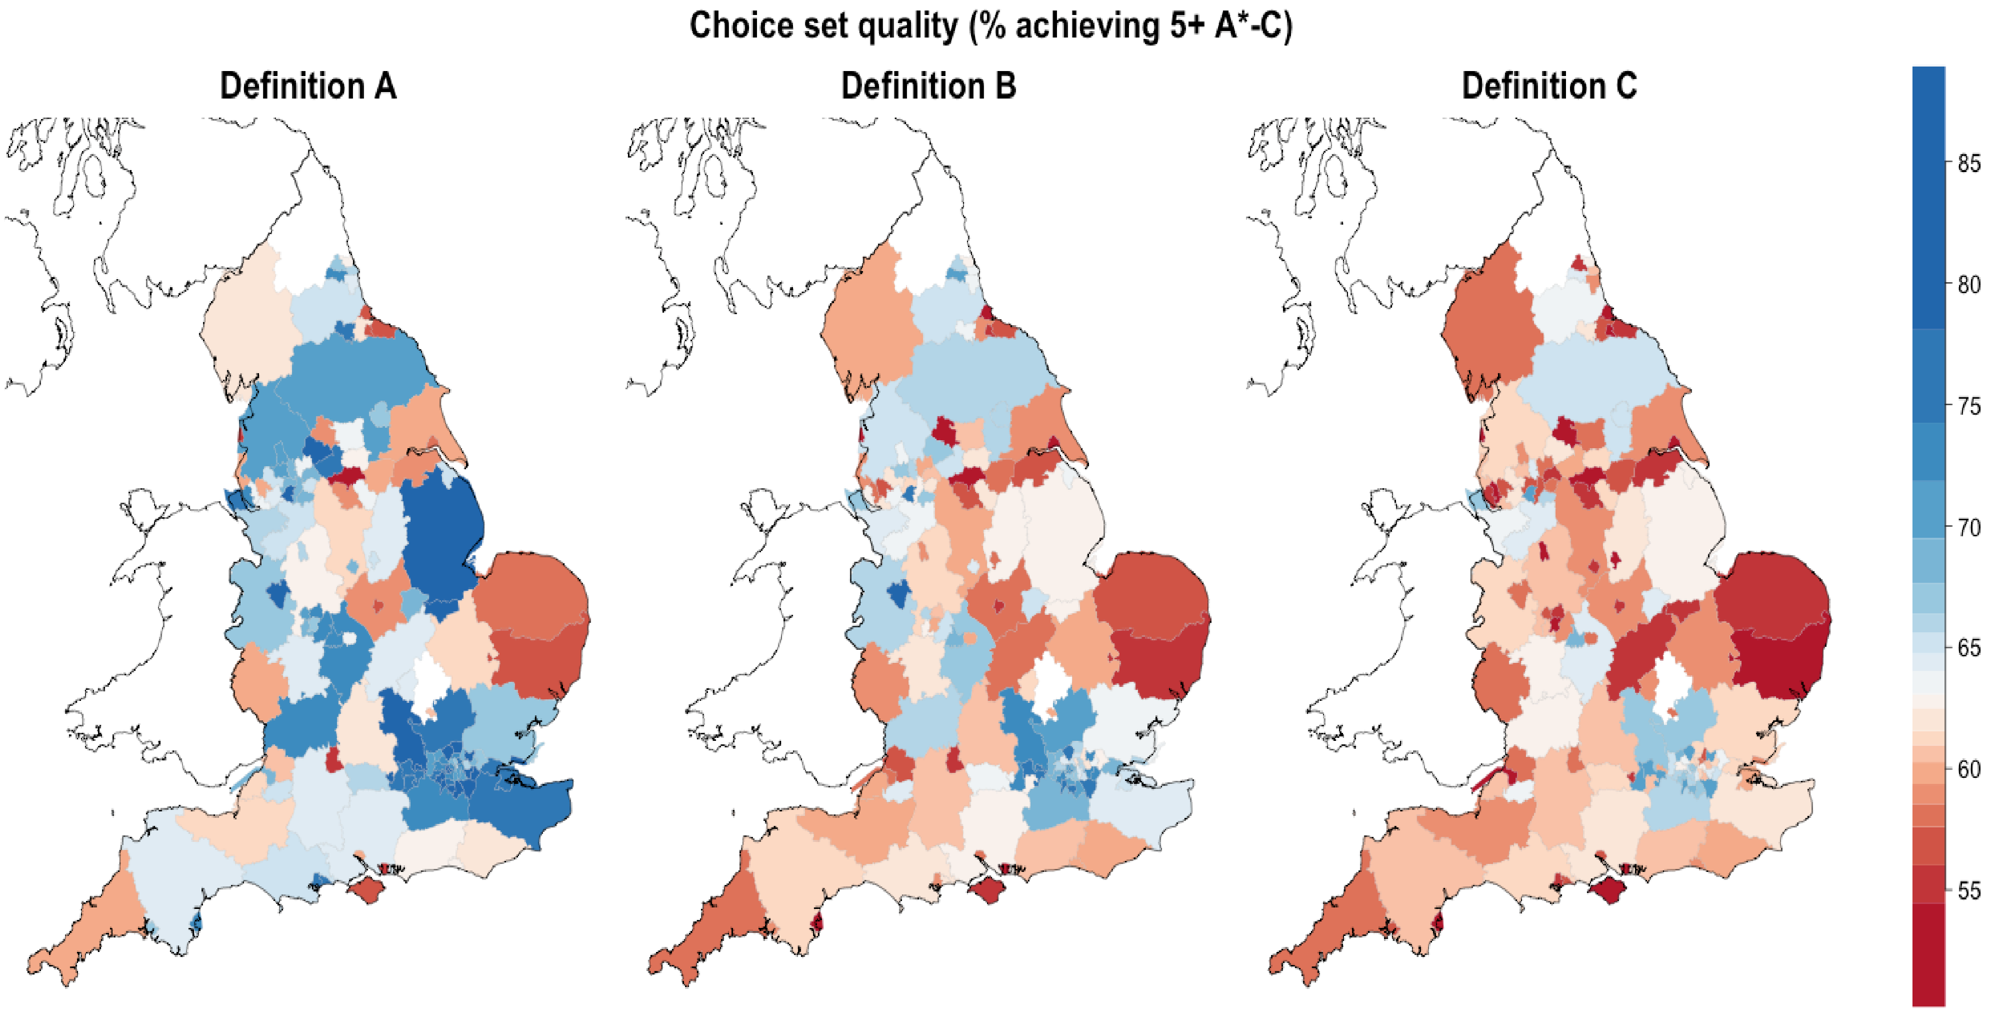

6.3 Policy application: could increasing list lengths improve welfare?

The choice set quality measures derived here can be used to evaluate the effects of hypothetical reforms to school choice. The advantage of our 'Definition B' measure over other measures is that it is sensitive to factors affecting the probability of admission to a school, such as the school's capacity, admission oversubscription rules and the design of the admission market. To give an example of the application of this measure, we evaluate a hypothetical reform, of increasing the maximum allowed length of rank-order lists that parents can choose.

The purpose of this hypothetical exercise is to demonstrate that our estimates could be used to simulate the effects of reforms – in principal any reforms that work through the admission mechanism. We can simulate the effects of reforms on how children get segregated across schools, and hence the mix of children within schools. Here, however, we focus only on the welfare implications associated with the specific reform that increases list lengths. This particular example is interesting in that it has an effectively zero cost of implementation. However, we expect this reform to show only modest effects on average for two important reasons.

