Preferences and Heterogeneous Treatment Effects in a Public School Choice Lottery

by

Justine S. Hastings
Department of Economics
Yale University and NBER

Thomas J. Kane
Graduate School of Education
Harvard University and NBER

Douglas O. Staiger
Department of Economics
Dartmouth College and NBER

May 2007

ABSTRACT

We use data from a public school choice lottery to estimate the effect of attending a first-choice school. For the average student, attending a first-choice school is not associated with improvements in test scores or other academic outcomes. However, academic achievement is only one goal families may have when choosing a school and, depending on their preferences, parents may trade-off academic achievement against other desirable school traits. We estimate the implicit weight families attached to school test scores as revealed by their choices, and test for interactions between those preferences and the impact of winning the lottery. We find that those whose parents placed high weights on school test scores in their school choices experienced significant gains in test scores. Therefore, the impact of winning the lottery on academic achievement depends upon parents' objectives when choosing schools.

I. Introduction

The recent federal No Child Left Behind Act (NCLB) includes a school choice component requiring that students at persistently under-performing public schools be given the option to choose a to attend a higher-achieving public school outside of their neighborhood. Public school choice plans are intended to improve public education for students in disadvantaged communities by offering them the immediate option to attend a higher-performing schools, and by giving their local schools a greater incentive to improve through the threat of losing students. However, there is little empirical evidence that giving disadvantaged students the opportunity to attend a higher-performing school actually raises their achievement.

A number of papers have used random assignment to estimate the effect on academic outcomes of attending a school other than the local public school. Most of these papers use either lottery assignment to over-subscribed schools in public school choice plans or randomization of private school vouchers to identify the treatment effect (Witte et. al 1995; Greene et al. 1997, Rouse 1998, Peterson et al. 1998, Mayer et al. 2002, Krueger and Zhu 2004, Cullen et al. 2006, Hastings et al. 2006b).[1] Taken together, these papers have been unable to find robust impacts on average academic outcomes. The inability to find average academic benefits, even for subgroups, has lent support to the claim that "measurable school inputs have little causal impact on student outcomes" (Cullen et al., 2006).

However, as noted in many of these papers, parents may have a variety of reasons for choosing schools. The child of an academically oriented parent who chooses to send their child across town to attend a high-performing school may experience larger gains from school choice than the parent who chooses a school primarily for convenience. Differences in the underlying reasons that parents choose schools and differences in the trade-offs they face may lead to heterogeneous impacts of attending a first choice school on academic achievement. The average effect of attending a first choice school may be zero even if schools do have an impact on student outcomes.

[1] An earlier non-experimental literature compared academic outcomes of those students who chose to attend private, charter or magnet schools to those students who remained in their neighborhood school (Coleman et al. (1982), Bryk et al. (1993), Blank (1983), Gamoran (1995), Evans and Schwab (1995), Neal (1997), Altonji et al. 2002).

To date, researchers have not had a setting in which they could both estimate the underlying determinants of school choice and the resulting outcomes from attending a first choice school. Thus, they have focused on estimating impacts for different subgroups of students who may have arguably different reasons for selecting schools (impacts by race, income or baseline academic achievement may reveal heterogeneous treatment effects caused by differences in preferences and choices across subgroups). In this paper we use the choice rankings of parents from the introduction of district-wide school choice in the Charlotte-Mecklenburg School District (CMS) along with lottery assignment of students to first choice schools to shed new light on the factors that drive parental choice and subsequent student outcomes in a public school choice plan.

We start with administrative data on parents' school choices in the first year of the CMS school choice plan. In 2002, CMS introduced a district-wide choice plan to end three decades of busing for integration. All parents in the district were asked to provide their top three school choices and a lottery was used to determine admission to oversubscribed schools. We show that parents choices are very heterogeneous, and we outline a simple model that illustrates how the academic gains from attending a first-choice school should be positive for parents who place a high implicit weight on academics when choosing a school, and potentially negative for parents who place low weights on academics and high weights on other factors that are negatively correlated with academics (e.g. geographic proximity or fraction minority). We then use the rich choice data and lottery assignment to test this model in the data, allowing the treatment effects to vary with the underlying preferences that drove parents' school choices (Heckman, Smith and Clements, 1997; Heckman, 1997; Heckman, Urzua and Vitlacil, 2006).

In order to estimate the implicit weight each parent placed on academics when choosing a school, we first estimate the distribution of preferences for school characteristics using a mixed-logit model (McFadden and Train (2000),Train (2003)).[2] This random utility model allows parents to have heterogeneous preferences for school characteristics that depend on observable demographics as well as idiosyncratic factors.

[2] For a more complete discussion of the effect of school characteristics in determining parental demand, see Hastings, Kane and Staiger (2006a).

The estimates from this model can be used to calculate a posterior of the implicit weight that each parent placed on academics when choosing a school.[3] This posterior is effectively a non-linear index of baseline characteristics; combining all of the information about the student, their proximity to each school, the characteristics of each school and the schools their parents selected into a single summary statistic. It measures how unexpected a parent's choices were relative to what a typical parent would have chosen given their demographics and geographic location. For example, if a parent persistently chose high-performing schools, passing closer schools that most other parents would have preferred, they most likely place a high implicit weight on academics.

We then turn to the lottery assignments to investigate if students experience academic gains from attending a first-choice school. We find no average gain in test scores from attending a first choice school, consistent with the prior literature. We do, however, find significant positive gains among white students and students of higher-income families and negative but insignificant impacts for African Americans and children of lower-income families. We then show that these patterns in the academic gains from attending a first-choice school across subgroups are driven by differences across subgroups in the underlying preferences and trade-offs parents face. We do this in two ways. First, we show that the average weight placed on academics is positively correlated with the average treatment effect across the subgroups focused on in the prior literature (although the subgroup impacts are generally insignificant). Second we allow the treatment effect of attending a first choice school to vary explicitly with the posterior weight that parent's place on academics, and find that the treatment effect varies significantly and positively with this weight. Students with estimated weights on school test scores in the top decile experienced significant rises in end of grade test scores of approximately 0.1 to 0.2 standard deviations. These students are most likely to be white and have higher income levels, thus generating the significant subgroup impacts by race and income. In contrast, students placing little value on academics actually experienced declines in standardized test scores. We show that this is strongest for African

[3] The multiple ranked choice responses, we are able to credibly identify idiosyncratic heterogeneity in preferences, based on systematic differences in the sequential choices made by parents of similar students facing similar school options (Berry, Levinsohn and Pakes (2004)).

Americans, who must trade-off academic gains against utility gains from attending a school with a high proportion of minority students.

An important innovation in this paper is offering an economic model that can explain heterogeneous treatment effects across subgroups of students. Our results suggest that differing impacts across subgroups can be largely explained by underlying differences in parents' willingness to trade off expected gains in academic achievement for gains in utility along other dimensions, such as proximity or school racial composition. More generally, this implies that the impact of school choice on academic outcomes will depend on both the willingness of parents to make these tradeoffs, and the extent to which the available school choices require such tradeoffs to be made.

This paper proceeds in five main sections. The first section lays the background for the data and estimation by describing key details of the CMS school choice plan. The second section outlines the relationship between expected academic outcomes and preferences in a school choice plan, where parents choose schools based on expected academic achievement and other school characteristics, and students are then granted admission to schools by lottery. In the third section we generate estimates of the preferences for academic achievement. We then incorporate these preference estimates into our final estimation of the effect of attending a first choice school on academic outcomes. The final section concludes.

II. Background: The CMS School Choice Plan

Before the introduction of a school choice plan in the fall of 2002, the Charlotte-Mecklenburg public school district (CMS) operated under a racial desegregation order for three decades. In September 2001, the U.S. Fourth Circuit Court of Appeals declared the school district "unitary" and ordered the district to dismantle the race-based student assignment plan by the beginning of the next school year. In December of 2001, the school board voted to approve a new district-wide public school choice plan.

In the spring of 2002, parents were asked to submit their top three choices of school programs for each child. Each student was assigned a "home school" in their neighborhood, often the closest school to them, and was guaranteed a seat at this school.

Magnet students were similarly guaranteed admission to continue in their current magnet programs. Admission for all other students was limited by grade-specific capacity limits set by the district. Students could choose any school in the district. However, transportation was only guaranteed to schools in a student's quadrant of the district (the district was split into 4 quadrants called 'choice zones'). The district allowed significant increases in enrollment in many schools in the first year of the school choice program in an expressed effort to give each child one of their top three choices. In the spring of 2002, the district received choice applications for approximately 105,000 of 110,000 students. Admission to over-subscribed schools was determined by a lottery system as described below.

Once the district was declared "unitary" and the court order requiring race-based busing was terminated, they could no longer draw boundaries based on the racial composition of a neighborhood. As a result, the former school assignment zones, which often paired non-contiguous black and white neighborhoods, were dramatically redrawn. Under the choice plan, 43 percent of parcels were assigned to a different elementary grade 'home school' than they were assigned to the year before under the busing system. At the middle school and high school levels this number was 52 and 35 percent respectively. Therefore, the 2002-2003 home school for many students is often <u>not</u> the school they would have been assigned at the time they chose their residence. This dramatic change in school assignment zones, the simultaneous introduction of a sweeping school choice plan, and the assignment of students to high-demand schools by lottery provides a unique opportunity to estimate parental preferences for schools and model the heterogeneous impact of attending a chosen school on academic outcomes.

Lottery Assignments

Approximately one third of the schools in the district were oversubscribed. The district implemented a lottery system for determining enrollments in those oversubscribed schools. Under the lottery system, students choosing non-home schools were first assigned to priority groups and student admission was then determined by a lottery number. The priority groups for district schools were arranged in lexicographic order based on the following priorities:

Priority 1: Student who had attended the school in the prior year. (Students were subdivided into 3 priority groups depending upon their grade level, with students in terminal grades—grades 5, 8 and 12—given highest priority.)

Priority 2: Free-lunch eligible student applying to school where less than half the students were free-lunch eligible.

