

Blueprint Labs

Discussion Paper #2026.11

Transportation as a Barrier to Education Access: Evidence from Chicago Public Schools

Cecilia Moreira
Steven Puller
Ini Umosen

October 2026

The views expressed in this paper are those of the authors and do not necessarily reflect the views of MIT Blueprint Labs, the Massachusetts Institute of Technology, or any affiliated organizations. Blueprint Labs working papers are circulated to stimulate discussion and invite feedback. They have not been peer-reviewed or subject to formal review processes that accompany official Blueprint Labs publications.



MIT Department of Economics
77 Massachusetts Avenue, Bldg. E53-390
Cambridge, MA 02139

Transportation as a Barrier to Education Access: Evidence from Chicago Public Schools

Cecilia Moreira — Stanford University

Steven L. Puller — Texas A&M University and NBER

Ini Umosen* — Ithaka S+R

October 2026

Abstract

School choice weakens the link between neighborhood and educational opportunity only if students can reach schools available to them. Using student-level data for Chicago high schoolers, we find that Black students are less likely to rank top-rated schools first and face longer travel times to them. We use a discrete choice model to separate school value and travel-time costs. Counterfactually, we give Black students white students' travel times and estimate that demand for top-rated schools increases 23 percentage points. Finally, express bus service that targets underserved neighborhoods increases demand for top-rated schools. Our counterfactuals quantify how transportation expands school access.

*Moreira: Stanford University, cecilia.moreira@stanford.edu. Puller: Texas A&M University and NBER, spuller@tamu.edu. Umosen: Ithaka S+R, ini.umosen@ithaka.org.

We are grateful to MIT Blueprint Labs and Chicago Public Schools for their support and data access. This research was supported by funding from the National Bureau of Economic Research and U.S. Department of Transportation research initiative "Transportation Economics in the 21st Century". We thank Josh Angrist, Andrew Barr, Ben Klopack, Maya Mikdash, Parag Pathak, Brittany Street, Chris Walters and seminar participants at Asia-Pacific IO Conference, Berkeley, Carnegie-Mellon, Chicago Fed, International IO Conference, LSU, MIT Blueprint Labs, NBER Economics of Transportation, the Federal Reserve Bank of Philadelphia, Urban Economic Association Meetings, University of Illinois-Chicago, University of Maryland, University of Missouri, University of Texas and Wharton. We thank Francisco Barba for excellent research assistance. Parts of this paper were prepared with the assistance of generative artificial intelligence tools. The authors reviewed, edited, and take full responsibility for the final version of the manuscript.

1 Introduction

Many public education systems have adopted school choice policies with a goal of weakening the link between where a student lives and the quality of their educational opportunities. These policies expand a student's choice set beyond an assigned neighborhood school. However, the burden of reaching an alternative school is left to the family. In large urban districts where the school system does not provide transportation, the public transit system becomes an important determinant of how much access school choice actually provides. Therefore, the extent to which school choice policies 'level the playing field' will depend on whether transportation policy makes it feasible for students to reach schools beyond their neighborhoods.

This paper studies the complementarities between education and transportation policy in the context of high school students in Chicago Public Schools. In this setting, there is a citywide high school choice system where students apply via a centralized ranked-choice system, and most students rely on public transit rather than district-provided yellow bus service. Students are able to rank schools throughout the city, however transit time within the public transit network will shape which schools are practically reachable. These transit times will depend not only upon distance but also the routes, transfers and service frequency of the transit system. Therefore, two students who live similar distances from the same school may face very different travel times. Because these transit constraints can vary significantly across neighborhoods, the transit network can have localized impacts on students, suggesting that policies to reduce travel times need to be highly targeted.

We ask whether transportation frictions limit the equalizing potential of school choice. Using ranked-choice application data for the first cohort of eighth grade students who applied to high school through a new platform, we link each student's submitted rankings, demographics, test scores, and public transit travel times from the student's home area to each high school. These data allow us to study whether students who are *formally* allowed to choose schools throughout the city face unequal *practical* access to the schools they value.

We begin by documenting a striking empirical regularity. Black students are much less likely to rank a top-rated school as their first choice: 38 percent do so as compared to 62 percent for both

white and Hispanic students. Black students also face a longer public transit time to top-rated schools. For Black students, the median transit time to the nearest Level 1+ school is 25 minutes, as compared to 15 minutes for Hispanic students and 20 minutes for white students. These patterns raise a key policy question: do differences in demand for top-rated schools reflect differences in what families value, or differences in the time cost of reaching the schools they value?

To address this question, we estimate a discrete choice model of school demand that separates preferences for individual schools from the disutility of travel time to reach those schools. The model uses student-level ranked-choice data and public transit travel times from each student's residential neighborhood to each school. Chicago's urban setting coupled with the layout of the public transit network generates substantial variation in both the value of school options and the time required to reach them. Students living in different parts of the transit network face different travel times to the same schools, and each student faces a set of school options that vary in both their value to the student and the travel time to the school. This variation allows our model to distinguish between schools that families most value and schools that are chosen because they are easier to reach.

Our estimates show that transportation access is central to how school choice operates. The 126 high schools vary substantially in how much families value them, and there is a high degree of agreement across racial groups in the ranking of schools. However, the time cost of reaching those schools is large and varies significantly across neighborhoods; thus the estimated disutility from travel is large enough to have a substantial effect on which schools families choose. For example, we estimate that an additional 30 minutes of travel time generates disutility equivalent to the difference in school value between the 20th- and 75th-percentile school. This quantifies the importance of transportation time in access to schools.

We then use our model parameters to conduct a simple thought experiment. If Black students faced the same travel times as white students, how would their school choices change? To simulate this counterfactual, each Black student retains their valuation of each school, but we assign that student the travel times of a randomly selected white student and then estimate choices implied by the discrete choice model. The effect on school preferences is substantial. The share of Black students

predicted to prefer a top-rated school increases by 23 percentage points, which substantially closes the observed racial gap in first choice rankings. This thought experiment implies that transportation is not a secondary friction; rather it is a key mechanism through which residential neighborhood affects educational opportunity.

A natural identification concern is that the patterns of transportation time and school access reflect residential sorting rather than the role of transportation time alone. Under this concern, families with stronger unobserved preferences for school quality may move to neighborhoods with better transit access to top-rated schools, and this could cause our model to attribute differences in rankings to travel time rather than unobserved school preferences. We address this concern in two ways. First, using student address histories, we examine whether families move closer to top-rated schools, or to neighborhoods with faster access to highly-valued schools, in the years before starting high school. We find no evidence of systematic movement toward neighborhoods with better access. Second, we estimate utility model specifications that allow for unobserved preferences for school quality to be correlated with the disutility of travel time. The estimated travel time parameter changes very little. Therefore, these results suggest that residential sorting is unlikely to explain our main findings.

We then use our model parameters as a policy design tool to identify where transportation improvements would most increase demand for highly-valued schools. We estimate the impacts of counterfactual express bus routes that target students in predominantly Black neighborhoods on the South Side with poor transit access to top-rated schools. Holding fixed the valuations of schools, we ask how demand changes when travel times are reduced to driving times, which approximate direct express bus service. We estimate that first-choice demand for multiple top-rated schools would increase by 9–15 percentage points. The counterfactuals therefore show how ranked-choice and travel-time data can identify the neighborhood-school connections where limited transit access most constrains demand.

This policy design tool can be used to target a wide variety of schools selected by education policymakers. Our baseline counterfactuals target top-rated schools designated as "Level 1+" by the school district. These ratings are also strongly correlated with observed family demand. However,

the same transportation intervention can be directed toward schools selected according to different policy objectives. We also present counterfactual transportation interventions that target schools with high estimated value-added, and we find that express bus interventions would increase demand for schools in the top quintile of mathematics value-added. This shows that transportation can be used as a policy lever to shift demand toward schools with high contributions to achievement, even if those contributions are only weakly related to observed family demand.

These counterfactuals intentionally measure the impact of transportation on student *demand* rather than on final assignment to schools. Transportation enters upstream by affecting the set of schools students are willing to consider and rank in their applications. Assignment mechanisms then determine to which schools a student is offered admission based upon the rankings, available capacity, and admissions priorities. Both margins are first-order. A well-designed assignment mechanism can allocate seats transparently and limit strategic incentives, but it operates on the demand presented to the mechanism. It is the underlying topology of the urban transit system that acts as a structural gatekeeper, shaping the first margin of access – the schools that families are willing to rank. Our goal is to isolate the role of transportation on the first margin. Combining ranked-choice data with student-to-school travel times allows us to quantify this upstream transportation margin.

More generally, our model identifies where in a city transportation most sharply constrains the demand for highly-valued schools. This cannot be inferred from school locations or distance to school alone because the frictions depend both on how costly it is to reach a school and how much a student values the school once they reach it. And public transit access depends on specific routes, transfers and network connections rather than simply on geographic distance. We can identify particular transportation corridors where shorter travel time would most increase demand by combining estimated school mean utilities with neighborhood-to-school travel times. We find that these frictions are highly local.

This paper highlights a limit of school choice as an equalizing policy. School choice can weaken the link between residence and educational opportunity, but it does not eliminate the role of place when families must navigate an uneven transportation network. Families may be allowed to rank schools throughout the city and may value top-rated schools, but some of those schools may be too

costly to reach. If transportation prevents students from choosing highly-valued schools, then other policies – assignment reforms, information provision, or capacity expansions – cannot by themselves equalize access.

Our results inform multiple dimensions of education policy beyond decisions about where to add supplemental transportation. School districts choose where educational programs and capacity are located. Therefore, our framework can be used to inform the siting of new magnet or specialty programs, targeted capacity expansions, and decisions about school closures or consolidations and the impacts of those decisions on different populations in the city.

This paper contributes to the literature on school choice, household demand for schools, and the extent to which choice systems equalize access to top-rated schools. A large literature studies how families trade off school quality and other attributes when making school choices and how assignment systems affect outcomes, including Hastings, Kane, and Staiger [2008], Abdulkadiroğlu, Pathak, and Walters [2018], Abdulkadiroglu, Pathak, Schellenberg, and Walters [2020], Allende [2021], Beuermann, Jackson, Navarro-Sola, and Pardo [2023], Barseghyan, Clark, and Coate [2019], Campos and Kearns [2020], and Calsamiglia, Fu, and Güell [2020]. This literature shows that families value school quality but that demand also reflects peers, program attributes, distance, and other factors. We build on this approach but isolate a constraint that is particularly important in large urban districts and largely absent from this literature: travel time through the public transit network.

A closely related literature studies information frictions, search frictions, beliefs, and awareness sets in school choice. Hastings and Weinstein [2008], Hoxby and Turner [2013], Hoxby and Turner [2015], Allende, Gallego, and Neilson [2026], Kapor, Neilson, and Zimmerman [2020], Arteaga, Kapor, Neilson, and Zimmerman [2022], Ainsworth, Dehejia, Pop-Eleches, and Urquiola [2023], Agte, Allende, Kapor, Neilson, and Ochoa [2026], and Campos and Kearns [2020] show that families may have incomplete information, limited awareness, search frictions, or biased beliefs about schools. Corradini [2024] and Corradini and Idoux [2025] show that racial differences in school choices reflect differences in information and beliefs as well as preferences over school attributes. Our paper focuses on a different margin. While this literature shows that information and beliefs affect which schools

families consider, we show that even when families value top-rated schools, they may not prefer them if those schools are too costly to reach.

Our paper also contributes to work on geography and school choice. Laverde [2026] is the closest comparison. Using Boston Public Schools, she shows that geography remains a major barrier to equalizing access under school choice. Other work has shown that residential location, proximity, and assignment rules shape school choices and assignment outcomes, including Oosterbeek, Sóvágó, and van der Klaauw [2021], Calsamiglia, Fu, and Güell [2020], Barrow, Sartain, and de la Torre [2020], and Park and Hahm [2022]. We differ by measuring access using travel time through the public transit network rather than geographic distance and by using the estimated demand model to identify where transit access most constrains demand.