Firstly, recall that Figure 2 showed that there are a significant minority r of LA's (outside of London) where the number of parents who are using ALL of their available choices– we think of this as showing that there is an appetite for greater choice, especially in those areas. That is, we would not expect much welfare improvement in increasing choice in LAs where there is no appetite for it. Secondly, the analytical data were collected at a time (2013, 2014) when the size of the age 11 cohort was at a historically low point. That is, there was more choice available at that time than has been the case for quite a while. If we were able to implement this exercise for the latest cohort then we might expect to find large overall effects since school capacity constraints are now more important and there are now more schools that are oversubscribed.

As discussed in Section 2, LA's vary in the number of schools they permit families to rank. In London, for example, all boroughs allow parents to rank up to six schools on the application form. Many other LA's allow parents to rank a maximum of only three schools. We imagine that these constraints affect the efficiency of choices in two ways. First, if parents were to continue to rank schools according to their true preferences, the constrained list would reduce the probability of children being allocated to one of their parents' top preferences. This is because if the child is not allocated to any school on the list, the child will be allocated to a nearby school with spare capacity. This may, therefore, increase the probability of the child being allocated to a school much further down the parents' *true* list of preferences. Secondly, parents may compensate for the constrained list by adjusting their *stated* preferences, to make them less ambitious. This would compensate for some of the loss of efficiency caused by the list length constraint, but would compromise the mechanism's *strategy-proofness* property.

Table 1 shows the results of a counterfactual exercise, calculating the expected quality (definition B) for children, based on the actual list size, and then based on scenarios where the list size is increased to ten[17]. Increasing the list length to ten schools improves expected quality (using the expected value of the 5 GCSE A*-C measure), but only by 0.15 percentage

[17] We also increased the list size to twenty, but additional welfare gains were negligible.

points. However, for minority ethnicities the measure increases by 0.2 percentage points or more. This difference is likely due to the fact that minority ethnic families tend to have lower average probabilities of admission into their first-choice schools. This, in turn, may be partly due to being more likely to live in urban areas that are near to over-capacity schools, but not near enough to be admitted, and partly due to them being less likely to have priority at faith schools.

Table 2 shows the same analysis broken down by terciles of prior attainment at primary school. The gains in test scores of the allocated school are disproportionately made by those in the top tercile of the KS2 distribution. It seems that the highest ability children are those whose potential is most constrained by the limited list lengths.

In summary, the list length reform had variable effects, but in almost all cases it either left the expected allocation unchanged or increased the expected test scores and distance. In most LA's the average increase was less than 0.2 percentage points for test scores (Q), and by a tint proportional difference in D (less than 20 metres).

Table 1: Expected % with 5+ GCSEss of allocations (definition B) with different counterfactual maximum list sizes, by ethnic group

	WHITE	BLACK	OTHER	SOUTH ASIAN	All
Status quo	61.76	65.03	66.25	64.64	62.75
List 10 schools	61.92	65.22	66.46	64.85	62.87
% Change in W_B	0.15	0.20	0.22	0.20	0.12

Notes: Definition B denotes the expected quality (in % achieving 5+ A*-C) that can be achieved with a given list length and a given set of preferences. The bottom row shows the % differences between row 2 and row 1. All units are percentage points

Table 2: Expected % with %+ GCSEs of allocations (definition B) with different counterfactual maximum list sizes, by KS2 attainment tercile

	Bottom 3rd	Middle 3rd	Top 3rd	All
Status quo	58.50	63.34	70.37	62.75
List 10 schools	58.58	63.51	70.59	62.87
% Change in W_B	0.08	0.17	0.22	0.12

Notes: Definition B denotes the expected quality (in % achieving 5+ A*-C) that can be achieved with a given list length and a given set of preferences. The bottom row shows the % difference between row 2 and row 1. All units are percentage points.

7 Policy implications

In this section, we discuss four policy themes suggested by the analysis: list length; admission mechanisms; choice architectures; and using the analysis here as a step towards incorporating school choice into wider education policy issues – starting with the relationship with social mobility.