Priority 3: Student applying to a school within their choice zone.

Students listing a given school as their first choice were sorted by priority group and a randomly assigned lottery number.[4] Any slots remaining after home school students were accommodated were assigned in order of priority group and random number.[5] If a school was not filled by those who had listed it as a first choice, the lottery would repeat the process with those listing the school as a second choice, using the same priority groups as above. However, for many oversubscribed schools, the available spaces were filled up by the time the second choice priority groups came up.

Students who were not assigned one of their top choices were placed on a waiting list. About 19% of students winning the lottery to attend their first choice schools subsequently attended a different school, with 13% choosing to attend their home school instead and another 6% choosing to attend a different school entirely, with most of these students changing address. When slots became available, students were taken off the wait list based on their lottery number alone, without regard for their priority group.

Potential for Strategic Choice

The lottery mechanism used by the Charlotte-Mecklenburg schools was a 'first-choice-maximizer,' in which slots were first assigned to all those listing a given school as a first choice before moving to those listing the school as a second or third choice. In such a mechanism, parents with poor home school options may have an incentive to misstate their preferences – not listing their most preferred school if it had a low

[4] The random number was assigned by a computer using an algorithm that we verified with CMS computer programmers.

[5] Once any sibling was admitted to a school, other siblings could choose to attend the school. We dropped those who were admitted to a school because of a sibling preference.

probability of admission (Glazerman and Meyer (1994), Abdulkadiroglu and Sonmez (2003), Abdulkadiroglu et. al (2006)). Instead, they may have hedged their bets by listing a less preferred option with a higher probability of admission in order to avoid being assigned to their home school. Such strategic behavior would imply that student choices would not reflect true preference orderings for schools—to the extent that students are *not* listing their preferred match due to strategic hedging on quality.

However, there were a number of reasons why such strategic behavior was probably rare in the first year of the choice plan that we are studying. First, parents did not know the details of how the lottery system would be operated. The handful of district officials who knew the lottery details were not allowed to communicate them to parents. Parents were never given their actual lottery numbers. The district also told parents that they would make every attempt to give each student admission to one of their chosen schools, and *instructed them to list what they wanted*. In order to accommodate demand, the district substantially expanded capacity at popular schools. In addition, the district gave a 'priority boost' to low-income students choosing to attend schools with low concentrations of low income students. Hence, choices for top schools by students with under-performing home schools would be given top priority. This would counteract the incentive for these students to hedge their choices as outlined above.

If there were widespread strategic behavior by parents, we would expect those with low-quality default schools to hedge their bets and list less desirable schools for which they might have a higher probability of admission. In another paper, Hastings, Kane and Staiger (2006a), we test whether parents with *exogenous* changes to the quality of their default school produced by the redistricting had lower preferences for high-quality schools as would be predicted if parents were behaving strategically. Perhaps because of the uncertainty about the mechanism and the newness of the system, we did not find evidence that strategic behavior played a significant role in this first year of school choice.[6]

[6] For the details of this test, as well as for further specification checks on the mixed-logit demand estimation, please see Hastings, Kane and Staiger (2006a), Section VIII, pp. 21-27.

III. Preferences, Choices and Expected Treatment Effects

We have access to administrative data for all students in CMS for the year before and after the implementation of school choice. Throughout the analysis, we focus on students entering grades 4 through 8 since we have baseline test scores for this group of students for North Carolina End of Grade Tests. For each student, we have the choice forms submitted to CMS, allowing a student to specify up to 3 choices for their school. In addition to the student choices our data contain student characteristics for the years before and after school choice, including geo-coded residential location, race, gender, lunch-subsidy recipient status, and student test scores for standardized North Carolina end-of-grade exams in math and reading, and school assignment. Coupled with these data are data on lottery number, lottery outcomes and student assignments for the 2002-2003 school choice lottery.

It is clear that parents have very heterogeneous preferences over school characteristics. Figure 1 shows that approximately 20% of students chose schools that had lower test scores than the school they had guaranteed admission to. In addition, among those with the same elementary home school for 2002-03, parents on average listed 10 different elementary schools as their first choice.[7] The range of choices made suggests that heterogeneous preferences may play a key role in school selection, and may therefore generate differential gains in academic achievement.

Expected Treatment Effect Given Choice

Suppose that parents choose schools for both the expected academic gain for their child, but also for other reasons, such as proximity or racial composition. Consider the following utility function of parent *i* for school *j*

(1) $$U_{ij} = \beta_i A_{ij} + V_{ij}$$

where A_{ij} is the expected academic achievement of student *i* if she attends school *j*, V_{ij} is the utility for student *i* from attending school j along non-academic dimensions, and β_i is

[7] This statistic excludes heterogeneity in choices generated solely by heterogeneity in prior-year school assignment under the bussing system. If we include choices driven by preferences for prior-year schools by students with different prior-year schools under bussing, but the same new home-school assignment area under choice, this statistic increases to 14.6.

the weight that parent *i* places on academic achievement relative to non-academic dimensions. The utility gain from attending the first choice over the alternative school is:

(2) $$\Delta U = \beta_i \, \Delta A + \Delta V$$

where delta denotes the difference in variables between school alternatives *k* and *j*. A student will choose an alternative school over their home school only if the utility gain is positive, *i.e.* $\Delta U > 0$. Among students choosing an alternative school over their home school, the expected academic gain of a student randomized into their 1st choice school is given by:[8]

(3) $$E(\Delta A \mid \beta_i \Delta A + \Delta V > 0)$$

In this simple framework, students with high β_i have a positive expected treatment effect (gain in academic achievement from attending the first-choice school). In fact, as β_i gets very large, the expected treatment effect alone determines choice and, therefore, must be positive for all students who choose an alternative school. For a student with low β_i (near zero), the expected treatment effect is ambiguous. If ΔA and ΔV are independent and ΔA is mean zero, then the expected treatment effect is zero, i.e. $E(\Delta A \mid \Delta V > 0) = 0$. If ΔA is negatively correlated with ΔV, as may be the case for some non-academic dimensions such as proximity and percent African American, then the treatment effect will be on average negative for students placing a near zero weight on academic outcomes. That is, test scores of students choosing for a school characteristic that is negatively correlated with academics will tend to fall if they are admitted to their first choice school. Hence, this basic framework generates the prediction that the expected treatment effect is positive for all students with a strong preference for academic achievement. Among students with weaker preferences for academic achievement, the expected treatment effect will depend on the tradeoffs that parents face. The treatment effect could even be negative if expected academic achievement is sufficiently negatively correlated with other valued school characteristics.

[8] As noted earlier, the lottery was run as a 'first-choice maximizer'. Because of this, most students who did not win the lottery for their first choice school were assigned to their home school.

Estimating Preferences Using a Random Utility Model

We use the data on student choices and demographics to infer preferences for academic achievement using a random utility framework. Let U_{ij} be the expected utility of individual *i* from attending school *j*. Suppose that utility is a linear function of the academic achievement of student *i* at school j, A_{ij} , and other school-student characteristics, Z_{ij} , such as distance from home, busing availability, and racial composition. Thus, expected utility is given by:

$$U_{ij} = \beta_i A_{ij} + \gamma_i^* Z_{ij} + \omega_{ij} \tag{4}$$

where β_i and γ_i^* represent preference parameters for person i, and ω_{ij} represent an unobserved idiosyncratic preference of student *i* for school *j*.

Furthermore, suppose that the expected academic achievement for student i attending school j depends on the average test score at school *j* (S_j with a coefficient normalized to one), other observable characteristics of the school (Z_{ij}), plus a mean zero deviation that is known to the parent (υ_{ij}).

$$A_{ij} = S_j + \alpha Z_{ij} + \upsilon_{ij} \tag{5}$$

Thus, parents base their expectations of academic achievement on observable school characteristics plus idiosyncratic factors affecting their child. This specification allows for the possibility that non-academic factors such as proximity may affect academic achievement (for example, through longer bus rides) and also allows for the possibility that parents adjust school test scores for racial composition of the school (the "value-added" approach). This adjustment can be different for parents with different observables (such as race) if preferences for school characteristics are allowed to vary with student observable characteristics.

Using equation (5), we can re-write the indirect utility function as:

$$U_{ij} = \beta_i S_j + \gamma_i Z_{ij} + \varepsilon_{ij} \tag{6}$$

where $\varepsilon_{ij} = \beta_i \upsilon_{ij} + \omega_{ij}$ and $\gamma_i = \gamma_i^* + \beta_i \alpha$. Assuming that ε_{ij} follows an independent extreme value distribution, we get the typical logit functional form for the probability of choosing school *j*.[9] With distributional assumptions on the preference parameters, we have a mixed-logit utility model. The mixed logit can approximate any random utility model, given appropriate mixing distributions and explanatory variables (Dagsvik (1994), McFadden and Train (2000)).

Even though the expected academic achievement for student *i* attending school *j* (A_{ij}) is not observed directly, the weight placed on academic achievement can be estimated. The weight on academic achievement (β_i) is identified in the mixed logit model because school test scores are assumed to influence utility only through their effect on academic achievement, whereas other school characteristics in Z_{ij} may affect utility directly as well as indirectly through expected academic achievement. Obtaining direct estimates of A_{ij} would require additional structural assumptions to identify how academic achievement depends on observable school characteristics and idiosyncratic factors affecting a child's academic performance at each school ($\alpha Z_{ij} + \upsilon_{ij}$).[10] Rather than impose additional assumptions, our analysis focuses on the more fundamental implication that parents placing a high weight on academics (high β_i) should choose schools that increase their child's academic achievement. An important benefit of this approach is that we do not have to completely specify how school and student characteristics combine to produce academic achievement, we just need to know that academic achievement was important to the parent in choosing a school.