Our paper also contributes to a growing literature on transportation and school choice. Costa, Naddeo, and Seiler [2026] use New York City administrative data matched to public transit travel times and show that longer commutes substantially reduce applications to several of New York City's highly selective exam high schools. Trajkovski, Zabel, and Schwartz [2021] study kindergarten choice in New York City and show that school bus eligibility affects school choices. We build on this work by studying transportation across the much broader set of schools available through a citywide high school choice market and examining how transportation access contributes to differences in demand across demographic groups and neighborhoods. Our demand model separates school valuations from travel-time costs and identifies the particular neighborhood-school connections where transportation most constrains demand.

Finally, our paper speaks to the literature on neighborhood effects and place-based inequality. A large body of work shows that childhood neighborhoods affect long-run outcomes, often at fine geographic levels; see for example Chetty, Friedman, Hendren, Jones, and Porter [2026], Chyn and Katz [2021], Deutscher [2020], Laliberté [2021], Nakamura, Sigurdsson, and Steinsson [2022], and Sharkey [2016]. School choice could weaken the link between place and opportunity by allowing families to access stronger schools beyond their neighborhoods. Our results identify a mechanism through which that link may persist: poor transit access can make top-rated schools difficult to reach from some neighborhoods.

The remainder of the paper proceeds as follows. Section 2 describes the Chicago high school choice setting and Section 3 describes the data. Section 4 presents descriptive evidence on racial and socioeconomic differences in access to top-rated schools. Sections 5 and 6 present the model, identification, estimates, and robustness exercises. Section 7 quantifies how travel time contributes to racial differences in choices, and Section 8 presents targeted transportation counterfactuals. Section 9 concludes.

2 Setting: High School Choice in Chicago Public Schools

2.1 District Background and High School Choice

Chicago Public Schools (CPS) is one of the largest public school districts in the United States, serving approximately 380,000 students during the period we study. This racially diverse district primarily educates minority students: 47% of enrolled students are Hispanic, 38% are African-American, 10% are white, and 4% are Asian. The district tends to serve lower-income students as 80% of students are eligible for free or reduced-price lunch.

Like many large urban districts, CPS has gradually expanded school choice over the last several decades, weakening the traditional link between residential location and school assignment. Students may apply to schools throughout the district through a citywide choice system rather than being restricted to their assigned neighborhood school. During the period we study, students could apply to over one hundred high schools that differed in academic performance, programs offered, student demographics, and location.

Beginning with admissions for the 2018-2019 school year, CPS introduced GoCPS, a centralized application platform for most district high school programs. Previously, applications to many programs were administered separately by individual schools, requiring families to navigate multiple application systems and deadlines. GoCPS consolidated these processes and standardized information about available schools. We focus on this first GoCPS admissions cycle.

GoCPS distinguishes between *choice* and *selective enrollment* programs. Choice programs generally have no minimum eligibility requirements, although some have eligibility restrictions or admissions

priorities. Selective enrollment programs require an admissions test and are ranked on a separate application list. We focus on choice programs because they account for most applications and represent the relevant margin for studying broad access to high school options. This restriction retains 95 percent of applicants in our sample. Students apply in the fall of 8th grade for enrollment the following school year and retain the option of attending their assigned neighborhood high school.

2.2 Information Available to Families and School Quality Ratings

GoCPS also serves as an information source for families. School profile pages provide information on academic performance, graduation rates, college enrollment outcomes, programs offered, admissions requirements, and student demographics. Thus, families observe a common set of school characteristics before submitting their ranked lists.

During the period we study, CPS summarized school performance using the School Quality Rating Policy (SQRP), which assigned schools to five ordered performance categories: Level 1+, Level 1, Level 2+, Level 2, and Level 3. The ratings were based on a weighted combination of academic performance, student growth, graduation, dropout, attendance, college enrollment, and school climate measures. Level 1+ was the highest rating. SQRP ratings were prominently displayed within GoCPS and provided families with a simple summary of school performance. Among the 126 choice high schools in our sample, 32 received a Level 1+ rating, 25 received Level 1, 35 received Level 2+, 33 received Level 2, and 1 received Level 3. Level 1+ schools differ substantially from lower-rated schools in average test scores, graduation rates, and college enrollment rates.

This common information environment is useful for interpreting our results because families could observe and compare school quality before submitting their rankings. Of course, information frictions or differences in attention could still arise across families.[1] However, a salient measure of school quality was available when families ranked schools.

[1] As we note in the Introduction, a rich literature studies the role of information frictions in school choice [e.g. Corradini, 2024].

2.3 School Assignment Through GoCPS

Students submit separate rank-ordered lists for choice and selective enrollment programs through GoCPS and can rank up to 20 choice programs and 6 selective enrollment programs. Our analysis focuses on Round 1 choice applications because they contain the initial rankings submitted before offers, waitlists, and subsequent revisions to the choice set.[2] Section 3 describes how we aggregate program-level rankings to the school level and construct each student's feasible choice set.[3]

CPS assigns students to programs using a student-proposing deferred acceptance algorithm. Students receive an offer from the highest-ranked program for which they qualify and have sufficient priority. Priorities for oversubscribed programs may depend on factors such as sibling status, proximity, or program-specific eligibility criteria. Students receive one offer from the choice list and, if applicable, a separate offer from the selective enrollment list.

This assignment mechanism matters for our empirical strategy because it gives students incentives to rank programs in their true order of preference. Ranking a less-preferred program above a more-preferred one does not increase the probability of receiving an offer. GoCPS explicitly instructed families to rank programs in their true order of preference and explained that "ranking a school higher does not increase your chances of getting an offer." Figure A.1 shows these instructions.

Our empirical approach uses the ranked lists as revealed preference data. When a student ranks one program above another, the ranking indicates that the student preferred the higher-ranked option. For this reason, our analysis does not require us to model the full assignment process or to combine choice and selective enrollment offers into a single final allocation. Instead, we use the Round 1 choice rankings to estimate the value families place on each school and the disutility associated with travel time from home to school as we describe below in section 5.

[2]Students may apply in later rounds if they are unsatisfied with their Round 1 placement or do not receive an offer. However, only programs with remaining seats after Round 1 participate in later rounds.

[3]Students apply to programs rather than directly to schools, and a single school may offer multiple programs. For the purposes of the analysis, we aggregate choices from the program level to the school level. A student is coded as applying to a school if they applied to at least one program offered by that school. If a student applies to more than one program at the same school, we assign the school the rank of the highest-ranked program at that school.

2.4 Transportation and School Access

A distinctive feature of the Chicago high school choice system is that high school students generally do not receive district-provided yellow bus transportation.[4] High school students in CPS therefore typically rely on Chicago's public transportation network to travel to school. As a result, access to schools depends not only on geographic distance but also on the structure of the public transit network. Two students living similar distances from a school may face very different commute times because of differences in bus routes, subway ('L') routes, transfers, and service frequency. Schools that are formally available through the citywide choice system may therefore be much easier to reach from some neighborhoods than from others.

Figure 1 shows one-way public transit time to the nearest top-rated (Level 1+) school. Some neighborhoods can reach a Level 1+ school in 10 to 15 minutes, while others, particularly on the South Side, face transit times of 30 minutes or more. The map also shows that straight-line distance is not a sufficient measure of access because transit times depend on the structure of the public transportation network. We describe the construction of these travel times in the next section. Although all students participate in the same citywide choice system, schools may be much easier to reach from some neighborhoods than from others. In the next section, we document these differences in access to highly rated schools and their relationship with school choice behavior.

3 Data

This section describes the data on students' school choices, school characteristics, and travel times. For each student i and school j, we observe whether the student listed the school and, if so, where the school appeared on the student's rank-ordered list, the school's characteristics, and the travel time from the student's residential area to the school.

[4] CPS facilitates public transit use through Ventra cards, the fare cards used to ride CTA buses and trains. Yellow-bus service for high school students is generally limited to students with a legally protected transportation entitlement, such as transportation specified in an IEP or 504 plan, homelessness/STLS status, or another state or federal legal basis. General education students choosing GoCPS programs are therefore typically expected to rely on CTA, walking, family transportation, or other modes.

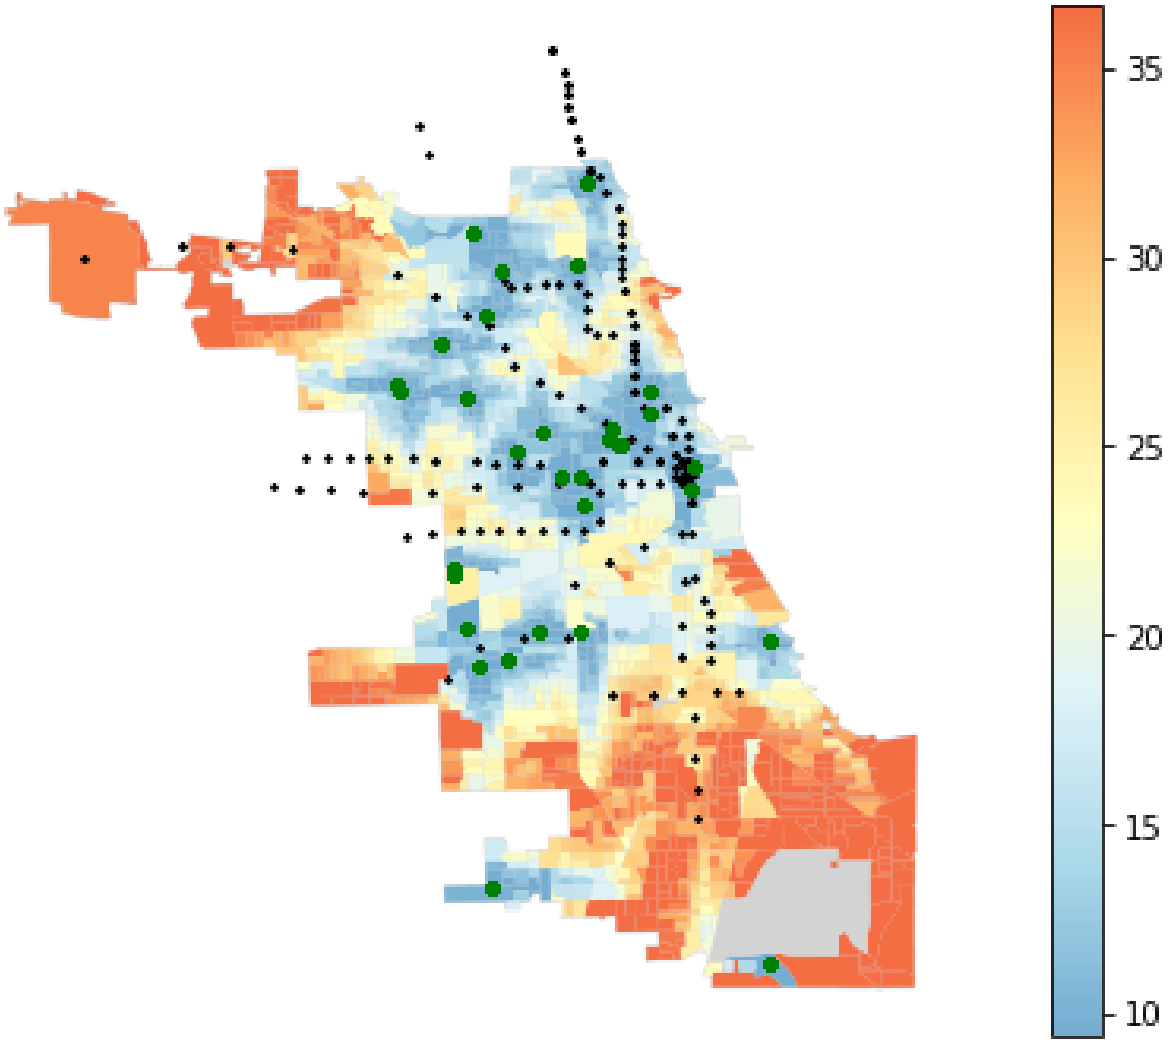


Figure 1: Transit Time to Nearest Level 1+ High School

Notes: This figure shows the one-way public transit travel time, in minutes, from each census block group to the nearest Level 1+ high school. Level 1+ schools are shown with green dots and L stops are shown with black dots. We describe the construction of the travel-time data in section 3.