7.1 List length

A discussion of policy recommendations from our research must begin with what we consider to be the easiest reform to implement: increasing the allowed number of preference slots for each parent. Currently, around half of LA's allow parents to rank no more than three schools. Although in some authorities this appears to be sufficient for the majority of parents, in other local authorities it is likely to not be enough – and this forces parents to make difficult strategic decisions about which schools to include in their three choices.

In London, on the other hand, all families are allowed to rank up to six schools, but it is not clear whether even this number is sufficient, and our analysis suggests that some groups of parents would have a greater chance of achieving a good allocation for their children if they were allowed to rank more schools. An increase in the maximum size of lists to, say, ten schools, would only affect those parents for whom the existing limit is binding; and for parents who currently do not fill their lists the additional slots would make no difference.

It may be argued that increasing list lengths would have unintended psychological consequences: for example, perhaps parents would feel compelled to 'fill the space' and therefore rank schools about which they have only second-hand or vague information (say, because they have not been able to visit schools 5-10 on their list). However, this must be evaluated against the status quo, where families, if they are not admitted to their first, second, or third preferences, have no say at all in where their children will eventually be assigned. Arguably, parents should be trusted to invest 'enough' time in obtaining information for decision-making, where 'enough' should be defined broadly. Parents ranking schools about which they have only second-hand or vague information can be rational if the probability of assignment to a low-ranked school is not high enough to justify the additional effort to acquire information.

It is likely that, with more choices, parents' first choices become more ambitious, so headline statistics on the number allocated to their first preference would probably decrease, whereas the number allocated to *any* of their choices would increase. Preferences would more straightforwardly signal demand for schools, and would therefore be more useful in planning and forecasting. Additionally, to the extent that school choice creates incentives for schools to improve to attract pupils, the reform would make market demand signals more trustworthy, and possibly improve the effectiveness of the market as a mechanism for improvement.

For parents, therefore, this would be an innocuous reform. For schools, however, the additional applications might present an additional administrative burden. However, this would only affect own-admissions schools, which rank their own applications according to their own criteria. For these schools, the additional administrative burden would be proportional to the complexity and level of subjective judgement required in ranking applications. A reform might, therefore, have the beneficial side-effect of incentivising schools to make their admission rules more transparent and less subjective.

A reform to increase the length of preference lists would affect those families who currently complete their preference lists (and presumably would be able to make good use of additional choices). Since parents of minority ethnicities are more likely to rank additional schools, they are more likely to benefit from the reform. The results of our analysis bear this out: the improvement in expected quality for minority ethnic groups, from moving to ten slots, is larger than it is for white families.

7.2 Admission mechanisms

In 2007, the Admissions Code outlawed a previously popular school allocation mechanism, and forced many LAs to adopt new allocation mechanisms. According to Pathak and Sonmez, (2013) and Coldron (2007) the vast majority of local authorities adopted variants of the Deferred Acceptance algorithm. However, there does not appear to be any central oversight of the details, implementation, or effects of allocation mechanisms used by authorities.

This may be a problem, since it is unlikely that local authorities themselves possess the technical expertise to evaluate their own admission algorithms. LAs usually outsource this function, and it appears that many use admission algorithms provided as part of software packages by database vendors, such as Capita, SAP etc. While there is no reason to believe that algorithms provided by these vendors are not fit for purpose, it does raise the question of who provides oversight of these very technical mechanisms. Who validates, for example, that allocation mechanisms intended to guarantee some property such as stability or optimality, actually fulfil this criterion?

The problem here is that, as the law does not stipulate that an admission mechanism should satisfy any particular property, there is no commonly-agreed set of standards upon which central or local governments can base an evaluation of allocation mechanisms and their implementation in software.

Indeed, although strategy-proofness and stability are two properties that are often invoked in the theoretical literature on matching, there is by no means an academic consensus that school choice mechanisms should always satisfy either or both of these properties, and there are several examples of real-world allocation mechanisms that, by design, do not meet either or both of these criteria.