Given a preference distribution, we estimate the underlying preference parameters in this random utility model using simulated maximum likelihood techniques (Train 2003). The probability that individual *i* chooses schools (j^1,j^2, j^3) is given by:

[9] Note that estimation involves normalizing the variance of ε_{ij}. Since Var(ε_{ij})is an increasing function of β_i, normalization will reduce the estimate of β_i for high- β_i types. While this will act to understate the estimated variation in β_i in the final model, it does not affect the relative rankings of individuals with respect to β_i – which is the information we use to estimate heterogeneous treatment effects.

[10] For example, if $A_{ij} = X_j\beta_i + v_{ij}$ and $V_{ij} = Z_{ij}\gamma_i + \omega_{ij}$, with v_{ij} i.i.d. normal and ω_{ij} i.i.d. extreme value, and no common variables in X and Z, then one can estimate $E(\Delta A|\Delta U>0)$ directly from the random utility model. We estimated models of this form and found that they performed poorly in terms of predicting the magnitude of the treatment effect, suggesting that either our assumptions were too restrictive or the necessary student-choice level idiosyncratic parameters were poorly identified.

$$P_i\left(j^1, j^2, j^3\right) = \Pr\left\{\left(U_{ij^1} > U_{ik} \forall k \in J_i^1\right) \cap \left(U_{ij^2} > U_{ik} \forall k \in J_i^2\right) \cap \left(U_{ij^3} > U_{ik} \forall k \in J_i^3\right)\right\}$$

$$(7) \qquad = \int \prod_{c=1}^{3} \frac{e^{X_{ij^c}\beta}}{\sum_{k \in J_i^c} e^{X_{ik}\beta}} f(\beta \mid \mu, \theta) d\beta$$

We assume that $\beta \sim f\left(\beta \mid \mu_\beta, \theta\right)$, where $f(\cdot)$ is a joint-normal mixing distribution, μ denotes the mean, and θ represents the variance parameters. The term inside the integrand represents the probability of observing the three ranked choices conditional on the preference coefficients (β): this is the product of three logit probabilities evaluated at β_i, corresponding to the probability of making each choice from among the remaining options.[11] This conditional probability is integrated over the distribution of β to yield the unconditional probability of observing the ranked choices. Estimation was by the method of maximum simulated likelihood, using 100 draws of β from $f(\cdot)$ for each individual in the data set. The results were not sensitive to the number of draws used. We assume that all random parameters are drawn from a normal or log normal distribution, and allow for correlation among some of the main preference parameters as reported in the tables.

The maximum likelihood results yield parameter estimates for the mean and variance of preferences in the population. We then use Baye's rule to calculate posterior estimates of the weight each student placed on school scores in the following way (Revelt and Train (1998) and Train (2003)):

$$(8) \qquad E\left(\beta_i^A \mid y_i, X_{ij}, \theta\right) = \frac{\int P\left(y_i \mid X_{ij}, \beta\right) f(\beta \mid \theta) d\beta}{P\left(y_i \mid X_{ij}, \theta\right)}$$

Where y_i denotes the choices the student made, X_{ij} denotes the student and school characteristics that enter the indirect utility function, θ denotes the parameters that describe the density of β, and β_i^A represents the weight student *i* placed on school test scores (including the estimated effect of income and student baseline scores). This equation is the expected value of student *i*'s preference for academics given her characteristics, the choices she made, the characteristics of the schools given her location, and the estimated distribution of preferences in the population. We calculate this

[11] For students submitting fewer than three choices, the likelihood is modified in an obvious way to reflect only the probability of the submitted choices.

posterior for each student in our randomized lottery admission group using 1000 draws from the estimated preference distributions in from the mixed logit demand estimation.[12]

The posterior estimate effectively calculates how different a student's preferences must have been from the average to generate the observed sequence of choices given her characteristics and the choice set she faced. Thus estimating β_i^A allows us to succinctly incorporate all of the relevant choice information for each student into one statistic – the estimated value the student places on a school's academic performance.

IV. Demand Estimation Results

We follow Hastings, Kane and Staiger (2006a) and present mixed logit results from that model here. The model includes key observable school characteristics. To measure proximity, we included the travel distance (measured in miles along roads) from each student to each school, an indicator if the student was eligible for busing to the school, and an indicator if the school was designated as the student's neighborhood school. We included a binary measure indicating whether the child attended at given school in the prior year. To capture the racial composition of a school, we included the percent black in the school in Spring 2003 and its square. When the quadratic term has a negative coefficient, this specification yields an implied bliss point (where the quadratic peaks) for preferences over racial mix of a school. To capture the academic quality of the school, we included a measure of average test scores in the school (the school level average of all students' standardized math and reading scores in spring of 2003).[13] Table I lists the independent variables in the indirect utility function and describes how they were constructed.

We allow the mean preference for academic achievement (the coefficient on school test scores) to vary linearly with the student's standardized baseline test score (from the spring of the prior year, standardized by grade level across the district) and the

[12] See Train (2003) p. 270 for Monte Carlo Simulations of the accuracy of individual-level parameter estimates and the number of observed choice situations.

[13] We use the average test scores at the end of the first year of choice instead of those at the end of the year before school choice was implemented. Hastings, Kane and Staiger (2006a) present demand estimation results for various measures of school academic achievement. They find for example that value added measures do a poor job of explaining choices, and that the choice data imply that parents are choosing schools based on levels instead of changes in academic achievement.

median household income in the student's neighborhood for the student's race (measured in \$1000's, using their census block group in 2000, and de-meaned with the countywide median of \$51,000). Preferences for distance are constrained to be negative, following a lognormal distribution. We allow preferences for proximity and academic quality to be correlated. All other preference distributions are assumed to be independently and normally distributed. We estimate the parameters of the preference distribution separately by race and lunch-subsidy status. This permits the preference distributions and logit-normalization to vary across the four socio-economic categories.

Several aspects of the CMS school choice data and experiment are helpful for identifying preferences in our demand estimation. We will mention them briefly here and refer the reader to Hastings, Kane and Staiger (2006a) for further detail. First, the large scale redistricting that occurred with the introduction of school choice helps identify preference parameters separately from residential sorting. Without redistricting and the multiple-choice responses, residential sorting would potentially confound the preferences for proximity with preferences for other desired school attributes. [14]

Second, historic placement of schools for busing in CMS provides wide variation in school characteristics for families in all socio-economic groups, dampening collinearity problems that may be present in other settings.[15] Third, approximately 95% of the 110,000 students submitted choices for the choice plan. Thus we have data for nearly the entire student population—whereas most work using school choice data has been dependent on limited and potentially non-representative subgroups of students.

Fourth, the multiple responses create variation in the choice set by effectively removing the prior chosen school from the subsequent choice set. This choice-set variation allows us to estimate the distribution of preferences for school characteristics from observed substitution patterns for each individual – a stronger source of variation for identification than cross-sectional changes in the choice set based on geographic location (Train (2003), Berry, Levinsohn, and Pakes (2004)). Intuitively, when only a

[14] In addition, multiple choices listed by those selecting their home school first further separates preferences for school characteristics from residential sorting by simulating the unavailability of the neighborhood school. For a comparison of preference estimates for the redistricted sub-sample versus the full sample, please see Hastings, Kane and Staiger (2006a). They show that preferences are very similar for the redistricted subsample of students relative to the population.

[15] Hastings, Kane and Staiger (2006a) show that average distance to a top-tier school is the same across all socio-economic groups.

single (1st) choice is observed for every individual, it is difficult to be sure whether an unexpected choice was the result of an unusual error term (ε_{ij}) or unusual preferences by the individual (β_i) for some aspect of the choice. However, when an individual makes multiple choices that share a common attribute (e.g. high test scores) we can infer that the individual has a strong preference for that attribute, because independence of the additive error terms across choices would make observing such an event very unlikely in the absence of a strong preference.

The final estimation sample includes 36,816 students entering grades 4-8. Estimation is limited to these grades because of the lack of test scores (either baseline or school test scores) in other grades. The means and standard deviations of these variables across the 2.4 million school, student, and choice rank interactions available to our sample of students and schools are reported in Table II.

Table III presents the results from the mixed logit demand estimation by race and lunch-recipient status. All of the point estimates were precisely estimated and statistically different from zero at less than the 1 percent level. We report the estimates for the means, standard deviations, and correlation coefficients (where appropriate) for the preference distributions. The discussion of results is focused around the parameters most relevant for our final estimation of the effect of attending a first choice school on academic achievement. For a further discussion of the results and their implications for student sorting and competition on quality in public school choice, please see Hastings, Kane and Staiger (2006a).

The first four rows of coefficients in Table III report the preferences for school test scores by race and lunch-recipient status. The first row of coefficients reports the mean preference for school scores for the average student. It is positive for all four demographic groups, implying that school test scores have a positive effect on choosing a school for the average student. For a student with average baseline test scores and average income, the mean preference for school scores is larger for non-white students (1.80) than for white students (1.17) among students not receiving lunch subsidies. These coefficients imply that a 0.1 increase in average test scores at a school (one tenth of a student-level standard deviation) is associated with a 10%-20% increase in the odds of choosing that school for an average student not receiving lunch subsidies. Preference for school scores among students receiving lunch subsidies are lower for both whites and nonwhites, but the difference between whites and nonwhites is similar.

To allow for the possibility that preferences for school scores would vary with student baseline academic ability as well as student income level, we included interactions between school test scores, student's baseline test scores, and neighborhood income level.[16] The third and fourth rows of parameter estimates report the coefficient on the interaction of school scores with income and the student's baseline score respectively. Recall that both neighborhood income and the student's baseline score are "de-meaned", so that the coefficient on the main effect of school score measures the value of school test score for a student with average income and baseline test score (both equal to zero).

The coefficients on the income interaction imply that mean preferences for a school's test score (conditional on its racial composition) are increasing with income. The magnitudes of these parameters are roughly consistent with the differences in the mean preferences for test scores between lunch-recipients and non-lunch recipients within race. Similarly, the mean preference for school scores is increasing in the student's baseline test score. The coefficient on the interaction between the student's baseline test score and the school mean test score is positive - implying that those with higher test scores relative to their baseline peer group value a school's test scores more. The effect of a student's baseline score on the preference for school test scores is similar in magnitude to the effect of income. A one standard deviation increase in the baseline test score is associated with a 0.3-0.6 increase in the mean preference for school test scores, while a one standard deviation increase in neighborhood income (about $25,000) is associated with a 0.3-0.4 increase in the mean preference for school test scores.