3.1 Student Application Data

Our primary data source is the administrative record of high school applications submitted through GoCPS for the first cohort using the centralized platform. The data include 22,902 incoming ninth graders. For each student, we observe the submitted ranked list, residential census block group, race and ethnicity, free or reduced-price lunch status, test scores, and assigned neighborhood school.

We focus on Round 1 applications to choice programs. As discussed in section 2, these rankings capture students' initial preferences before offers, waitlists, and later-round revisions to the choice set. Restricting attention to students with choice program applications retains 95 percent of applicants in our sample. These students are more likely to be low-income and minority than students who apply only to selective enrollment programs. The application data provide the revealed preference input for the demand model. In the descriptive analysis, we use the rankings to measure how many schools students list and whether their first-ranked school is top-rated. In the model, we use the ordinal rankings to estimate how families trade off school value against travel time.

3.2 School Programs, School-Level Choices, and School Characteristics

Students apply to programs rather than directly to schools, and a single school may offer multiple programs. We conduct the analysis at the school level because travel time varies at the school rather than program level. We therefore aggregate program rankings to school rankings. A student is coded as applying to a school if the student ranked at least one program at that school. If the student ranked multiple programs at the same school, we assign the school the rank of the highest-ranked program.

We construct each student's feasible choice set as the set of schools offering at least one choice program for which the student is eligible. We then merge these school-level choices with school characteristics, including SQRP ratings, average SAT scores, graduation and college enrollment rates, value-added measures, and student demographic composition. CPS used SQRP ratings to assign schools to performance levels. In the empirical analysis, we refer to Level 1+ schools as

top-rated schools.[5]

3.3 Travel Time Data

A key data contribution is a matrix of travel times between student residential areas and high schools. We use the Google Maps Directions API to measure one-way travel time from every Chicago census block group centroid to every high school. Travel time is calculated for arrival at 8:30 a.m. on a weekday, corresponding to the relevant morning commute for high school students.

We collect travel times for public transit and driving. Public transit time is the main measure used in the demand model because CPS high school students generally do not receive district-provided yellow bus transportation. Driving time is used in counterfactual exercises that approximate faster point-to-point transportation, such as targeted express bus service. This procedure generates approximately 300,000 census-block-group-by-school travel-time pairs. After assigning each student to a census block group based on home address, we measure public transit and driving time to each school in the student's choice set. Google Maps did not return public transit times for 4 percent of student-school pairs, so we remove those school alternatives from the student's choice set.[6]

4 Descriptive Evidence on School Quality, Geography, and Application Behavior

This section presents descriptive evidence on school quality, residential location, application behavior, and public transit access. We document that demographic groups differ in the ratings of the schools they attend and that the groups least likely to rank top-rated schools first also face longer public transit times to reach those schools. These patterns do not establish that transportation causes differences in school choices. However, they show that differences in application behavior coincide with large differences in access to top-rated schools.

[5] We use SQRP ratings from the 2017–2018 school year because these ratings were in place when students submitted their GoCPS applications.

[6] When transit time to a school is missing, we make no revealed-preference comparisons involving that school for the student. We are unaware of any systematic pattern in missing travel times that would introduce bias. If transit time to the student's first-choice school is missing, we drop the student from the analysis; this affects 0.2 percent of students.

Figure 2 shows the distribution of the SQRP index of the school attended, separately by race and income. These attendance patterns reflect the combined effects of application behavior, school preferences, admissions priorities, capacities, and transportation costs. More than half of Black students attend Level 2 or Level 2+ schools, while more than half of white students attend Level 1+ schools. Hispanic students also attend higher-rated schools than Black students, with nearly half attending Level 1+ schools. The FRPL panel, using free or reduced-price lunch eligibility as a proxy for low income, shows a similar pattern: FRPL-eligible students are more likely to attend lower-rated schools than non-FRPL students.

We next describe the geography of schools and students. Figure 3 plots high schools by SQRP rating over majority-ethnicity zones. Level 1+ schools are more concentrated in the downtown Loop area and to the north, while the South Side has fewer Level 1+ schools relative to the number of students living there. This spatial pattern is relevant because many Black students live on the South Side, Hispanic students are more concentrated in the southwest and northwest, and white students are more concentrated north of the Loop.

The geographic distribution of schools is informative but does not fully measure access. In this setting, the relevant measure is public transit travel time rather than straight-line distance. A nearby school may require a long commute because of indirect routes or transfers, while a more distant school may be easier to reach along a direct transit route.

Table 1 summarizes application behavior and travel-time access by demographic group. Panel A shows that 56 percent of students rank a Level 1+ school first, but this rate differs substantially by race. Among Black students, 38 percent rank a Level 1+ school first, compared with 62 percent of Hispanic students, 62 percent of white students, and 73 percent of students in the other-race category. The same pattern appears when we combine the two highest SQRP categories: 58 percent of Black students rank either a Level 1 or Level 1+ school first, compared with 81 percent of Hispanic students, 94 percent of white students, and 91 percent of students in the other-race category. Black students also rank more schools on average than other groups, listing 7.52 schools compared with 5.89 for Hispanic students and 3.02 for white students. Thus, their lower rate of ranking a Level 1+ school first does not reflect lower overall application intensity.

Students routinely consider schools well outside their immediate neighborhoods. The median driving distance to the farthest school on a student's rank list is 8.2 miles, and at the 75th percentile it is 13.0 miles. Students' first choices also involve substantial commutes: the mean public transit time to the first-ranked school is 32 minutes overall, 37 minutes for Black students, 29 minutes for Hispanic students, and 30 minutes for white students. Because students include distant schools on their rank lists, their choices do not appear to be heavily constrained by a lack of awareness of out-of-neighborhood options.

Panel B shows that Black students also face longer travel times to top-rated schools. The mean public transit time to the nearest Level 1+ school is 25 minutes for Black students, compared with 17 minutes for Hispanic students, 21 minutes for white students, and 19 minutes for students in the other-race category. The larger differences in public transit time indicate that unequal access is not only a function of where schools are located, but also of how those locations are connected to students' neighborhoods by the transit network.

The gap between public transit and driving time to the nearest Level 1+ school is also largest for Black students: 15 minutes, compared with 10 minutes for Hispanic students, 13 minutes for white students, and 11 minutes for students in the other-race category. This gap suggests that transportation policies that move commute times closer to driving times, such as targeted express buses, may have the greatest potential to change demand for top-rated schools among Black students. The magnitude of that response depends on the disutility of travel time estimated in the demand model.

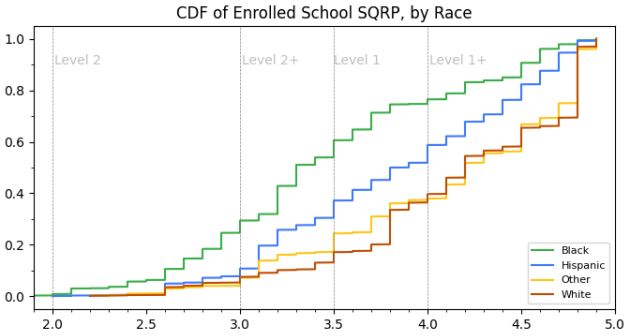
The mean travel times in Table 1 show that access differs across groups, but they do not indicate whether these differences extend throughout the distribution. Figure 4 plots the cumulative distribution of public transit time to the nearest Level 1+ school by race. The distribution for Black students is shifted toward longer travel times relative to the other groups. Nearly one-third of Black students face a one-way public transit trip of more than 30 minutes to the nearest Level 1+ school.

Taken together, the descriptive evidence shows that racial differences in demand for top-rated schools coincide with racial differences in public transit access: Black students are less likely to rank a Level 1+ school first and face longer transit times to the nearest Level 1+ school. However,

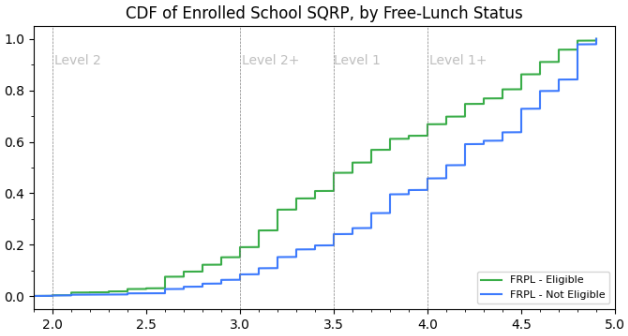
these patterns cannot determine whether differences in rankings reflect travel time, preferences over school characteristics, information, or other factors. The demand model in the next section separates the value families place on schools from the disutility of travel time, allowing us to quantify the role of transportation in demand for top-rated schools.

Figure 2: Distribution of SQRP Index of School Attendance by Demographics

(a) By Race

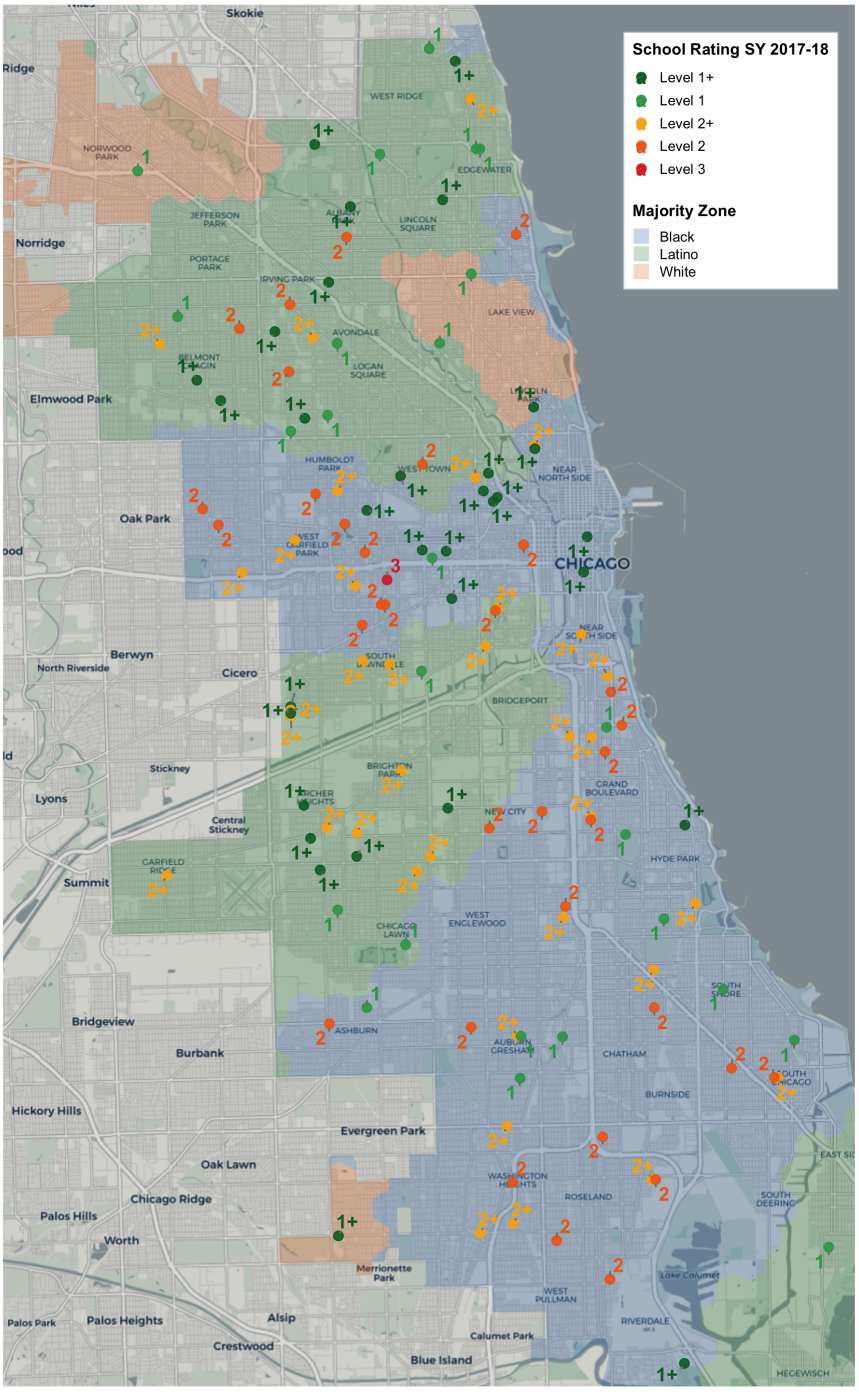


(b) By Free/Reduced Price Lunch Status



Notes: These figures are the cumulative distribution functions of the index underlying the SQRP status of the school attended.