A promising approach to raising the profile of mechanism design in local authorities, that has been popular in the US, would be for local authorities to provide researchers with the necessary data[18] to simulate the effects of different matching mechanisms, that would facilitate deliberation within local authorities about the relative merits of different mechanisms, and the trade-offs between the properties of these mechanisms.

7.3 Choice architectures: information and nudge

Inspection of the existing choice architectures suggest that the advice provided to applicants is very limited. Many seem to encourage conservatism in applicants - to strategize by making sure that they list schools that they are likely to be admitted to. None seem to provide simple diagrammatic visualisations of the possible choices and their respective qualities and

[18] Parents' ranked preferences data is now available nationally, but schools' ranked priority data and capacities are still not routinely made available to researchers. Given that ranked preferences, linked to National Pupil Database, are arguably much more sensitive and personal than schools' oversubscription criteria, the case for sharing this data seems strong.

proximities – the simple things that we have come to expect from hotel listing websites that provide far less crucial services to the population.

Indeed, very few even seem to offer crow-fly distances – something that is not readily available online to parents that could at least approximate the tie-breaker calculations that LAs do. Providing such information could make a large difference – it could ensure that parents were well informed about the good schools where they are likely to gain admission.

Our own research could provide the basis of a recommender algorithm – indeed if the websites were linked to NPD it maybe be possible to provide customised advice, of the kind that users will be familiar with from choosing within streamed TV services, that might produce better matches at the individual level. Even in the absence of a link to NPD it would be possible, with a supplied postcode, to provide less specific advice on choices that might appear to be driven by idiosyncratic decisions – for example, an algorithm could be designed to point out that a better AND closer school was available. Or it could easily advise that, in previous years, your child would have sufficient priority to be admitted to a set of particular schools. Indeed, if NPD were used to identify older sibling, school, and postcode, it would be possible to say which school(s) one could be (almost) guaranteed a place at. This would have the effect of lessening the impact of limited list lengths on the incentive to strategize.

7.4 Towards wider educational policy: school choice and social mobility

Changing admission criteria will change who gets into which school – and hence changes each school's social mix, or "balance". One suggestion is to include a priority for disadvantaged families whereby a number of places are reserved for applicants from less well-off backgrounds, for example based on eligibility for the Pupil Premium. But many other ideas are possible. There are two channels by which admission criteria can change outcomes. First, it can change the "matches" between schools and children. For example, after some reform, a given child might get matched to a different school. This affects equality of opportunities (to attend a good school) that, in turn, affects social mobility (SM). To the extent that a reform results in lower achieving children getting matched to "better" schools we might expect it to improve SM. Secondly, even if a particular child is post-reform matched to the same school as would occur under the status-quo admission criteria there may still be an effect on outcomes. This could arise if the mix of peers at that school (its "balance") changed and there were peer effects on outcomes.

"Balance" depends on how children get distributed across schools which determines who goes to school with who – and, hence, the within school mix of peers. Balance can be defined with reference to observable characteristics such as parental background, child ability, ethnicity, parental background, gender, etc. The within school mix may affect outcomes for each child through peer effects - outcomes such as the child's educational attainment and, ultimately, her income. Therefore, the mix of children within a school contributes to the outcomes of all children within the school via peer effects. Thus, it seems natural to merge the analysis here with a model of peer effects so we can model the effect of admission criteria on the outcomes for the children via peer effects. Of course, other factors will matter for child outcomes. A model of peer effects could capture these additional considerations via school "fixed effects" (FE)– an estimated parameter for each school that reflects the school's "value added" *abstracting from the effects of the composition of pupils*, where each FE measures the effectiveness of a school in teaching those children admitted. Longitudinal Education Outcomes (LEO) would be a useful vehicle for such research.