The coefficients on the interactions of income and baseline score with school scores demonstrate that there is considerable observable heterogeneity in preferences for academics. Parameter estimates for the standard deviation in idiosyncratic preferences for academics are reported in Row 2. While differences in baseline test scores and income each generate a standard deviation in preferences of roughly 0.3-0.6 based on the calculations from the previous paragraph, the estimated standard deviation in

[16] For students who are eligible for lunch subsidies, we did not include the interaction with neighborhood income because all of these students are presumably very low income. In initial specifications using a conditional logit, income interactions with the preference for school scores were generally insignificant for the lunch-recipient segments.

idiosyncratic preferences for school test scores is also around 0.3 for non-whites and 0.65 for whites. Hence, there is substantial unobserved heterogeneity in preferences for test scores. Taken together, the coefficients imply that academics are on average a significant determinant of school choice. Furthermore, the substantial variation across students in the weight placed on academics suggests that we may expect to see strong school choice selection on academic outcomes for some students and not for others. The fact that much of the heterogeneity in preferences is unobservable implies that the traditional approach of allowing the treatment effect to vary with observable characteristics, such as race or lunch status, may not completely capture heterogeneous treatment effects by preferences for academics.

The parameter estimates for the remaining coefficients indicate that parents face important trade-offs between academic and non-academic factors when choosing schools. Rows 5 and 6 report the parameter estimates for the lognormal distribution of preferences for distance. Rows 7 and 8 report the mean preference and standard deviation for the neighborhood (or 'home') school.[17] Parents dislike distance and prefer their neighborhood school. These coefficients indicate that the average parent must trade-off utility for proximity in order to gain utility from expected academic outcomes. For most students, attending a high-achieving school will require them to choose a school that is farther than their home school and a school that is not their home school. Hence there is a negative correlation between school characteristics that measure proximity and those that capture academic achievement. This implies that parents of all races must, on average, trade-off utility for academic gains against utility gains for proximity.

In addition to trading-off proximity for academics, African American parents must trade-off academic gains against the racial composition of peers. The preference coefficients on percent black imply that the average African American parent prefers schools where approximately 70% of the student body population is black, while the district as a whole is approximately 45% African American. However, the percent black at a school is negatively correlated with average test scores (correlation is around -0.65).

[17] Hastings, Kane and Staiger (2006a) discuss the interpretation of the neighborhood school. They test if this coefficient represents a non-linearity in the preference for proximity or if it is potentially consistent with a default effect. They provide evidence that the preference for the neighborhood school is a neighborhood preference that is not generated by default behavior.

The negative correlation between test scores and racial composition implies that African American parents must value academic achievement much more than their white counterparts in order to induce them to choose a higher performing school that also has, on average, fewer African American students. Given the coefficients for the quadratic term in racial preferences, the loss in utility for black families is highest when percent black is low (less than 40%), which is precisely the range in which school average test scores are highest.

V. Estimating the Impact of Attending a 1st Choice School on Academic Achievement

In this section we estimate the causal relationship of attending ones 1st choice school on academic achievement, allowing heterogeneity in the treatment effect with the revealed preference for school test scores. In order to exploit the randomized admissions, we focus on the subset of students choosing schools that were over-subscribed and limit our sample to the marginal priority groups within those schools for whom lottery number alone determined initial admission. We ignore members of priority groups in which all students were either admitted or denied admission—since the assignment of lottery numbers had no impact on their admission status. In some schools, the marginal priority group will consist of those who attended the school the year before, or free-lunch eligible students, or students from the choice zone. The marginal priority group may also be different for different grade levels in a school.

Within the marginal priority groups, we estimate the impact of *attending* a first-choice school on academic achievement. Since not all of those who won the lotteries actually chose to attend their first choice school, and some of those who lost the lotteries were subsequently admitted off the waiting lists, we used the randomized lottery outcome as an instrumental variable to estimate the impact of attending one's first choice school in following regression:

$$Y_{ij} = X_i \alpha + \gamma_1 Attended1stChoice_{ij} + \gamma_2 Attended1stChoice_{ij} * \hat{\beta}_i^A + \delta_j + \varepsilon_{ij} \tag{9}$$

Winning the lottery and winning the lottery interacted with $\hat{\beta}_i^A$ serve as instruments. Note that all of the information used to derive the preference weights was observed prior to randomization. Since $\hat{\beta}_i^A$ depends only on baseline data that is independent of whether the student won the lottery, its interaction with winning the lottery is a valid instrument once one has conditioned on baseline data. We include as regressors: $\hat{\beta}_i^A$, gender, race/ethnicity, free lunch status, home school dummy variables, baseline test scores, income, absences, suspensions, and grade retentions. Random assignment by lottery implies that the impact of winning the lottery, γ_1, is consistently estimated even without these control variables, but the additional control variables greatly improve precision. Finally, note that coefficient estimates for terms involving $\hat{\beta}_i^A$ are not attenuated by the usual measurement error bias – the measurement error ($\beta_i^A - \hat{\beta}_i^A$) is uncorrelated with the posterior estimate $\hat{\beta}_i^A$ by construction (Hyslop and Imbens, 2001).

The fixed effects, δ_j, are included for each school and grade, to account for the fact that the probabilities of winning the lottery varied across lotteries. We report robust standard errors, allowing for correlations in outcomes among students with the same first-choice school (which may include more than one grade with a lottery). As long as winning the lottery has an impact on student outcomes only through the likelihood that one attends a first choice school, then the IV estimates of γ_1 and γ_2 using the lottery outcome and its interaction with $\hat{\beta}_i^A$ as instruments, will be consistent estimates of the impact of attending one's first choice on various outcomes.[18]

In equation (9) the effect of attending one's first choice school is $\gamma_1 + \gamma_2\hat{\beta}_i^A$. If the dependent variable is the student's own test score, we expect $\gamma_2 > 0$, implying that students who place more weight on test scores experience a larger treatment effect. The parameter γ_1 gives the treatment effect for a student that places no weight on test scores in their

[18] After the initial lotteries some students were taken off the waitlist according to lottery number. Adding the waitlisted students to our sample (in addition to the marginal priority groups), we estimated specifications similar to the one above, using as instruments both whether or not a student won the lottery and the randomly assigned lottery number interacted with being placed on the waitlist. The results were quite similar to the results we report.

school choice decision, and could in principal be negative as such a student would trade off other school attributes for lower test scores.

Lottery Data and Characteristics of the Randomized Subpopulation

We began with the choice forms submitted by 105,706 students in the first year. Reflecting the district's intensive outreach efforts, choice forms were received for over 95% of all the students enrolling that fall. After dropping students who were not in grades 4-8, who had special disabilities needs, and students who were admitted because of siblings, we were left with a sample 37,115. Of these, 22,872 listed their guaranteed home school (n=19,669) or magnet continuation school (n=3,203) and, therefore, were not subject to randomization. Another 7,583 students were in groups sufficiently high on the priority list that they were not subject to the randomization. There were 3,065 students in marginal priority groups, described above as those priority groups within the schools where slots were allocated on the basis of a random number. Finally, there were 3,595 students in priority groups that were sufficiently low on the priority list that all members of the priority group were denied admission and placed on the waitlist.

Our outcome measures include data on student absences, suspensions, and standardized test scores for all students in the district for the years surrounding the implementation of the choice program. Because students in kindergarten through 2nd grade do not take the state exams, and because high school students only take the end-of-course exams in the subjects they choose, we had reading and math scores for students in grades 3 through 8. We standardized the test scores by grade level and year to have mean zero and a standard deviation of one. In addition the testing data in North Carolina also include student self-reports on the number of hours of home work they did each week.

Empirical Results: Summary Statistics

We focus on the 3,065 sample members in the marginal priority groups that were subject to the randomization. Table IV compares descriptive statistics on the baseline characteristics for these students and for students in the district as a whole. Students in the marginal priority group were more likely to be African American, and more likely to receive lunch-subsidies, reflecting their higher probabilities of choosing non-guaranteed schools and the priority boosts given to these students. Students in the randomized group also tended to have lower test scores, higher absences, and more suspensions, and have home schools with lower average test scores and higher percent minority. It is important to keep these differences in mind, since we are only able to estimate the impact of the school choice program for the population of students in the randomized group.

In order to verify the validity of the randomization of lottery numbers, we examine the baseline characteristics of lottery winners and losers within the randomized group. Table V reports these baseline characteristics for our estimation sample. Our estimation sample excludes 181 students who were in marginal priority groups but missing needed baseline characteristics (such as address, which was used in the choice model). The table reports unadjusted differences, as well as differences from an OLS regression including fixed effects for the school program and grade for which the lottery is being conducted. Before adjusting for lottery block fixed effects, there are some differences in baseline characteristics between lottery winners and losers (although none are statistically significant). However, these differences were largely due to a correlation between the characteristics of lottery participants and the lottery odds. After including a fixed effect for each school program and grade, all such differences were smaller and were not significantly different from zero.

Impact of Winning Lottery on the Characteristics of School Attended

Before presenting the effects of winning the lottery on student outcomes, we test whether winning the lottery had any effect on the characteristics of the school attended. Table VI reports the results of winning the lottery on the characteristics of the school attended, based on OLS estimates of the following equation:

$$(10) \qquad Y_{ij} = X_i\beta + \gamma WonLottery_{ij} + \delta_j + \varepsilon_{ij}$$

In these regressions, we control for student baseline characteristics, and home school and choice-grade fixed effects. The regression results give the average impact of winning the lottery on the characteristics of the school attended, Y_{ij}.