Figure 3: Distribution of Schools by Majority Ethnicity Zone



Racial majority zones are constructed by first assigning schools to 1000-meter hexagonal grid cells based on their student demographic composition (>50% threshold for Black, Latino, or White students), then applying an iterative spatial smoothing algorithm that assigns empty cells and resolves isolated cells based on their neighbors' majority status. The resulting continuous zones reflect predominant racial patterns while maintaining consistency with the underlying school-level demographic data.

Table 1: Descriptive Evidence: Means of Student Rankings and Travel Times

	All	Black	Hispanic	White	Other race	Missing race	FRPL	Non-FRPL
Panel A: Application Behavior								
Number of Schools Ranked	6.05	7.52	5.89	3.02	4.86	3.97	6.53	4.13
Percent with First Choice a Level 1+ School	55.5	38.3	62.2	61.7	73.1	77.2	51.9	70.0
Percent with First Choice a Level 1 School	19.6	19.4	18.5	31.9	17.5	13.2	19.3	20.5
Panel B: Travel-Time Access								
Transit Time to Neighborhood School (minutes)	17.3	20.1	15.1	18.7	18.7	17.0	17.2	18.1
Transit Time to Nearest Level 1+ School (minutes)	20.3	25.4	16.9	21.4	18.7	19.4	20.1	21.3
Drive Time to Nearest Level 1+ School (minutes)	8.2	10.4	6.6	8.5	7.9	8.0	8.1	8.5
Transit Minus Drive Time to Nearest Level 1+ School	12.2	15.0	10.3	12.9	10.8	11.4	12.0	12.8
Drive Time to Neighborhood School (minutes)	6.6	7.5	5.9	6.8	7.0	6.7	6.6	6.8
First Choice Transit Time (minutes)	32.2	37.3	29.0	30.4	31.0	33.5	31.8	33.8

Notes: The table reports student-level means for the full sample of 2017 applicants. Percent variables are multiplied by 100.

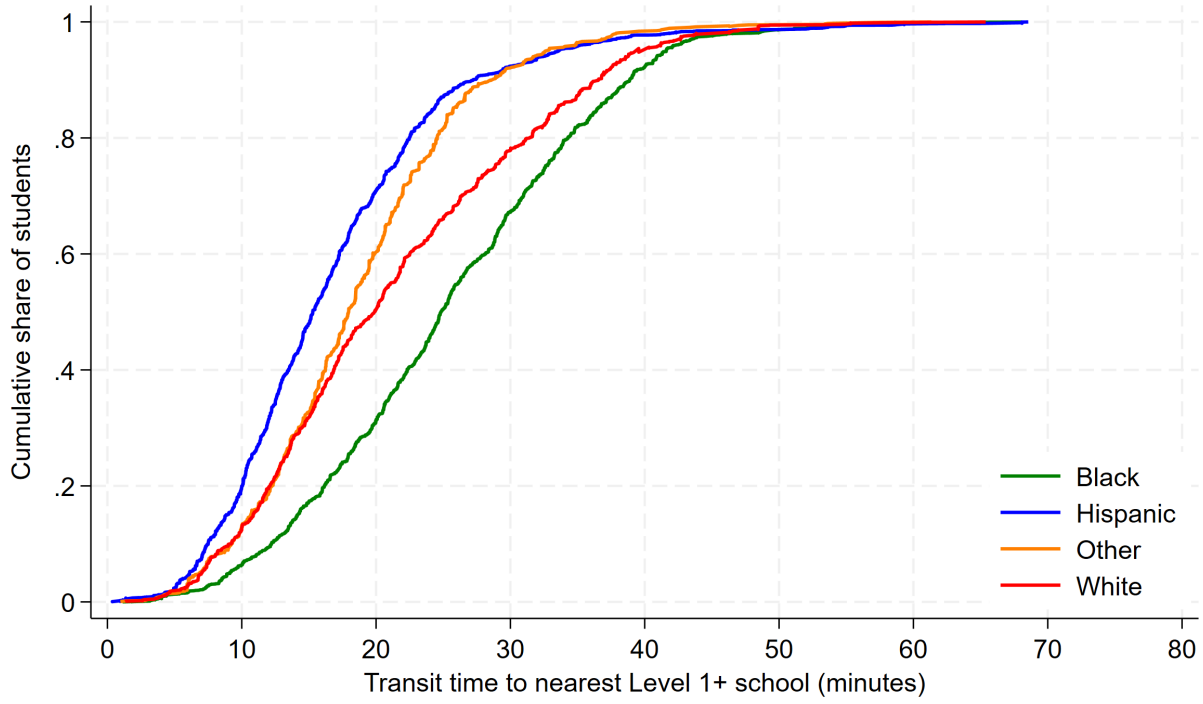


Figure 4: Distribution Across Students of Transit Time to Nearest Level 1+ School

Notes: This figure shows cumulative distribution across all students in our sample of the one-way public transit travel time, in minutes, from the student's residential census block group to the nearest Level 1+ high school. We describe the construction of the travel-time data in section 3.

5 Model

We model demand for schools as revealed by the ranked lists that students submit through GoCPS. The model separates the value families place on each school from the time cost of reaching it. Using ranked-choice data and student-school public transit times, we estimate how families trade off the value of a school against the cost of reaching it when forming their preference lists. As described in sections 2 and 3, we aggregate program-level rankings to the school level, construct student-specific choice sets, and interpret the resulting rankings as revealed-preference data. CPS uses a student-proposing deferred acceptance algorithm, which gives students incentives to rank programs in their true order of preference [Abdulkadiroglu and Sonmez, 2003].

5.1 Utility Specification

Each school provides a bundle of characteristics that families may value. Some are observed in our data, including SQRP ratings, graduation rates, college enrollment rates, peer composition, programs offered, and location. Others may be observed by families but are difficult for the researcher to measure, such as school climate, safety, reputation, extracurricular offerings, or the extent to which a school develops non-cognitive skills. Rather than specify which characteristics enter utility and how families value them, we use a flexible school-specific mean utility. This allows the model to recover the average value families place on each school after accounting for travel time.

Let k denote a demographic group such as race. We allow both preferences over schools and the disutility of travel time to vary across groups because families in different demographic groups may value school attributes and commute time differently. The average utility that school j provides to students in group k is denoted by δ_{jk}. This parameter captures the value of school quality, peers, programs, location, reputation, and any other school attributes that are valued by students in group k and common across students in the group. Student i's utility from attending school j is

$$U_{ij} = \delta_{jk} - \lambda_k d_{ij} + \epsilon_{ij}.$$

The variable d_{ij} is the public transit travel time from student i's home to school j. The parameter λ_k measures the disutility of travel time for students in group k. The idiosyncratic term ϵ_{ij} captures student-specific factors that affect the value of school j for student i and are observed by the family but not by the researcher. We assume that ϵ_{ij} is independently distributed Type I extreme value.

5.2 Choice Sets and Estimation

Let J_i denote the student-specific choice set. We begin with the 126 CPS high schools available through the choice process and adjust this set to reflect student-specific eligibility. Because only a small number of schools have requirements based on academic performance, gender, or other criteria, J_i includes nearly all CPS high schools for most students: the 25th percentile of $|J_i|$ is 114

schools, the median is 119, and the 75th percentile is 121.[7] The data do not allow us to model the decision to leave CPS for private or parochial schools. Given the income composition of the estimation sample, we assume that this outside option is not relevant for a meaningful share of students and focus on substitution across public high schools.

If student i ranks school j first, we interpret this ranking as revealing that school j provides higher utility than every other school in J_i:

$$U_{ij} \geq U_{i\ell} \quad \forall \ell \neq j, \ell \in J_i.$$

Given the utility specification above and the Type I extreme value assumption on ϵ_{ij}, these pairwise utility comparisons imply the following probability that student i ranks school j first:[8]

$$P_{ij} = \frac{\exp(\delta_{jk} - \lambda_k d_{ij})}{\sum_{\ell \in J_i} \exp(\delta_{\ell k} - \lambda_k d_{i\ell})}.$$

Our baseline estimates use each student's first-ranked school. We estimate the parameters by maximum likelihood using the first-ranked choices of all students in the estimation sample.[9] Let 1_{ij} be an indicator equal to one if student i ranks school j first. The log-likelihood is

$$LL(\boldsymbol{\delta}, \lambda) = \sum_{i=1}^{N} \sum_{j \in J_i} 1_{ij} \log(P_{ij}).$$

[7] Because students submit rankings at the program level but the model is estimated at the school level, we include a school in J_i if student i is eligible for at least one choice program at that school. We also include student i's assigned neighborhood high school in the choice set.

[8] We also estimate mixed-logit specifications that allow preferences for schools and the disutility of travel time to vary across students. In these specifications, the preference parameters $\gamma \equiv (\lambda, \delta)'$ are drawn from a distribution $f(\gamma \mid \mu, \theta)$. This allows for richer preference heterogeneity, consistent with evidence that families may trade off school quality and distance differently [Abdulkadiroglu et al., 2020]. The first-choice probability is then

$$P_{ij} = \int \frac{\exp(-\lambda d_{ij} + \delta_j)}{\sum_{\ell \in J_i} \exp(-\lambda d_{i\ell} + \delta_\ell)} f(\gamma \mid \mu, \theta) d\gamma.$$

[9] In robustness specifications, we also use information from students' second-ranked schools. If student i ranks school j first and school m second, the second ranking implies that school m provides higher utility than all schools in the choice set other than the first-ranked school:

$$U_{im} \geq U_{i\ell} \quad \forall \ell \neq j, m, \ \ell \in J_i.$$

Thus, the second-ranked school provides additional information about family preference parameters.

The baseline specification allows both school mean utilities and the disutility of travel time to vary by demographic group. We use alternative specifications to evaluate whether the estimated role of travel time is sensitive to assumptions about preference heterogeneity or to the use of information beyond students' first-ranked schools.

5.3 Identification

The model identifies the disutility of travel time from variation in public transit times across students, holding fixed each school's mean utility. For a given school, students in different neighborhoods face different travel times because of their locations in the public transit network. Some have direct bus or L connections, while others require transfers or indirect routes. Because travel time is not simply a smooth function of distance, nearby neighborhoods can face different travel times to the same school. This local variation helps separate the effect of commute time from broader spatial differences in residential location.

The model also uses variation across schools within a student's choice set. Each student faces schools that differ in mean utility and travel time from home. The ranked list reveals how families trade off these two dimensions. A student who ranks a farther school above a closer school provides information about how much additional school value is needed to offset the additional travel time. A student who ranks a closer school above a more distant school with higher mean utility provides information about the cost of travel time. Combining these comparisons across students, schools, and neighborhoods identifies both the school mean utilities and the disutility of travel time.

5.4 Threats to Identification and Interpretation

The first maintained assumption is that submitted rankings correspond to families' true ordinal preferences over schools. This assumption is standard in empirical analyses that use rank-order lists from centralized school choice mechanisms to estimate preferences, particularly when the mechanism gives families incentives to rank schools truthfully [Abdulkadiroğlu et al., 2017, Fack et al., 2019, Agarwal and Somaini, 2020]. The deferred acceptance mechanism has this property if families understand the incentives created by the mechanism [Abdulkadiroglu and Sonmez, 2003].