8 Conclusion

This report summarises our research on school choice based on using the newly available secondary school choice data. This provides the raw data on the rank-order lists constructed by parents of 10 year olds destined for secondary schools. We are able to link the data to NPD information on prior ability, and subsequent secondary school attended. We use the data to model how parents, as revealed by the preferences they submit to admission authorities, trade-off school proximity against school quality and the probability of admission. The latter is an important element of our story – something that is missing from previous research. Previous research has relied heavily on the theoretical presumption that the Deferred Acceptance algorithm **is** *strategy-proof* – that is, it frees parents from having to engage in strategizing behaviour. However, the proof depends on being ale to list all schools, without limit. While its not clear how long lists have to be for the theorem to hold its seems very likely that three will be too little.

On average, we estimate that parents place a considerable weight on school performance (our proxy for quality). By observing their choices we suggest that parents are prepared to allow their child to travel an additional 0.9 km (when the mean distance is around 2.5 km) to achieve a 10 percentage point better quality school. This is a considerable burden that households seem to be willing to pay. We also find large differences around the average in the willingness to pay (i.e. travel for) for school performance across types of parents. For white British parents the willing less is around 0.8km while it is approximately 1.1km for non-white minority parents on average – a difference is highly statistically significant.

We also find considerable variation across parental types in the extent to which parents engage in making *strategic* choices – choices that reflect aversion to risk.

We explore the simplest and cheapest possible intervention – extending the length of lists that parents are able to specify. The raw data suggests that many Local Authorities restrict the ability to list sufficient schools – and this is something that causes parents to be too conservative in their choices. While this enables LAs to say that it has a high proportion of children attending their first choice – if their first choice is an unduly safe one, this is a hollow achievement.

There are many instances in the data that are consistent with the rational use of the school choice system. Further reforms could improve the choices that parents make – through better information provided by tools that are familiar to us as consumers of media, transport, and accommodation services.

In work in progress we show that reforms to admission criteria offer the possibility of manipulating the allocation of children to improve the chances of disadvantaged children to attend more effective schools and benefit from high ability peers. Additional social mobility might be obtained at minimal cost.

We view this latter activity as the likely to become the most important aspect of our research. It is now beginning to be possible to develop an equivalent to the IFS tax-benefit model – but for education policy, not tax and welfare policy. Where the levers are not tax and benefit parameters but the school admission criteria, school locations, school types and capacities; and the trade-offs are not work vs “leisure” but proximity vs quality. Linked to LEO this would be an important framework for planning policy – and an important way of holding policymakers to account.

References

Abdulkadiroglu, Atila, Nikhil Agarwal, and Parag A. Pathak. 2015. ‘The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match’. *The American Economic Review* 107(November): 3635–3689.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. ‘School Choice: A Mechanism Design Approach’. *American Economic Review* 93(3): 729–47.

Agarwal, Nikhil, and Paulo Somaini. 2018. ‘Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism’. *Econometrica* 86(2): 391–-444.

Ajayi, K. 2013. *School Choice and Educational Mobility: Lessons from Secondary School Applications in Ghana*. . Working Paper.

Ajayi, K., and M. Sidibé. 2016. *An Empirical Analysis of School Choice Under Uncertainty*. . Working Paper.

Allen, Rebecca, Simon Burgess, and Leigh McKenna. 2014. *School Performance and Parental Choice of School: Secondary Data Analysis*. Department for Education. Report.

Allen, Rebecca, and Rob Higham. 2018. ‘Quasi-Markets, School Diversity and Social Selection: Analysing the Case of Free Schools in England, Five Years On’. *London Review of Education* 16(2): 191–213.

Andrews, Jon, and Natalie Perera. 2017. *Access to High Performing Schools in England*. Education Policy Institute. Report. https://epi.org.uk/publications-and-research/access-high-performing-schools-england/.

Burgess, Simon, Ellen Greaves, and Anna Vignoles. 2019. ‘School Choice in England: Evidence from National Administrative Data’. *Oxford Review of Education* 45(5): 690–710.