The first row of estimates in Table VI shows that lottery winners were 53 percentage points more likely to attend their first choice school than the lottery losers. This is the first stage regression for the instrumental variables regression of the impact on test scores of attending a first choice school. This estimate is not equal to 100 percent for two reasons: first, some of those who were given the opportunity to attend their first choice did not do so and, second, some of those who were originally waitlisted at their first choice were subsequently called off the waitlist. Overall, approximately 75% of lottery winners and 25% of lottery losers attended their first choice school.

The second row of estimates in Table VI show the effect of winning the lottery on whether or not the student was enrolled in any CMS school in the 2002-2003 school year. This estimate gives the differential attrition rate between lottery winners and losers. Average attrition rates were fairly low at 9.8%, and consistent with estimates of inter-county mobility rates from the Census.[19] The estimated effect of winning the lottery on attrition is small in size (-0.018) and not significantly different than zero, indicating that there was no significant differential attrition by the end of the 2002-2003 school year.

The remaining rows of Table VI report the impact of winning the lottery on average student characteristics in the school attended. Winning the lottery was associated with approximately one-tenth of a student-level standard deviation increase in the average combined reading and math scores at the attended school. In addition, winning the lottery implied that students attended a school with a significantly lower concentration of free-lunch recipients.

The Effect of Attending a First-Choice School on Student Outcomes

In order to estimate the marginal impact of allowing parents to switch to their first choice school, we used an indicator of whether or not a student won their lottery as an instrument for attending one's first choice school to estimate:

[19] Approximately 6% of school age children living in the south moved to a different county between March 2002 and March 2003. Mobility rates tend to be somewhat higher in urban, high-poverty populations (Schachter, 2004).

$$(11) \qquad Y_{ij} = X_i \alpha + \gamma_1 Attended1stChoice_{ij} + \gamma_2 Attended1stChoice_{ij} * \hat{\beta}_i^A + \delta_j + \varepsilon_{ij}$$

Estimates of the average treatment effect (equation 9 with γ_2=0) for various student-level outcome measures, Y_i, are reported in Table VII. [20]

The estimates in Table VII are broken down by academic and non-academic outcomes. For non-academic outcomes we include the impact of winning the lottery on absences, suspensions, retentions, and homework time. Among these outcome measures, the average treatment effect is significant and negative for retention rates. Winning the lottery to attend a first choice school causes a dramatic reduction in retentions – a 2.3 percentage point decrease off of an average base of 2.2%. We do not find a significant impact on absences or suspensions, however. In addition, we find that students who are randomized into their first choice school report spending more time on homework. The outcome measure is an indicator if the student reports spending more than 3 hours per week on homework on self-reported surveys given to students with the end of grade exams. Even though students who attend their first choice school report a significant increase in homework hours, we find no measurable average effect on standardized test scores. The final row of estimates in Table VII shows no significant impact of attending a first choice school on standardized test scores.[21] The point estimate is nearly zero, but there is a relatively large standard error. The results are consistent with the current literature, and while they exploit randomization into first choice schools to create credible counterfactual comparisons, the average treatment effect may mask important heterogeneity.

Heterogeneous Treatment Effects on Test Score Outcomes

Prior studies, which have not had data for estimating preferences, have used observable characteristics to identify subgroups of students who on *a priori* grounds are believed to have different underlying reasons for choosing schools that may be correlated with the expected treatment effect. Table VIII shows estimates of the average treatment

---

[20] Some readers may prefer to see the reduced form impact of winning the lottery on various student outcomes. Recall from Table V that lottery winners were roughly 50 percentage points more likely to attend their first choice school than lottery losers. To obtain a rough estimate of the reduced form impact of winning the lottery, simply divide the estimates of γ_2 and its standard error in Table VI by 2.

[21] Regression estimates show the same effect on math and reading scores when run separately, so we use the combined score to improve precision.

effect on student test scores in various subgroups of students defined on the basis of student demographics or characteristics of the school chosen. Estimates for most of the subgroups remain insignificant. However, the estimated treatment effect is positive and significant for two of the subgroups (whites, and students with above median income) and there is an apparent pattern of positive treatment effects for higher SES students and students applying to more academically oriented schools. Taken together, the subgroup estimates suggest that there might be heterogeneous treatment effects, although there is no clear story for the pattern of effects and the power to detect these effects is fairly low (especially given the large number of potential subgroups one could consider).

The pattern of subgroup impacts is strongly related to differences across subgroups in the underlying determinants of choice. As can be seen in Figure 2, the subgroup estimates from Table VIII are strongly positively associated with the average weight that students in the subgroup place on school test scores (correlation=0.89). Table IX summarizes the variation in our posterior estimates of the preferences for academic achievement both across and within subgroups. There is considerable variation in the sample as a whole, with a standard deviation in $\hat{\beta}_i^A$ of 0.81, ranging from a value near zero at the 5th percentile to a value near 3 at the 95th percentile. The posterior estimates vary across subgroups in ways we expect: students with higher income and higher baseline achievement, as well as students choosing more academically oriented schools have higher average $\hat{\beta}_i^A$. But within each subgroup, there is considerable variation in preferences as well.

The evidence from Tables VIII and IX highlight three advantages of using estimates of $\hat{\beta}_i^A$ to identify heterogeneous treatment effects, rather than subgroup estimation based on observables such as race and income. First, using a single index, rather than estimating differences in impacts for an arbitrary number of subgroups, increases the precision with which we can identify heterogeneous treatment effects by exploiting all of the within and between subgroup variation in preferences. Second, the $\hat{\beta}_i^A$ incorporate information on the choice set, distinguishing between students who pick a good school because it is convenient versus students who pick it for its academics. The $\hat{\beta}_i^A$ incorporate this information as guided by utility theory, rather than controlling for characteristics of the choice set in add hoc way. Third, the $\hat{\beta}_i^A$ can be used to evaluate the impact of school choice outside of the estimation sample. If we believe that the $\hat{\beta}_i^A$ captures the relevant dimensions of parental preferences, we can infer the impact of choice even for those who were not in the marginal priority groups.

Table X incorporates the weights placed on academic achievement when choosing a school into the estimated impact of attending a first choice school on standardized test scores. The coefficients imply that the effect of attending one's first choice school on a student's test scores is significantly increasing with the weight that a student placed on test scores in choosing a school. The regression estimates imply that a one standard deviation increase in the weight that an individual places on school test scores raises the treatment effect on the student's own test score by 0.062 standard deviations. For students who place no weight on test scores in their school choice, the coefficient on attending one's first-choice school implies a negative (although not significant) treatment effect – their test scores fall by 0.105 standard deviations if they attend their first-choice school. These estimates imply a small negative impact (-0.002 standard deviation score gain) of attending a first-choice school on test scores for an average student with a $\hat{\beta}_i^A$ of 1.34, and a large positive effect on test scores (about 0.10-0.20) for students in the top decile of the $\hat{\beta}_i^A$ distribution.

A 0.1 standard deviation increase in a student's test score results is equivalent to a 3-4 percentile rank gain in test scores. Child development psychologists suggest that a 5 percentile rank gain in a student's test score translates into a significant cognitive gain in academic aptitude. Alternatively, estimates of the impact that test scores have on future earnings suggest that a 0.1 standard deviation in increase in test scores is worth $10,000 to $20,000 in net present value of future earnings (Kane and Staiger, 2002).

Parents Who Face Significant Tradeoffs

These estimates are consistent with our general prediction that students with high $\hat{\beta}_i^A$ should have a positive expected treatment effect (gain in academic achievement from attending the first-choice school). However, the treatment effect for a student with low $\hat{\beta}_i^A$ (near zero) is theoretically ambiguous and depends on whether parents face trade-offs – if expected academic achievement is negatively correlated with other valued school characteristics. Since the percent black at a school is negatively correlated with average test scores in CMS schools (correlation is around -0.65), the racial composition of a school is an important trade-off that many African American parents face. We estimate (from the mixed-logit results) that the average African American parent prefers schools where approximately 70% of the student population is black. Parents that prefer a school with a high proportion of students African American must value academic achievement more in order to induce them to choose a higher performing school that also has, on average, fewer African American students. Thus, all students with strong academic preferences (high $\hat{\beta}_i^A$) will have a positive gain in academic achievement from attending the first choice school, but among students with weak academic preferences (low $\hat{\beta}_i^A$) we might expect a negative treatment effect among students that prefer a school with a high proportion African American. In other words, the interaction effect between $\hat{\beta}_i^A$ and winning the school choice lottery should have a negative intercept and a steeper slope for students who have strong preferences for predominantly African American schools.

Table XI presents the results from specifications identical to those in Table X, but estimated separately for students who prefer a school that is less than 55 percent black (primarily white students) and students who prefer a school that is more than 55 percent black (primarily non-white students). Posterior estimates of student-level preferences for school racial composition were calculated in the same way as the $\hat{\beta}_i^A$'s were. The average treatment effect is positive for students who prefer a predominantly white school, and there is no significant interaction with the weight that the student places on test scores in their school choice. For these students, both high and low $\hat{\beta}_i^A$ students experience academic gains from attending their first choice school. In contrast, for students who prefer a predominantly black school there is a significant interaction between their estimated preference for academics and the treatment effect. High $\hat{\beta}_i^A$ students experience academic gains from attending their first choice school that are similar to students who prefer a predominantly white school. In contrast, low $\hat{\beta}_i^A$ students who prefer a predominantly black school experience a negative effect on academic performance from attending their first choice school. This evidence suggests that the relationship between preferences and treatment effects may depend importantly on the trade-offs that parents face given their preferences and their school choice options. These results also highlight the potential importance of the underlying decision-making process to understanding the heterogeneous impacts that public school choice has on student academic outcomes.