It is critical that families do not skip over more-preferred schools because they believe doing so creates a strategic advantage in gaining admission to a less-preferred school. The GoCPS website attempted to make this incentive clear. The instructions provided examples of how the assignment mechanism operates and stated: "Ranking a school higher does not increase your chances of getting an offer. Make sure you rank programs in the true order of preference in attending." (see Figure A.1) Fack et al. [2019] show that even under deferred acceptance, applicants may omit highly preferred programs they perceive as effectively unattainable in a behavior they call "skipping the impossible." This concern is likely limited in our setting because we study choice programs rather than the separate selective enrollment list; choice programs generally have no minimum academic eligibility requirements, and nearly all choice schools are feasible for most students.

A related concern is that families may omit more-preferred schools if they are constrained in the number of programs they can rank, which can affect incentives and submitted rankings [Calsamiglia et al., 2010, Agarwal and Somaini, 2020]. This is unlikely to be important in our setting because families could rank up to 20 choice programs and the average family ranked just over six programs. Moreover, only 5 percent of students submitted 15 or more choices, and only 0.2 percent submitted the maximum of 20 choices.

The central threat to identification is residential sorting. Families with stronger unobserved preferences for school quality may choose neighborhoods with better transit access to high-value schools. If so, the estimated travel-time parameter may partly capture unobserved preferences for school quality. This concern is common in the school choice and urban economics literature because residential location is itself a choice. In subsection 6.3, we assess this concern directly using address histories before high school application and specifications that allow unobserved preferences for school quality to be correlated with the disutility of travel time.

A separate concern is that the estimated disutility of travel time may partly absorb a lack of information about distant schools. However, as documented in section 4, the typical student ranks at least one school more than 8 miles from home, suggesting broad geographic awareness. This provides some confidence that the travel-time parameter primarily captures the friction of the commute rather than severe information deficits about the broader set of public high schools.

The final interpretive issue concerns the school mean utilities. The parameters δ_{jk} should be interpreted as demand-side valuations, not causal measures of school effectiveness. They capture the average value that students in group k place on school j, after accounting for travel time. These values may reflect academic quality, peers, programs, safety, reputation, location, or other characteristics that families observe or value. They therefore summarize demand for schools rather than the causal effect of attending those schools.

6 Model Estimates

This section presents estimates from the school demand model. We estimate the model on a random subsample of 2,500 incoming freshmen from the approximately 22,000 students in the first GoCPS cohort.[10] We begin with a pooled specification and then estimate the model separately by race. The race-specific estimates are used in the counterfactuals below to study racial differences in access and place-based transportation policies. We then present specification checks and assess the central threat to identification: residential sorting.

6.1 Benchmark Estimates

Figure 5 presents estimates from the pooled specification. The vertical bars show each school's estimated mean utility, $\hat{\delta}_j$, relative to the omitted benchmark school and are ordered from lowest to highest. Because utility units have no standalone interpretation, the horizontal dashed line reports the estimated disutility of one hour of one-way public transit travel time, $\hat{\lambda}$.

The estimates show substantial variation in the value families place on schools. Mean utilities vary smoothly across most of the school distribution, indicating that families distinguish among CPS high schools even after accounting for travel time. Four schools have especially low estimated mean utilities, which we return to below.

The estimated disutility of travel time is economically large. One hour of one-way public transit travel time is equivalent to 4.47 utils, with a standard error of 0.09. To interpret this magnitude,

[10] In the mixed-logit robustness specifications with random coefficients, we limit the estimation sample because of computational constraints on the secure data server. To maintain a consistent estimation sample across specifications, we use the same random subsample for all specifications.

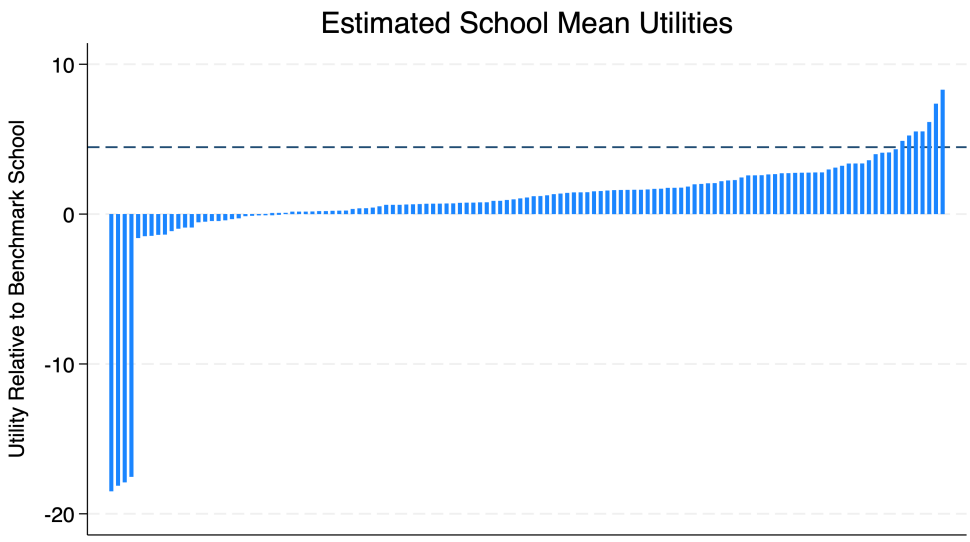


Figure 5: Pooled Model Estimates

This figure presents estimates from the benchmark model. The horizontal line represents the estimated disutility of travel time ($\hat{\lambda}$), measured in utils per hour of one-way transit travel. The vertical bars show each school's estimated mean utility $\hat{\delta}_j$ (in utils) relative to the excluded benchmark school, with schools ordered by estimated mean utility.

compare two schools with mean utilities that differ by 4.47 utils. A family with no idiosyncratic preference for one school over the other would be indifferent between the lower-utility school with no commute and the higher-utility school with a one-hour one-way commute. In Figure 5, this difference is approximately the gap between the 22nd-lowest and 8th-highest mean utility schools. Similarly, 30 minutes of one-way travel time is equivalent to the difference in mean utility between the 20th and 75th percentile school. The magnitude of the travel time estimate implies that the differences in transit access documented in section 4 are large enough to change which schools families rank first.

We also estimate a mixed-logit version of the benchmark model that allows the disutility of travel time to vary across students. In this specification, the travel-time coefficient is drawn from a Normal distribution with mean μ_λ and standard deviation σ_λ. This allows students to differ in how sharply they trade off school mean utility against commute time. Appendix Figure A.2 compares the fixed-coefficient and random-coefficients estimates in the pooled sample. The estimated mean disutility of travel time and the school mean utilities are very similar across the two specifications.

We use this random-coefficients specification for the race-specific estimates below.

Importantly, the school mean utilities are estimated only from the revealed-preference information in students' ranked lists, after accounting for travel time. We do not construct them from observed school characteristics or require the researcher to specify which characteristics families value. This flexibility is useful because families may value many attributes that are difficult to measure, including safety, reputation, school climate, peers, programs, or non-academic outcomes. Because we also observe several school characteristics that families are likely to value, we can assess whether the estimated mean utilities line up with observable measures of school quality.

Table 2 reports descriptive regressions of the estimated school mean utilities on school characteristics. Schools with higher estimated mean utilities tend to have lower shares of students eligible for free or reduced-price lunch, higher average SAT scores, higher four-year graduation rates, higher college enrollment rates, and higher math value-added[11]. The estimates are not intended to recover causal effects of these characteristics on demand. Rather, they show that the schools assigned higher mean utility by the model also tend to have observable characteristics that families plausibly value.

As an additional check on interpretation, we compare the estimated mean utilities to external school ratings. Schools rated poorly on GreatSchools.org fall in the bottom quartile of the estimated mean-utility distribution in both the pooled and Black-student models. Highly-rated choice schools also tend to rank highly in the estimated mean-utility distribution in both the pooled and Black-student models (Chicago Tribune, 2024). This provides additional reassurance that the estimated mean utilities are capturing meaningful differences in how families value schools, rather than only fitting idiosyncratic features of the application data.

We next estimate the random-coefficients model separately for Black and Hispanic students. To do so with sufficient precision while staying within the computational constraints of the secure data server, we draw separate random samples of 2,500 Black students and 2,500 Hispanic students rather than using only the Black and Hispanic students included in the pooled 2,500-student sample. These race-specific estimates are useful for two reasons. First, students in different racial

[11] We describe the construction of value-added measures in Appendix A.3.

Table 2: Estimated Mean School Utility: Correlates with Characteristics

	$\hat{\delta}_j$	$\hat{\delta}_j$	$\hat{\delta}_j$	$\hat{\delta}_j$	$\hat{\delta}_j$	$\hat{\delta}_j$
Pct FRPL	−0.071					−0.032
	(0.011)					(0.013)
Avg SAT Score		0.012				0.006
		(0.001)				(0.003)
4-year Cohort Grad Rate			0.082			0.053
			(0.009)			(0.013)
College Enrollment Rate				0.052		−0.001
				(0.009)		(0.011)
Math Value-Added					1.902	0.735
					(0.915)	(0.987)
Constant	6.680	−10.470	−5.666	−2.711	0.462	−6.618
	(1.019)	(1.244)	(0.689)	(0.593)	(0.146)	(2.762)
Observations	126	122	117	109	125	108

Notes: This table shows results from descriptive OLS regressions of the relationship between school mean utility and school characteristics. An observation is a school. Each column corresponds to a separate regression where the dependent variable is the school's estimated mean utility $\hat{\delta}_j$ from the benchmark model with fixed coefficients. Sample size can vary if characteristics are not available in 2017 for the school. Standard errors are in parentheses.

groups may value school attributes differently, including peer composition, location, programs, or other school characteristics. Second, the counterfactuals below use these estimates to study racial differences in access and place-based transportation policies targeted to neighborhoods that may be predominantly Black or Hispanic.

Because separately estimated logit models may place coefficients on different utility scales, we compare school valuations across groups using the ratio of each school mean utility to the estimated mean disutility of travel time.[12] Figure 6 reports $\hat{\delta}_{jk}/\hat{\mu}_{\lambda k}$ separately for Black and Hispanic students. This ratio expresses each school's mean utility in travel-time units and allows comparisons across groups without relying on the scale of utility in either race-specific model. The bars are ordered by the estimated ratios for Black students. For 78 schools, we report estimates for both Black and Hispanic students; these schools appear with two bars. For the remaining schools, we report an estimate only for the racial group with enough applicants to estimate the school mean

[12]In a logit model, coefficients are identified only relative to the scale of utility, which is normalized through the variance of the idiosyncratic utility shock. Dividing the estimated school mean utility by the estimated mean disutility of travel time expresses school valuations in travel-time units and removes this arbitrary scale difference.

utility precisely, so these schools appear with a single bar.[13] Note that when we estimate mean utility separately by race, the four very negative mean utilities in the pooled specification are no longer distinctly different from the other schools, which reflects the impact of constraining mean utility to be equal across racial groups for schools that are predominantly one race.