Burgess, Simon, Ellen Greaves, Anna Vignoles, and Deborah Wilson. 2015. ‘What Parents Want: School Preferences and School Choice’. *The Economic Journal* 125(587): 1262–1289.

Calsamiglia, Caterina, Chao Fu, and Maia Guell. 2014. *Structural Estimation of a Model of School Choices: The Boston Mechanism vs. Its Alternatives*. . FEDEA Working Paper no. 2014-21.

Chade, H., and L. Smith. 2006. ‘Simultaneous Search’. *Econometrica* 74(5): 1293–1307.

Coldron, J. et al. 2008. *Secondary School Admissions.* Department for Children, Schools and Families.

Deming, David J., Justine S. Hastings, Thomas J. Kane, and Douglas O. Staiger. 2014. ‘School Choice, School Quality, and Postsecondary Attainment’. *American Economic Review* 104(3): 991–1013.

Department for Education and Skills. 2007. *School Admissions Code*.

Fack, G., J. Grenet, and Y. He. 2015. *Beyond Truth-Telling: Preference Estimation with Centralized School Choice*. . Working Paper.

Gale, David, and Lloyd Shapley. 1962. ‘College Admissions and the Stability of Marriage’. *The Americal Mathematical Monthly* 69: 9–15.

Gibbons, Stephen, and Stephen Machin. 2006. 'Paying for Primary Schools: Admission Constraints, School Popularity or Congestion?' *The Economic Journal* 116(510): C77–C92.

———. 2008. 'Valuing School Quality, Better Transport, and Lower Crime: Evidence from House Prices'. *Oxford Review of Economic Policy* 24(1): 99–119.

Gibbons, Stephen, Stephen Machin, and Olmo Silva. 2013. 'Valuing School Quality Using Boundary Discontinuities'. *Journal of Urban Economics* 75: 15–28.

Gorman, Emma and Walker, Ian, Heterogeneous Effects of Missing Out on a Place at a Preferred Secondary School in England. *IZA Discussion Pape*r No. 13167.

Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger. 2005. *Parental Preferences and School Competition: Evidence from a Public School Choice Program*. National Bureau of Economic Research. NBER Working Paper. http://www.nber.org/papers/w11805.

Hoxby, Caroline. 2003. 'School Choice and School Productivity: Could School Choice by a Rising Tide That Lifts All Boats?' In *The Economics of School Choice*, ed. Caroline Hoxby. Chicago, IL: University of Chicago Press.

Hunt, Emily. 2018b. 'Secondary School Choice in England'. *Education Policy Institute*. https://epi.org.uk/publications-and-research/secondary-school-choice-in-england/ (October 11, 2019).

———. 2019a. 'Fair Access to Schools? The Impact of the Appeals and Waiting List System'. *Education Policy Institute*. https://epi.org.uk/publications-and-research/appeals-and-waiting-lists/ (October 11, 2019).

———. 2019c. 'What Have We Learned from National Offer Day 2019?' *Education Policy Institute*. https://epi.org.uk/publications-and-research/national-offer-day-2019/ (October 11, 2019).

Luflade, M. 2018. *The Value of Information in Centralized School Choice Systems*. Duke University. Working paper.

Machin, Stephen, and Kjell G. Salvanes. 2016. 'Valuing School Quality via a School Choice Reform'. *The Scandinavian Journal of Economics* 118(1): 3–24.

McKendrick, Andrew and Walker, Ian, 'The Role of Faith and Faith Schooling in Educational, Economic, and Faith Outcomes', IZA Discussion Paper No. 13192.

Pathak, P. and Tayfun Sönmez. 2013. 'School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation'. *American Economic Review* 103(1): 80–106.

Weldon, M. 2018. *Secondary School Choice and Selection: Evidence from New National Preferences Data*. . Department for Education Research Report.

Wilson, Deborah, and Gary Bridge. 2019. 'School Choice and the City: Geographies of Allocation and Segregation'. *Urban Studies* 56(15): 3198–3215.