Schools versus Student-school Interactions

Our results show that the school one attends has a causal impact on student academic outcomes, but not all parents choose a school that maximizes their child's academic gains. Students who choose schools primarily for academic gains are able to get those gains from attending their chosen schools, while others do not. An interesting secondary question is whether all students (even those who put little weight on academics) would gain academically from attending these same schools. Do students who place a high weight on academics simply choose better schools in the sense that any student attending such a school would receive a positive treatment effect? If so, efforts to identify these schools and steer students towards them may increase the impact of school choice on academic outcomes. Alternatively, academic gains may be student-school specific, with the impact of attending a school depending on the student's ability to gain academically at that school. In this case, students who choose for academic achievement at a particular school will have gains in academics that students who chose the same school for convenience will not.

In Table XII, we investigate the extent to which the heterogeneity in treatment effects associated with individual preferences for academics can be explained by characteristics of the school chosen (suggesting common school quality). The first column of the table replicates the base estimates for comparison. The next two columns add an interaction to allow the treatment effect to depend on the difference in average test scores between the 1st-choice school and the home school. This interaction is insignificant on its own, and has no impact on the interaction with $\hat{\beta}_i^A$, suggesting that the gap in test scores in not a good indicator of the likely treatment effect. Columns 4 and 5 add an interaction with the average test score in the 1st-choice school (not relative to the home school), and the results are similar. Since average test scores may be a poor proxy for academic quality of a school, the remaining columns of the table try interacting the treatment with the average weight ($\hat{\beta}_i^A$) among all students participating in the school's lottery (a high average indicates the school attracts students who care about academics) and a full set of school effects (allowing each school to have its own average impact on student outcomes). In both of these specifications, we find no significant evidence of heterogeneity in treatment effects across schools, and the coefficient on the interaction with $\hat{\beta}_i^A$ is little changed (although the standard errors are larger).

Overall, the heterogeneity in treatment effects associated with individual preferences for academics do not appear to be explained by characteristics of the school chosen, although the estimates are not very precise. The general pattern of results suggests that idiosyncratic match-specific school quality plays a more important role in determining academic outcomes.

VI. Conclusion

When given the choice to attend a public school other than the home school to which they have been assigned, the parents of 49 percent of the students in Charlotte took the opportunity and listed a school other than their assigned school as a first choice. In this paper, we evaluate the impact of switching schools on various academic and non-academic outcomes. On average, among those applying to the oversubscribed schools, winning the lottery had no discernable impact on students' own reading and math scores overall, even though lottery winners attended schools with higher math and reading scores than did lottery losers. Winning the lottery had only modest impacts on other outcomes, such as increasing homework time and reducing grade retentions.

However, parents seem to choose schools for many different reasons. Indeed, one quarter of parents who were willing to switch chose schools with lower mean test scores than their assigned schools. Overall, the results presented in this paper imply that the effect of attending one's first choice school on academic outcomes is significantly increasing with the value that a student placed on test scores in choosing a school.

Among students placing a high weight on school test scores, there was a positive effect of attending their first-choice school on their own test score. In contrast, for students who placed a low weight on school test scores, we found negative effects of attending their first-choice school on their own test score.

A number of recent papers have found no impact on average of attending a first-choice school on academic achievement. Our evidence suggests that the absence of any academic gains on average does not imply that school choice is ineffective. Instead, parents appear to get what they want. When parents want improved academic outcomes, they are able to get them. When parents value other school attributes, and are willing to trade off academic gains for utility gains on other dimensions, school choice will allow them to make that choice—even if maximizing parental utility does not maximize academic achievement. Ultimately, the trade-offs that parents are willing to make along these dimensions in choosing a school for their children will determine whether school choice programs will be successful at improving academic achievement.

References

Abdulkadiroglu, Atila and Tayfun Sonmez "School Choice: A Mechanism Design Approach", *American Economic Review,* 93(3):729-747, June 2003.

Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez "Changing the Boston School Choice Mechanism: Strategy-proofness as Equal Access," working paper, Harvard University*,* May 2006.

Altonji J.G., T.E. Elder, C. Taber. "An Evaluation of Instrumental Variable Strategies for Estimating the Effect of Catholic Schools." *National Bureau of Economic Research Working Paper* No. 9358, 2002.

Berry, Steven, James Levinsohn and Ariel Pakes "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112(1): 68-105, 2004.

Blank, R.K. Survey of Magnet Schools: Analyzing a Model of Quality Integrated Education. Washington, DC, ABT Associates for the Department of Education,1983.

Bryk, A., V. Lee and P. Holland. Catholic Schools and the Common Good. Cambridge, MA: Harvard University Press, 1993.

Cullen, Julie, Brian Jacob and Steven Levitt "The Effect of School Choice on Student Outcomes: Evidence from Randomized Lotteries" forthcoming, *Econometrica*, 2006.

Coleman, James S., Thomas Hoffer and Sally Kilgore, High School Achievement: Public, Catholic, and Private Schools Compared. New York: Basic Books, 1982.

Dagsvik, John K, "Discrete and Continuous Choice, Max-Stable Processes, and Independence from Irrelevant Attributes," *Econometrica*, 62(5):1179-1205, 1994.

Evans, William and Robert Schwab. "Finishing High School and Starting College: Do Catholic Schools Make a Difference?" *Quarterly Journal of Economics* 110 (4): 941-974, 1995.

Gamoran, A. "Student Achievement in Public Magnet, Public Comprehensive, and Private City High Schools." *Educational Evaluation and Policy Analysis* 18 (1): 1-18,1996.

Glazerman, Steven and Robert H. Meyer, "Public School Choice in Minneapolis," in T.A. Downes and W.A. Testa, eds., *Midwest Approaches to School Reform.* Proceedings of a conference held at the Federal Reserve Bank of Chicago, 1994, pp.110-126.

Green, J. P., P.E. Peterson, and J. Du (1997). “The Effectiveness of School Choice: The Milwaukee Experiment.” Cambridge, MA, Program on Education Policy and Governance, Harvard University.

Hastings, Justine S., Thomas Kane, and Douglas Staiger (2006a) “Parental Preferences and School Competition: Evidence from a Public School Choice Program,” *National Bureau of Economic Research Working Paper* No.11805, November 2005. Revised version (May 2006) available at: http://aida.econ.yale.edu/~jh529/HKS_QualityCompetition_NBERWP.pdf

Hastings, Justine S., Thomas Kane and Douglas Staiger (2006b) “Gender, Performance and Preferences: Do Girls and Boys Respond Differently to School Environment? Evidence from School Assignment by Randomized Lottery,” *American Economic Review Papers and Proceedings,* 96 (2): 232-236.

Heckman, James, “Instrumental Variables: A Study of Implicit Behavioral Assumptions Used In Making Program Evaluations,” *Journal of Human Resources*, 32(3):441-462, 1997.

Heckman, James J., Jeffrey Smith, and Nancy Clements, “Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts,” *Review of Economic Studies*, 64:487-535, 1997.

Heckman, James J., Sergio Urzua and Edward Vytlacil, “Understanding Instrumental Variables in Models with Essential Heterogeneity,” *Review of Economics and Statistics*, 88 (3):389-432, 2006.

Hyslop, Dean R. and Guido W. Imbens. "Bias From Classical And Other Forms Of Measurement Error," *Journal of Business and Economic Statistics*, 2001, v19(4), 475-481.

Krueger, Alan B. and Pei Zhu "Another Look at the New York City School Voucher Experiment," *American Behavioral Scientist*, January 2004, pp.658-98.

Mayer, Daniel P., Paul E. Peterson, David E. Myers, Christina Clark Tuttle, and William G. Howell, “School Choice in New York City After Three Years: An Evaluation of the School Choice Scholarships Program,” Washington, DC: Mathematica Policy Research, Inc., 2002.

McFadden, D. and K. Train (2000). “Mixed Multinomial Logit Models for Discrete Response.” *Journal of Applied Econometrics*, 15(5): 447-470.

Neal, D. “The Effects of Catholic Secondary Schooling on Educational Achievement.” *Journal of Labor Economics,* XV, 98-123, 1997.

Revelt, David and Kenneth E. Train, “Mixed Logit with Repeated Choices: Households’ Choices of Appliance Efficiency Level,” *Review of Economics and Statistics*, Vol. 80 (4): 647-657, Nov. 1998.

Peterson, P. E., D. Myers and W. G. Howell (1998). “An Evaluation of the New York City School Choice Scholarship Program: The First Year.” Cambridge, MA, Mathematica Policy Research and the Harvard Program on Education Policy and Governance.

Rouse, Cecilia. E. 1998. “Private school vouchers and student achievement: An evaluation of the Milwaukee Parental Choice Program,” *Quarterly Journal of Economics*, 113(2):553-602.

Train, Kenneth E. Discrete Choice Methods with Simulation. Cambridge University Press (2003).

Witte, John F., Troy D. Sterr and Christopher A. Thorn (1995). Fifth-Year Report: Milwaukee Parental Choice Program. Madison, Wisconsin, University of Wisconsin.

Figure 1: Distribution of Difference in Average Standardized School Score Between Student's First Choice School and Home School.