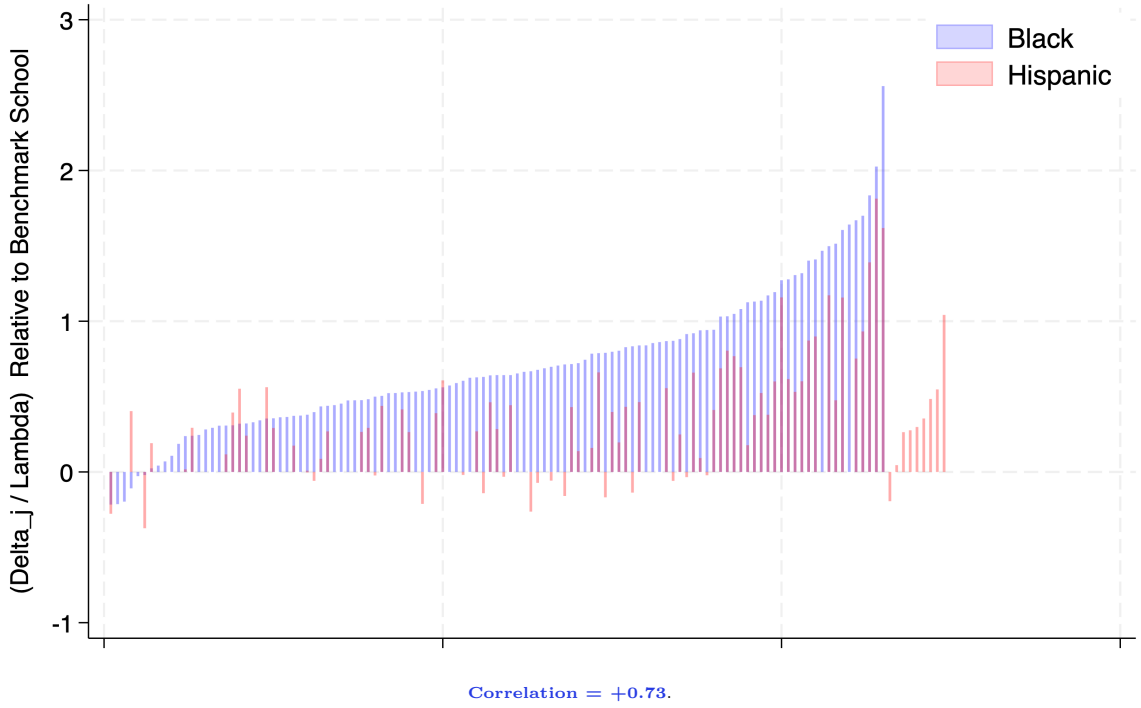


Figure 6: Race-Specific Estimates of School Mean Utilities Relative to Travel-Time Disutility

Notes: This figure reports $\hat{\delta}_{jk}/\hat{\mu}_{\lambda k}$ from separate random-coefficients models for Black and Hispanic students. Each bar is a school. The bars are ordered by the estimated ratios for Black students. No coefficient is reported for school-race cells with too few rankings by students of that race to estimate the school mean utility precisely. The estimated model is a mixed logit: $U_{ij} = \delta_{jk} + \beta_{ki} \cdot 1\{j = i\text{'s zoned school}\} - \lambda_{ki} d_{ij} + \epsilon_{ij}$. For each race k, school mean utilities δ_{jk} are constant across students, while the local-school preference β_{ki} and travel-time disutility λ_{ki} vary across students.

The estimates show substantial variation across schools within each racial group. Among Black students, some schools have much higher estimated value relative to travel time than others; the same is true among Hispanic students. At the same time, the ordering of schools is strongly related

[13]The single-bar schools also illustrate why race-specific estimates are important for the counterfactuals below. Some CPS high schools draw applicants overwhelmingly from one racial group. For example, Manley was 96.0 percent Black and 3.4 percent Hispanic among high school students in 2017, while Kelvyn Park was 90.2 percent Hispanic and 6.9 percent Black. In these cases, the data provide too little information to estimate the non-majority group's school mean utility precisely. When evaluating policies targeted to predominantly Black or Hispanic neighborhoods, we therefore use the race-specific estimates for the relevant group rather than imposing pooled valuations across all students.

across groups. Among the 78 schools, the correlation between the two sets of ratios is 0.73. Thus, the groups do not value every school identically, but schools with higher estimated value relative to travel time for Black students also tend to have higher estimated value relative to travel time for Hispanic students.

This strong correlation is important for interpreting the racial differences in first choices documented above. The lower rate at which Black students rank top-rated schools first cannot be explained simply by Black and Hispanic students valuing entirely different schools. The broad ranking of schools is similar for Black and Hispanic students. The transportation differences documented in section 4 therefore remain a central part of the explanation. The counterfactuals below quantify how much demand for top-rated schools would change if Black students faced different travel times.

6.2 Robustness of Estimates

We next assess whether the pooled estimates are sensitive to several modeling choices. We vary the benchmark specification in five ways: allowing travel-time disutility to vary across students, adding a separate preference for a student's assigned neighborhood school, allowing travel-time disutility to be nonlinear, using both first and second choices in the likelihood, and including roadway distance as a separate spatial control. These exercises assess robustness to preference heterogeneity, functional form, additional information from ranked lists, and broader physical separation between students and schools.

Appendix Figures A.2–A.4 and Appendices A.1–A.2 report the results. Across specifications, the estimated school mean utilities and disutility of travel time are very similar to the benchmark estimates. Allowing travel-time disutility to vary across students produces nearly the same average travel-time cost and school mean utilities as the fixed-coefficient benchmark. Adding a separate preference for the neighborhood school allows students to value that school more highly than other students value the same school, for example because more middle school friends may attend. The resulting estimates are again similar to the benchmark model. Allowing travel-time disutility to be quadratic does not materially change the implied tradeoff between school value and travel time over the relevant range of commutes. Finally, incorporating both first and second choices into the

likelihood function produces estimates close to the first-choice benchmark. These results indicate that the central conclusion is not driven by a specific modeling choice: public transit travel time is large relative to school valuations and can therefore affect which schools students rank first.

A further question is whether $\hat{\lambda}$ captures frictions specific to the public transit network – including indirect routes, transfers, and waiting times – or broader spatial separation between a student's home and a school. Families may be reluctant to attend schools far from home for reasons beyond travel time, including unfamiliarity with distant parts of the city or attachment to their own neighborhoods. Two pieces of evidence presented earlier weigh against this interpretation: students routinely rank schools well beyond their immediate neighborhoods, and allowing a separate preference for the assigned neighborhood school yields estimates similar to the benchmark model.

As an additional check, we estimate a specification that includes both public transit time and roadway distance between a student's residential census block group and each school. The estimated penalty associated with one hour of public transit travel remains large at 2.81 utils, while the estimated penalty associated with roadway distance is 0.16 utils per mile. Thus, among student–school pairs separated by the same roadway distance, schools requiring longer public transit trips face substantially lower predicted demand. This provides additional evidence that the structure of the public transit network is an important friction in school choice. Appendix A.2 provides details.

6.3 Assessing Residential Sorting

As discussed in subsection 5.4, the main threat to interpreting the travel-time estimate is residential sorting with respect to school access. If families with stronger unobserved preferences for school quality choose residences with better transit access to highly-valued schools, then travel times would be correlated with unobserved demand for quality. The model could then attribute to travel costs what is partly sorting on preferences, potentially biasing our estimates of (λ, δ_j). This is a natural concern because households may choose where to live in part to gain access to preferred schools. Recent work shows that residential location and school choice are jointly determined and that access to schools can be an important determinant of neighborhood choice [Park and Hahm, 2022]. At the same time, research on low-income families shows that housing search and other mobility frictions

can substantially limit moves to more desirable neighborhoods, even when families would prefer to make such moves [Katz, 2026]. The concern for our setting is more specific: whether the families in Chicago that we study, particularly those with stronger unobserved demand for high-quality schools, move to neighborhoods with shorter travel times to those schools before submitting their GoCPS rankings.

We provide two pieces of evidence on this threat. First, in a descriptive analysis of pre-application moving behavior, we examine whether families who move in the years before high school application systematically relocate to neighborhoods with faster access to top-rated or highly-valued schools. Second, in a model-based test, we examine whether the estimated travel-time disutility is sensitive to potential sorting bias.

6.3.1 Moving Behavior in the Years Before High School

We begin with the address-history evidence. Families move for many reasons unrelated to school choice, so the relevant question is not whether they move, but whether they move in ways that would bias our parameter estimates. In particular, we ask whether families systematically move toward neighborhoods with faster access to schools that families with stronger unobserved demand for quality would be more likely to value. We test whether moves in the three years before application reduce travel time to top-rated or highly-valued schools.

We compare each applicant's residential location in fall 2017, the application year, to their location three years earlier, in fall 2014.[14] We geocode both addresses to census block groups and compute public transit travel times from each block group to every school. For each student, we record the transit time to the nearest Level 1+ school from the student's 2017 and 2014 block groups. If families systematically relocate toward top-rated schools, the 2017 travel time should be shorter than the 2014 travel time.

Figure 7 plots the change in travel time to the nearest Level 1+ school, measured as 2017 minus 2014 travel time. Negative values indicate movement toward faster public transit access. The

[14] CPS administrative records allow us to match 20,813 fall 2017 applicants to addresses in fall 2014, which represents 91% of our sample.

distribution is tightly centered at zero, with a mean change of +0.22 minutes, or about 13 seconds. Thus, there is no detectable tendency for applicants to move toward better access to a Level 1+ school in the pre-application period. This conclusion is not driven by students who did not move. Appendix Figure A.5 restricts the sample to applicants who moved to a different census block group between 2014 and 2017. The distribution is again centered near zero, with a mean change of +0.54 minutes, or about 32 seconds.

Next, we broaden the metric of access to highly-valued schools. Access to the nearest Level 1+ school is a transparent measure because it uses the district's administrative rating system and is independent of our estimated demand model. However, families may value access not only to the nearest top-rated school, but also to a broader set of nearby schools that differ in estimated value. We therefore construct a second measure that combines travel time with the model-implied value of nearby schools.

For each census block group, we identify the five schools with the shortest public transit times and denote this set $\mathcal{J}_i^{(5)}$. We then convert the estimated mean utilities, $\hat{\delta}_j$, of these schools to percentiles, q_j^{pct}, and summarize the value-adjusted accessibility of this set using the Minutes-per-Percentile metric:

$$\text{MPP}_i(5) \;=\; \frac{\sum_{j \in \mathcal{J}_i^{(5)}} \text{transit_time}_{ij}}{\sum_{j \in \mathcal{J}_i^{(5)}} q_j^{\text{pct}}}\,.$$

The unit of MPP is minutes per percentile point; lower values indicate greater value-adjusted accessibility among the five nearest schools. For movers, a decrease in $\text{MPP}_i(5)$ means the new neighborhood offers better access to nearby highly-valued schools.

Figure 8 shows the distribution of the *change* in $\text{MPP}_i(5)$, i.e., $\text{MPP}_i^{2017}(5) - \text{MPP}_i^{2014}(5)$. Because $\text{MPP}_i(5)$ is an index of value-adjusted access, the key prediction is directional: moves toward better access to highly-valued schools would reduce the index. In the data, the mean change is instead positive and close to zero, at +0.15 minutes per percentile point, and the distribution is centered near zero. Thus, even when we broaden the access measure to a portfolio of nearby high-$\hat{\delta}_j$ schools, we find no evidence that families systematically move to neighborhoods with better value-adjusted

school access.

Both descriptive measures provide little evidence that pre-application moves systematically improve access to top-rated or highly-valued schools. These patterns do not rule out all residential sorting. Families may have sorted before 2014, may sort on dimensions not captured by these measures, or may make moves whose effects differ across subgroups.[15] However, the address-history evidence weighs against the specific form of pre-application sorting most likely to bias our estimates: families with stronger unobserved preferences for school quality moving to neighborhoods with shorter travel times to the schools they value more.