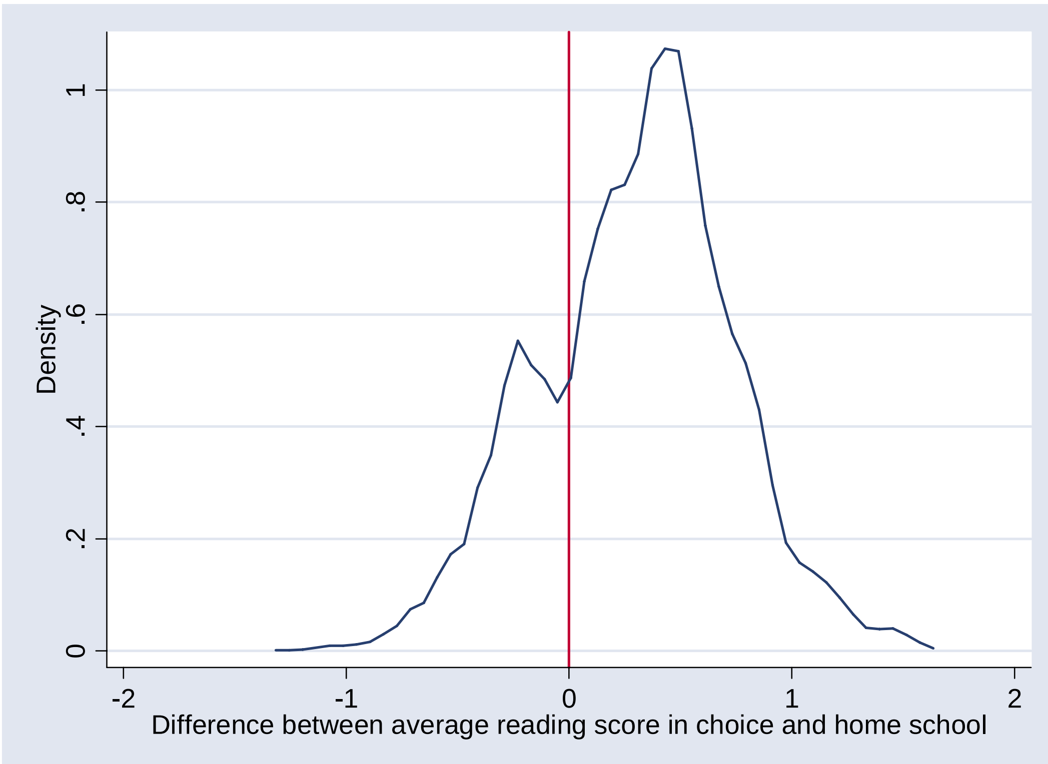

Figure 2. Subgroup Estimates of the Effect of Attending a 1st-choice School

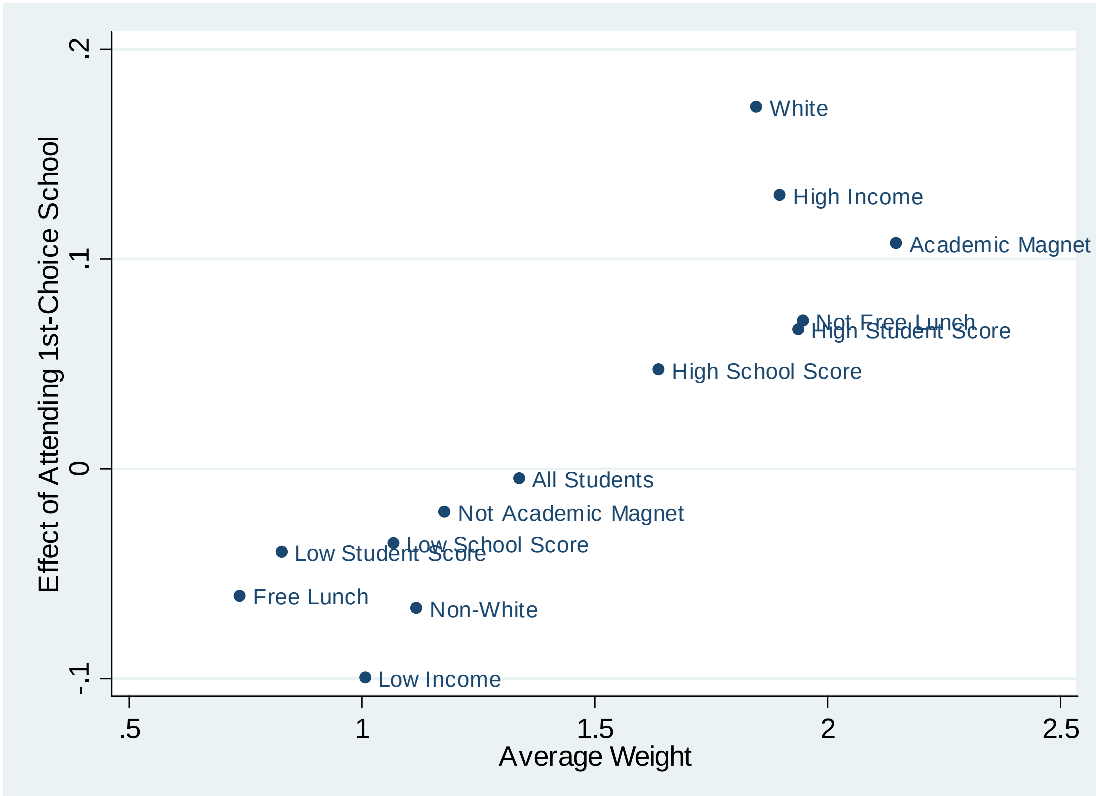

Table I: Key Explanatory Variable Definitions

Variable	Description
Distance	Driving distance from student *i* to school *j* calculated using MapInfo with Census Tiger Line files.
School Score	Average of the student-level standardized scale score for students in school *j* on math and reading End of Grade exams for the 2002-2003 school year. This is the average of the test score variable described below across all students in school *j*.
Test Score	The sum of student *i*'s scale score on End of Grade math and reading exams in baseline year 2001-2002 standardized by the mean and standard deviation of district-wide scores for students in his or her grade.
Income	The median household income reported in the 2000 Census for households of student i's race in student i's block group. Income is demeaned by the county-wide average of approximately $51,000 and is reported in thousands of dollars.
Percent Black	The percent of students in school *j* who are black according to 2002-2003 school year administrative data.

Table II: Key Explanatory Variable Summary Statistics

Summary Statistics Using First Choice Data					
Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
Distance	2434113	13.0071	6.7254	0.0010	42.4069
School Score	2434113	-0.1087	0.4487	-0.9537	1.9478
Test score	2434113	0.0567	0.9886	-2.9113	3.0255
Income	2434113	5.1226	27.5669	-48.5010	149.0010
Percent Black	2434113	0.5252	0.2507	0.0584	0.9801

Table III: Estimates from Mixed Logit Model

		Parameter Estimates*			
		Not Receiving Lunch Subsidies		Receiving Lunch Subsidies	
Variable	**Parameter**	**White**	**Black**	**White**	**Black**
Preferences for Scores					
School Score	**Mean**	1.1732	1.8035	0.3671	0.9396
	Std. Dev.	0.5674	0.2688	0.6175	0.2706
Income*School Score	**Mean**	0.0151	0.0126	--	--
	Std. Dev.	--	--	--	--
Baseline own score * School Score	**Mean**	0.5558	0.5734	0.2924	0.4995
	Std. Dev.	--	--	--	--
Preferences for Proximity					
Distance**	**Mean**	-0.3526	-0.2684	-0.3784	-0.2751
	Std. Dev.	0.0684	0.0413	0.1273	0.0639
Home School	**Mean**	2.1300	1.7373	1.9816	1.7710
	Std. Dev.	0.5130	0.6799	0.8248	0.7752
Preferences for Race					
Percent Black	**Mean**	3.3068	5.1340	1.9268	3.1409
	Std. Dev.	2.6417	1.6447	2.0795	0.8745
Percent Black Sqaured	**Mean**	-5.4580	-3.6790	-3.5385	-2.3005
	Std. Dev.	--	--	--	--
Implied Mean Preferred % Black		0.3029	0.6977	0.2723	0.6827
	Std. Dev.	0.2420	0.2235	0.2938	0.1901
Other Preferences					
Last-year School	**Mean**	3.7941	3.3837	3.5016	2.8495
	Std. Dev.	2.4977	2.7896	3.4651	3.3825
Choice Zone (busing)	**Mean**	1.1909	1.2484	1.9203	1.6132
	Std. Dev.	0.8285	1.2418	1.5083	1.2442
Estimated Correlation Coefficients:					
Corr(Distance, School Score)		0.4939	-0.1055	0.3379	-0.6355
Corr(Distance, Home School)		-0.0788	0.0007	-0.2623	-0.1122
Corr(School Score, Home School)		-0.7888	-0.6016	-0.8411	-0.5895

* All estimates are significant at the 1% level or higher

** Distribution of preference on distance follows a log normal distribution.

Table IV: Comparison of Student Characteristics

	All Students	Chose Guaranteed School	Chose Non-guaranteed School Admitted	Randomized	Waitlisted
Student demographics					
Black	44.3%	34.6%	62.5%	59.7%	54.8%
Free or reduced lunch	39.2%	31.3%	60.3%	51.3%	34.3%
Student's prior year performance					
Reading test score (SD units)	0.02	0.15	-0.26	-0.09	-0.11
Math test score (SD units)	0.02	0.16	-0.26	-0.12	-0.15
Absent 18 or more days	8.5%	6.8%	11.7%	10.8%	10.7%
Retained	1.5%	1.2%	2.0%	1.9%	1.9%
Suspended	12.2%	9.3%	17.7%	16.5%	15.4%
Choice school characteristics					
Average combined scores	0.05	0.09	-0.09	0.08	0.10
Percent free or reduced lunch	40.6%	38.6%	50.9%	36.6%	35.6%
Percent black or hispanic	49.4%	46.2%	59.8%	50.0%	47.0%
Home school characteristics					
Average combined scores	-0.08	0.03	-0.28	-0.23	-0.27
Percent free or reduced lunch	47.0%	40.7%	59.3%	53.3%	56.0%
Percent black or hispanic	53.6%	47.1%	65.3%	61.6%	63.8%
School assignment					
Assigned to 1st choice	85.4%	100.0%	100.0%	40.4%	0.0%
Assigned to guaranteed school	72.5%	100.0%	0.0%	44.6%	74.5%
School attendance 02-03					
Attended 1st choice	78.7%	92.1%	81.6%	45.4%	16.2%
Attended home school	58.8%	79.4%	9.7%	35.0%	51.3%
Number of students	37115	22872	7583	3065	3595

Notes: Data from Charlotte-Mecklenberg Schools (CMS). Sample includes all students in grades 4-8 who applied to a regular or magnet school as their 1st choice for the 2002-2003 school year and were enrolled in CMS in the 2001-2002 school year. Students guaranteed placement because of siblings and in ESL are excluded.