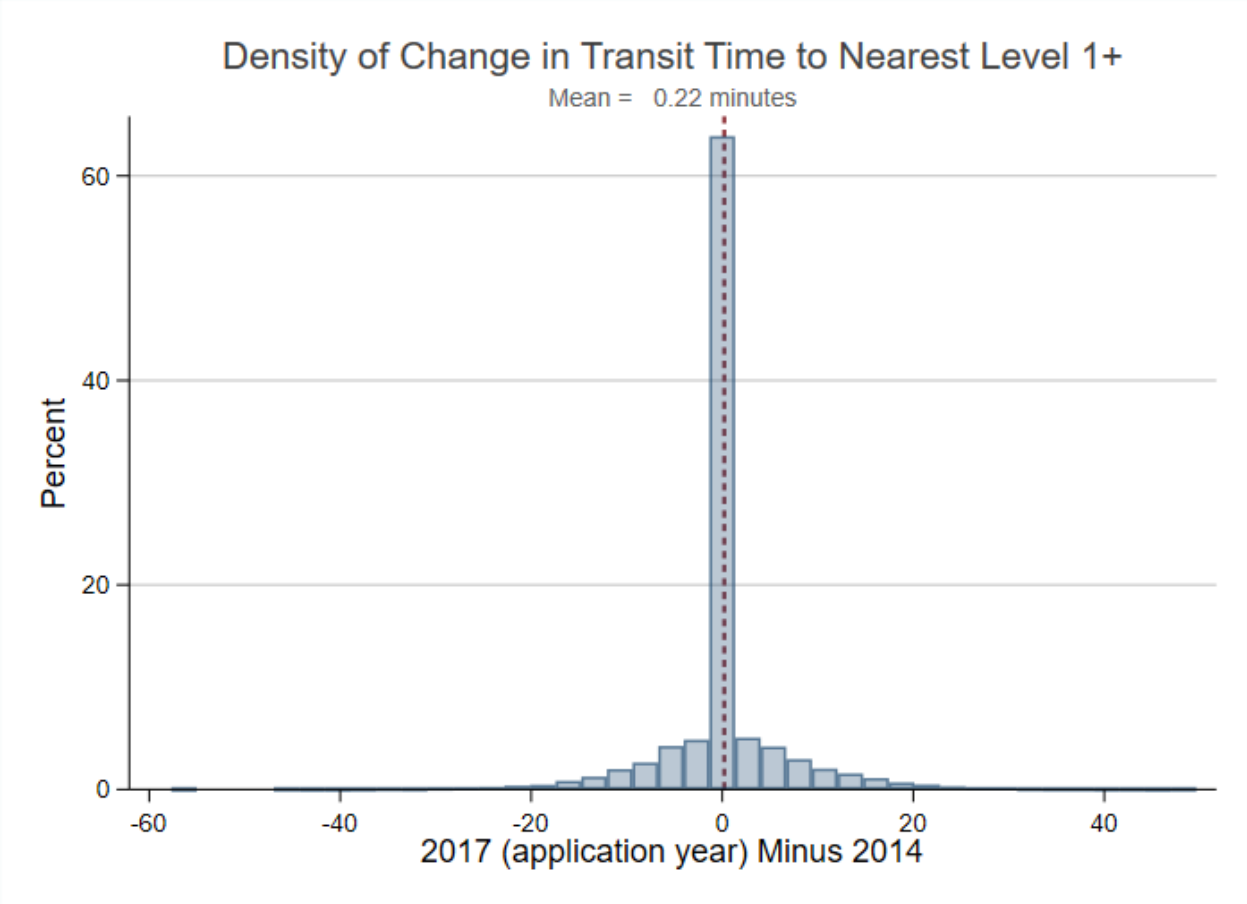
6.3.2 Allowing for Correlation in Taste for School Quality and Travel-Time Disutility

As a complementary model-based check, we ask whether the travel time estimate changes when we add an observed measure of school quality to the model and allow tastes for that measure to be correlated with travel time disutility. To motivate this specification, recall that the benchmark model includes school mean utilities, δ_j, and therefore does not require us to specify which school characteristics families value. This flexibility is useful, but it also means that demand for school quality is absorbed into the estimated school mean utilities. If residential sorting is partly driven by families with stronger demand for school quality locating in neighborhoods with better access to high-quality schools, then the travel time estimate could change once we explicitly allow for heterogeneity in tastes for quality. We therefore estimate a mixed-logit specification that adds a student-specific taste for school quality, measured by the school's SQRP rating, and allows that taste to be correlated with the student-specific disutility of travel time:

$$U_{ij} = \delta_j + \beta_i SQRP_j - \lambda_i d_{ij} + \epsilon_{ij} \quad where \begin{pmatrix} \beta_i \\ \lambda_i \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_\beta \\ \mu_\lambda \end{pmatrix}, \begin{pmatrix} \sigma_\beta^2 & \sigma_{\beta\lambda} \\ \sigma_{\beta\lambda} & \sigma_\lambda^2 \end{pmatrix} \right)$$

[15] An unconditional average change near zero could theoretically mask offsetting mobility across subgroups. However, substantial sorting toward better school access by some families would require comparably large declines in access among other movers. Such declines could reflect either purposeful sorting or displacement arising from housing market adjustments. The former is unlikely because access to highly-valued schools is a weakly positive amenity, while the latter would have to be sufficiently large and systematic to offset the gains from access-seeking moves. The near-zero mean therefore suggests that access-optimizing relocation is not pervasive during this three-year window.

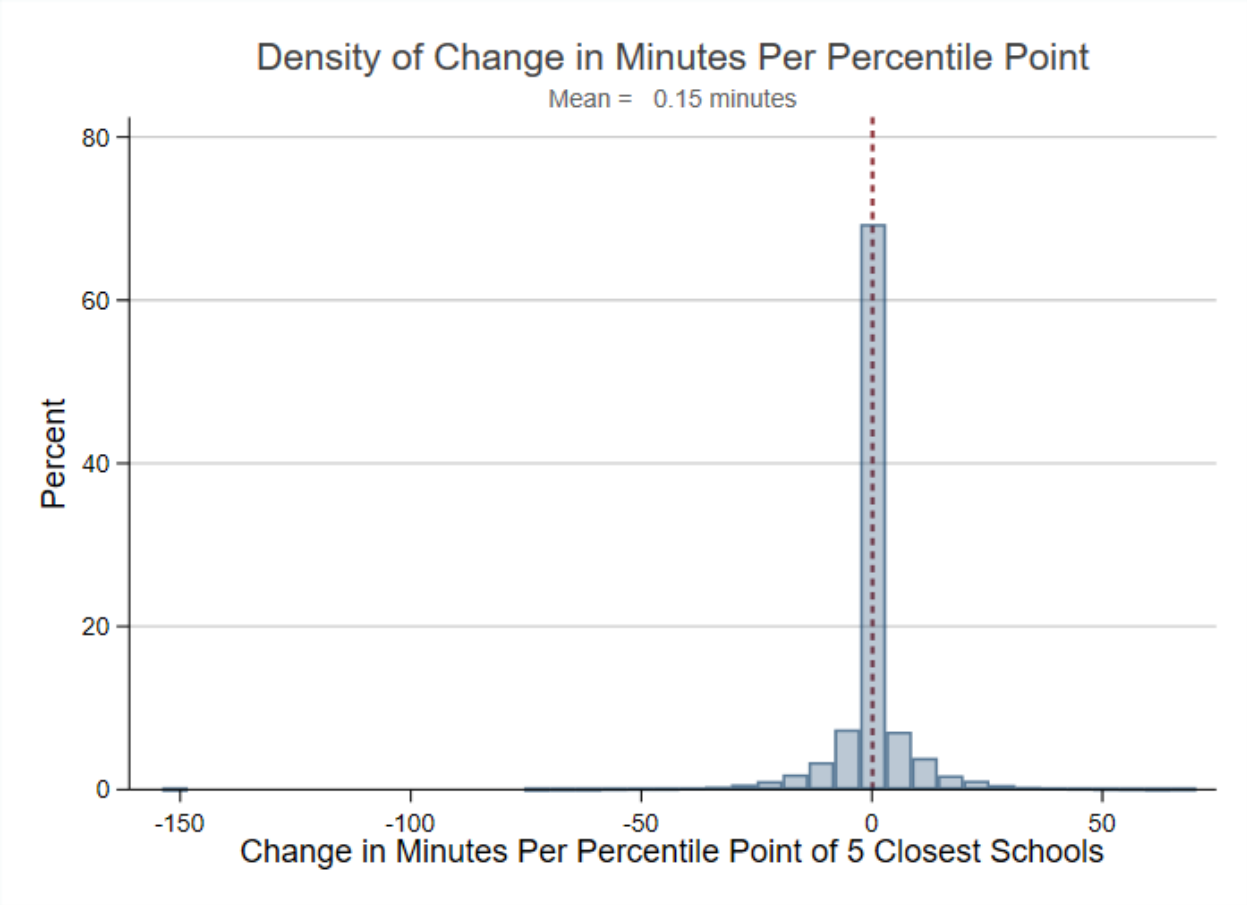
Figure 7: Pre-Application Moves and Transit Time to the Nearest Level 1+ School



This figure depicts the distribution of changes in transit time to the nearest Level 1+ school between a student's residential location in the year of high school application, 2017, and the student's residential location three years prior, in 2014. A negative value indicates having shorter public transit travel time to a Level 1+ school in 2017 than in 2014.

The estimates are nearly identical across the two specifications. In the benchmark random-coefficients model, the estimated mean disutility of travel time is $\hat{\mu}_\lambda = 4.71$, with a standard error of 0.11. When we add a student-specific taste for SQRP and allow that taste to be correlated with travel-time disutility, the estimate is $\hat{\mu}_\lambda = 4.74$, with a standard error of 0.12.

Figure 8: Pre-Application Moves and Access to Nearby High-Value Schools



This figure depicts the distribution of changes in a residential neighborhood's $\text{MPP}_i(5)$ between the year of high school application, 2017, and the student's residential location three years prior, in 2014. $\text{MPP}_i(5) = \frac{\sum_{j \in \mathcal{J}_i^{(5)}} \text{transit_time}_{ij}}{\sum_{j \in \mathcal{J}_i^{(5)}} q_j^{\text{pct}}}$. A negative value indicates moving to a neighborhood with better value-adjusted access to nearby high-$\hat{\delta}_j$ schools.

6.3.3 Summary of Tests for Sorting

Our evidence in subsection 6.3 does not rule out all residential sorting. Families may, for example, make residential choices earlier than the three-year period we observe. However, if sorting toward better school access were a widespread and ongoing feature of families' residential decisions, we would expect at least some evidence of it among moves in the three years before high school application. We find none: across both descriptive measures, the distributions are centered near zero, and the mean changes are slightly in the opposite direction. Moreover, allowing preferences for observed school quality to be correlated with travel-time disutility leaves the estimated travel-time parameter nearly unchanged. A sorting mechanism that materially biases our estimate of λ would therefore have to be concentrated before our address-history window, operate through dimensions of school preferences not captured by observed school quality, or take some other form that is consistent with both of these findings. We cannot rule out such mechanisms, but the evidence provides little support for residential sorting as an important source of bias in our estimate of λ.

7 Geography, Transportation Networks, and Racial Gaps in Access to Top-Rated Schools

We now use the estimated demand model to quantify how geography contributes to racial differences in demand for top-rated schools. These exercises are demand-side thought experiments, not assignment simulations: they ask how predicted first choices change when we alter the travel-time environment or the spatial placement of school value, holding preference parameters fixed. We focus on Black students because they are least likely to rank a Level 1+ school first and face the longest transit times to top-rated schools.

The counterfactuals isolate two related geographic channels. The first changes the travel-time environment Black students face by assigning them the residential locations and corresponding travel-time vectors of white students. The second keeps student home locations and travel times fixed but randomly reassigns estimated school mean utilities across school locations. Comparing the two exercises helps distinguish whether racial differences in demand for top-rated schools are

driven primarily by where Black students live in the transit network or by where highly-valued schools are located.

7.1 A Transit-Time Swap

The first thought experiment asks how Black students' first choices would change if they faced the travel-time environment of white students. This counterfactual is similar in spirit to Laverde [2026], who gives Black families the same distribution of distances as white families when studying the impact of assignment rules in Boston. In our setting, the relevant transportation friction is public transit travel time rather than distance.