Table V: Baseline Characteristics by Treatment and Control Group

Variable	Admitted	Waitlisted	Difference	Adjusted Difference
Student demographics				
Black	0.614	0.585	0.030	0.011
			(0.067)	(0.022)
Free or reduced lunch	0.467	0.531	-0.064	-0.015
			(0.078)	(0.012)
Median income ($1000s) by race and block-group in 2000 census	48.4	49.4	-1.0	-0.7
			(3.6)	(0.7)
Student's prior year performance				
Reading test score	-0.127	-0.069	-0.058	-0.025
			(0.110)	(0.031)
Math test score	-0.135	-0.113	0.023	0.025
			(0.106)	(0.030)
Absent 18 or more days	0.097	0.106	-0.009	-0.007
			(0.013)	(0.016)
Suspended	0.152	0.162	-0.010	-0.022
			(0.028)	(0.015)
Retained	0.019	0.018	0.001	0.001
			(0.005)	(0.006)
Home school characteristics				
Average combined score	-0.241	-0.213	-0.028	0.003
			(0.051)	(0.013)
Percent free or reduced lunch	0.543	0.524	0.019	0.001
			(0.034)	(0.007)
Percent black	0.625	0.607	0.018	-0.003
			(0.036)	(0.007)
Number of students	1175	1709	2884	2884

Notes: Sample limited to students in randomized priority groups with complete baseline data. Difference is between students admitted (won the lottery) and waitlisted (did not win the lottery). Each adjusted difference is from a separate regression of the given baseline characteristic on whether the student was randomly assigned to her first-choice school, controlling for lottery fixed effects. Standard errors adjust for clustering at the level of the first-choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table VI: The Impact of Being Randomly Assigned to 1st Choice School on Characteristics of School Attending at End of 2002-2003 School Year

Characteristic of School Attending	Mean	Estimated Impact
First choice school	0.460	0.533***
		(0.054)
Not attending CMS in 2002-2003 (Attrition)	0.098	-0.018
		(0.011)
School average combined score	-0.073	0.129**
		(0.040)
Percent free or reduced lunch	0.463	-0.070***
		(0.019)
Percent black or Hispanic	0.576	-0.049
		(0.026)
Total observations		2884

Note: Each entry in the table is from a separate regression of the given characteristic of the school a student was attending at the end of the year on whether the student was randomly assigned to her first choice school, controlling for lottery fixed effects, home school fixed effects, and the baseline covariates listed in Table V. Sample includes only students in the randomized priority group with complete baseline data. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table VII. Instrumental Variables Estimates of the Impact of Attending 1st Choice School on Student Outcomes in 2002-2003

Student Outcome	Mean	Average Treatment Effect
Non-academic Measures		
Absent 18 or more days	0.135	-0.001
		(0.023)
Suspended	0.201	0.012
		(0.032)
Retained	0.022	-0.023*
		(0.009)
> 3 hrs. homework per week	0.303	0.122*
		(0.050)
Academic Performance		
Combined test score	-0.086	-0.005
		(0.050)

Note: Each entry in the table is from a separate IV regression of the given student outcome on whether the student was attending her first choice school, using random assignment to the first choice school as an instrument. These regressions control for lottery fixed effects, home school fixed effects, and the baseline covariates listed in table V. Sample includes 2884 students in the randomized priority group with complete baseline data. Sample sizes for homework (N=2554) and combined test score (N=2581) are smaller due to missing data on the dependent variable for some students. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table VIII: Subgroup Estimates of the Effect of Attending a 1^{st}-choice School

Sample	IV Estimate of Effect of Attending 1^{st}-Choice School on Combined Test Score	Number of Students
All Students	-0.005	2581
	(0.050)	
Race:		
Non-White	-0.067	1790
	(0.058)	
White	0.172*	791
	(0.073)	
Income:		
Below Median	-0.100	1601
	(0.058)	
Above Median	0.130*	980
	(0.063)	
Free Lunch Eligibility		
Eligible	-0.061	1296
	(0.078)	
Not Eligible	0.070	1285
	(0.043)	
Baseline Test Score		
Below Average	-0.040	1386
	(0.055)	
Above Average	0.066	1195
	(0.064)	
1^{st}-Choice School Academic Magnet		
Not Academic Magnet	-0.021	2155
	(0.055)	
Academic Magnet	0.107	426
	(0.089)	
1^{st}-Choice School Combined Score		
Below Median	-0.036	1337
	(0.080)	
Above Median	0.047	1244
	(0.043)	

Note: Each row of the table reports estimates for a different student sub-sample, as indicated. The table reports IV estimates of the impact of attending the first choice school on the combined student test score, using random assignment to the first choice school as an instrument. Regressions control for lottery fixed effects, home school fixed effects, and the baseline covariates listed in table V. Sample includes only students in the randomized priority group with complete baseline data. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table IX. Summary Statistics for Posterior Estimate of Weight Students Place on Test Scores in School Choice Overall and in Student Subgroups.

Sample	Mean	Standard Deviation	5th Percentile	95th Percentile
All Students	1.34	0.81	0.19	2.74
Race:				
Non-White	1.12	0.71	0.15	2.41
White	1.85	0.77	0.46	3.06
Income:				
Below Median	1.01	0.61	0.13	2.14
Above Median	1.90	0.78	0.46	3.05
Free Lunch Eligibility				
Eligible	0.74	0.44	0.04	1.47
Not Eligible	1.95	0.62	0.92	2.95
Baseline Test Score				
Below Average	0.83	0.52	0.06	1.77
Above Average	1.94	0.65	0.94	2.99
1st-Choice School Academic Magnet				
Not Academic Magnet	1.18	0.74	0.16	2.52
Academic Magnet	2.15	0.63	1.04	3.24
1st-Choice School Combined Score				
Below Median	1.07	0.69	0.12	2.33
Above Median	1.64	0.82	0.34	2.93

Note: Each row of the table reports estimates for a different student sub-sample, as indicated. The first column reports the average weight that the students place on test scores (*Weight*) in the school choice decision. The second column reports the standard deviation of *Weight*. The final two columns report the 5th and 95th percentile of *Weight.*

Table X: IV Estimates of the Impact of Attending 1st Choice School with Heterogeneous Treatment by Weight Placed on Academics in Choice Decision

	Combined Score	**Combined Score**
Attended 1st-choice school	-0.005	-0.105
	(0.050)	(0.074)
Weight * attended 1st-choice school		0.077*
		(0.031)
P-value for interaction with *Weight*		0.016
Joint p-value on reported coefficients	0.924	0.031
Observations	2581	2581

Notes: Each column in the table is from a separate IV regression. The dependent variable is a student's combined standardized test score in the spring of 2003. Each specification reports the coefficients on attending the first choice school and its interaction with the weight that the student places on test scores (*Weight*) in the school choice decision, using random assignment to the first-choice school and its interaction with *Weight* as instruments. All specifications control for lottery fixed effects, home school fixed effects, the baseline covariates listed in Table V, and a direct control for the student's *Weight* estimate. Sample includes only students in the randomized priority group with complete baseline data. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table XI: IV Estimates of the Impact of Attending 1st Choice School with Heterogeneous Treatment by Weight Placed on Academics in Choice Decision, Estimated Separately by Student Preference for Racial Mix at School

Dependent Variable: Combined Score	**Students Who Prefer School Less Than 55% Black**		**Students Who Prefer School at Least 55% Black**	
Attended 1st-choice school	0.115 (0.058)	0.186 (0.158)	-0.054 (0.059)	-0.164* (0.078)
Weight * attended 1st-choice school		-0.041 (0.065)		0.098* (0.041)
P-value for interaction with *Weight*		0.533		0.019
Joint p-value on reported coefficients	0.052	0.097	0.250	0.053
Observations	870	870	1711	1711

Notes: Each column in the table is from a separate IV regression. The dependent variable is a student's combined standardized test score in the spring of 2003. Each specification reports the coefficients on attending the first choice school and it interaction with the weight that the student places on test scores (*Weight*) in the school choice decision, using random assignment to the first-choice school and its interaction with *Weight* as instruments. All specifications control for lottery fixed effects, home school fixed effects, the baseline covariates listed in Table V, and a direct control for the student's *Weight* estimate. Sample includes only students in the randomized priority group with complete baseline data. Student preference for racial composition in the school is each student’s posterior estimate of the value that maximizes their quadratic utility in %black at the school. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).

Table XII. IV Estimates of the Impact of Attending 1st Choice School with Heterogeneous Treatment Effect Associated With Characteristics of the School Chosen.

	Base Model	Interaction With Difference in Average School Sores		Interaction with Average Score in 1st-Choice School		Interaction With Average Weight of Students in 1st-Choice Lottery		Interaction With 1st-Choice School Fixed-Effects	
Attended 1st-choice school	-0.105 (0.074)	0.000 (0.057)	-0.100 (0.077)	-0.015 (0.049)	-0.089 (0.073)	-0.164 (0.132)	-0.166 (0.132)	n.a.	n.a.
Attended 1st-choice school *									
Weight	0.077* (0.031)		0.077* (0.032)		0.061 (0.041)		0.043 (0.053)		0.074 (0.071)
Difference in average school scores		-0.017 (0.072)	-0.019 (0.077)						
Average score in 1st-choice school				0.136 (0.090)	0.082 (0.114)				
Average *Weight* of students in lottery						0.121 (0.075)	0.081 (0.112)		
P-value for interaction with *Weight*	0.016		0.018		0.149		0.427		0.293
P-value for other interaction(s)		0.813	0.803	0.136	0.474	0.109	0.472	0.066	0.203
Joint p-value on reported coefficients	0.031	0.963	0.071	0.241	0.02	0.13	0.039	0.069	0.177
Observations	2581	2581	2581	2581	2581	2581	2581	2475	2475

Notes: Each column in the table is from a separate IV regression. The dependent variable is a student's combined standardized test score in the spring of 2003. All specifications include the same controls as in Table X. Standard errors adjust for clustering at the level of the first choice school. Asterisks indicate significance (*=.05, **=.01, ***=.001).