For each Black student, we randomly draw a white student and assign the Black student the white student's residential location and corresponding vector of public transit travel times to each school. We then recompute the Black student's predicted first-choice probabilities using the Black-student preference parameters. The student's school valuations and mean travel-time disutility therefore remain those estimated for Black students, while the travel times are those faced from the randomly selected white student's location.[16]

The effect is large. Under the baseline travel times, the model predicts that 45 percent of Black students prefer a Level 1+ school. When Black students are assigned the residential locations and travel-time vectors of white students, the predicted share rises to 68 percent. Thus, holding Black students' estimated school valuations fixed, changing only the travel-time environment increases predicted demand for Level 1+ schools by 23 percentage points.[17]

[16] For each Black student i, we compute predicted first-choice probabilities using

$$\hat{P}_{ij} = \frac{\exp(V_{ij})}{\sum_{s=1}^{J} \exp(V_{is})} \qquad \text{where} \quad V_{ij} = \hat{\delta}_{j,Black} - \hat{\mu}_{\lambda,Black} d_{ij}^{white}$$

where d_{ij}^{white} is the public transit travel time from the residential location of a randomly selected white student to school j, and $\hat{\delta}_{j,Black}$ and $\hat{\mu}_{\lambda,Black}$ are estimated from the Black-student model with random parameters on the disutility of travel time. For school-race cells for which we cannot estimate a precise Black mean utility because too few Black students ranked the school, we assign a large negative utility so that the predicted first-choice probability is effectively zero.

[17] In all counterfactuals in this section, we exclude several military high schools and a pre-professional arts high school whose specialized missions make them less comparable to the general high school choice set. As a result, the baseline predicted share of Black students preferring a Level 1+ school in these exercises is 45 percent, rather than the 38 percent reported in the raw descriptive statistics. We therefore focus on changes relative to this 45 percent counterfactual baseline.

This result indicates that public transit geography is quantitatively important for Black students' lower demand for top-rated schools. The counterfactual holds Black students' estimated school valuations and the set and quality of schools fixed, changing only the travel times they face. The large increase therefore shows that where students live in the transit network substantially shapes which schools they are willing to rank first.

7.2 Randomly Reassigning School Mean Utilities

The second thought experiment asks whether the spatial placement of highly-valued schools also matters. We keep student home locations, school locations, and public transit travel times fixed, but randomly reassign the estimated school mean utilities across existing school locations. A school's Level 1+ status moves with its mean utility, so the counterfactual changes the spatial placement of school value and top-rated status together. We then recompute predicted first-choice probabilities for Black students using the same mean travel-time disutility as in the Black-student model. We repeat the random reassignment 1,000 times and average the results across simulations.[18]

Randomly reassigning school mean utilities increases predicted demand for Level 1+ schools, but by less than the transit-time swap. The predicted share of Black students preferring a Level 1+ school rises from 45 percent to 53 percent. This 8 percentage point increase shows that the spatial placement of highly-valued schools matters: if these schools were located differently across the city, Black students' predicted demand for top-rated schools would increase.

The smaller effect is also informative. Randomly reassigning school mean utilities changes where school value is located but leaves Black students in the same neighborhoods facing the same public transit network. The comparison therefore suggests that the location of highly-valued schools contributes to the racial gap in demand for top-rated schools, but differences in travel times through

[18] In each simulation, we randomly assign estimated Black-student school mean utilities to existing school locations and compute

$$\hat{P}_{ij} = \frac{\exp(V_{ij})}{\sum_{s=1}^{J} \exp(V_{is})} \quad \text{where} \quad V_{ij} = \hat{\delta}_{j,Black}^{random} - \hat{\mu}_{\lambda,Black} d_{ij}$$

where $\hat{\delta}_{j,Black}^{random}$ is the Black-student mean utility randomly assigned to the physical location of school j, and d_{ij} is the public transit travel time from student i's residential location to that school location. The Level 1+ indicator is reassigned together with the school's mean utility, so predicted demand for Level 1+ schools is computed using the reassigned Level 1+ status in each simulation.

the transit network are the larger channel in these simulations.

Table 3: Counterfactual Demand for Level 1+ Schools Among Black Students

Scenario	Predicted share preferring Level 1+
Baseline counterfactual sample	45%
White-student residential locations and travel times	68%
Randomly reassigned school mean utilities	53%

Notes: This table reports predicted first-choice demand for Level 1+ schools among Black students under the baseline and two counterfactuals. The baseline differs from the raw descriptive share because the counterfactual sample excludes several military high schools and a pre-professional arts high school with specialized missions. The transit-time swap assigns Black students the residential locations, and corresponding travel-time vectors, of randomly selected white students while holding Black-student preference parameters fixed. The random-utility reassignment keeps student and school locations fixed but randomly reassigns estimated Black-student school mean utilities, together with Level 1+ status, across school locations.

Table 3 summarizes the two counterfactuals. Black students' predicted demand for Level 1+ schools rises by 23 percentage points under the transit-time swap, compared with 8 percentage points when school mean utilities are randomly reassigned across locations. Thus, geography matters through both the spatial distribution of school value and the travel times created by the transit network, but the latter is the larger channel in these simulations. This distinction motivates the targeted transportation counterfactuals below: reducing travel times from specific neighborhoods to highly-valued schools may directly increase which schools families are willing to rank first.

8 Transportation Policy and Targeted-School Access

The previous section shows that travel times are a central mechanism behind racial gaps in demand for top-rated schools. We now ask whether targeted transportation policies can increase demand when aimed at neighborhoods where public transit access most sharply constrains which schools students rank first. Because public transit access is not a smooth function of distance, some neighborhood-school pairs involve especially long or indirect commutes. A transportation policy therefore need not operate citywide to matter: a direct connection from a particular neighborhood to a particular school may substantially change which schools students are willing to rank first.

The counterfactuals in this section use the demand model to identify where reductions in travel

time would most change students' rankings. This focus on demand is deliberate: under deferred acceptance, students must rank a school before they can be assigned to it.

We approximate a targeted express bus by replacing the public transit travel time between a neighborhood and a target school with the driving time between the same locations. For student i in targeted neighborhood S and target school j, the predicted change in the probability that the student ranks school j first is

$$\Delta_{ij} \equiv \hat{P}_{ij}\left(d_{ij}^{drive}, \hat{\boldsymbol{\delta}}_{Black}, \hat{\mu}_{\lambda,Black}\right) - \hat{P}_{ij}\left(d_{ij}^{transit}, \hat{\boldsymbol{\delta}}_{Black}, \hat{\mu}_{\lambda,Black}\right).$$

We then aggregate this change across all students in the targeted neighborhood. The predicted increase in demand for school j from students in neighborhood S is

$$\Delta_j = \sum_{i \in S} \Delta_{ij}.$$

Thus, Δ_j is the predicted number of additional students from the targeted neighborhood who would rank school j first if an express bus reduced travel time to that school from public transit time to driving time. We calculate this separately for each school, as if only one express route were created at a time. The exercise asks whether reducing travel time along a specific neighborhood-school connection would make the target school more likely to be ranked first.

The large gaps between public transit and driving times suggest substantial scope for this type of intervention to increase demand. Among Black students in our sample, the median difference between public transit and driving time to the nearest Level 1+ school is 14 minutes. For one-quarter of Black students, public transit takes at least 19 minutes longer than driving, and for one-tenth it takes at least 24 minutes longer. Thus, direct transportation improvements could substantially reduce travel times to top-rated schools for many Black students. Appendix Figure A.6 reports the full distribution by race.

We apply this exercise to a predominantly Black neighborhood on the South Side. The neighborhood has 348 applicants in the 2017 cohort and is 95 percent Black. Under the status quo, only 27 percent of applicants in this neighborhood rank a Level 1+ school first, compared with 56 percent

across all neighborhoods. The nearest Level 1+ school is 35 minutes away by public transit.[19] Because the exercise is designed to study demand among Black students, we use the Black-student model parameters estimated above.

Figure 9 reports the ten schools with the largest predicted increases in first-choice demand from an individually targeted express bus route. The shaded area is the targeted neighborhood, and the numbered schools are ordered by Δ_j. The table reports the corresponding increase in first-choice demand as a share of applicants in the neighborhood.

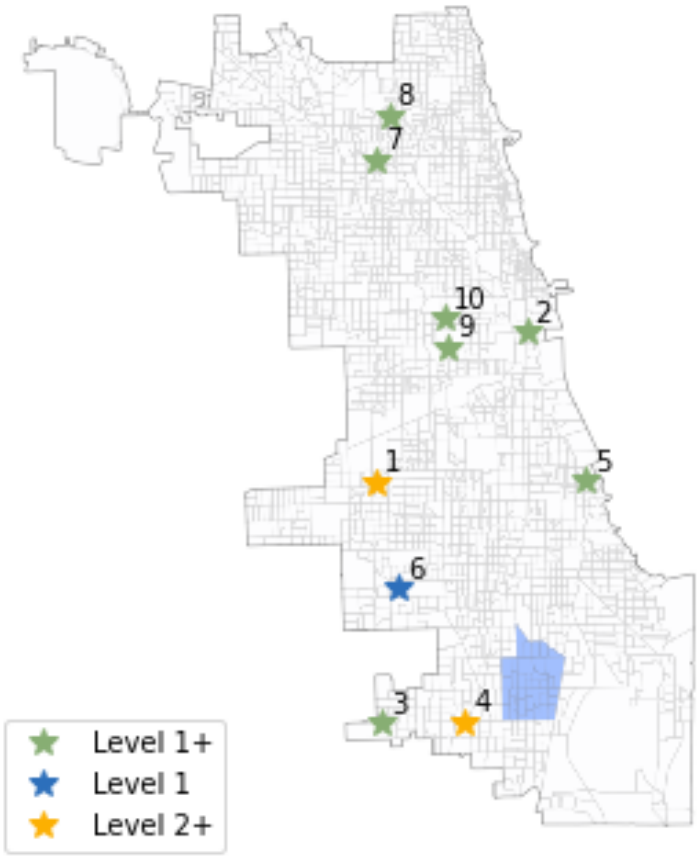
The largest predicted responses are sizable. A route to the highest-impact school would increase the share of students ranking that school first by 14.8 percentage points. Routes to the next two schools would increase first-choice demand by 14.7 and 13.7 percentage points. Each of the ten largest individually targeted routes increases predicted first-choice demand for the target school by at least 8.7 percentage points.

The rankings in Figure 9 also illustrate how the model combines transportation and school demand. Reducing travel time to any school increases the probability that students rank that school first, but the magnitude depends on both the travel-time reduction and the school's estimated mean utility. The largest responses therefore identify neighborhood-school pairs where transportation constrains demand for schools that students otherwise value. Seven of the ten largest predicted responses are to Level 1+ schools. The remaining schools have lower accountability ratings but are sufficiently valued, or experience sufficiently large reductions in travel time, to generate sizable changes in first-choice demand.

This example evaluates one express bus route at a time, but the same demand system can evaluate a network of targeted routes. The combined demand response need not equal the sum of the single-route effects because routes may substitute for one another when they make several attractive schools easier to reach, or complement one another by improving access from different neighborhoods to different schools. Because the model predicts first-choice probabilities under any vector of

[19] As an external descriptive check, we compare the targeted neighborhood with tract-level measures of economic mobility from the Opportunity Atlas (Chetty et al., 2026). The measure is the predicted mean adult household-income percentile rank of children who grew up in a tract and whose parents were at the 25th percentile of the national income distribution. All 12 tracts in this targeted neighborhood fall in the bottom quartile of Chicago's distribution: the highest predicted rank among them is 33.0, slightly below the citywide 25th percentile of 33.1.

Figure 9: Targeted Express Bus Counterfactual from a Predominantly Black Neighborhood



Rank	SQRP Rating	Switch share
1	Level 2+	14.8%
2	Level 1+	14.7%
3	Level 1+	13.7%
4	Level 2+	11.6%
5	Level 1+	10.9%
6	Level 1	10.5%
7	Level 1+	9.6%
8	Level 1+	9.5%
9	Level 1+	9.4%
10	Level 1+	8.7%

Notes: The shaded region is the targeted neighborhood. The map labels the ten schools with the largest predicted increases in first-choice demand from an individually targeted express bus route. Each route is evaluated separately by replacing public transit time from the targeted neighborhood to the labeled school with driving time. The table reports the corresponding predicted increase in first-choice demand as a share of applicants in the targeted neighborhood. Calculations use Black-student model parameters.

travel times, these interactions can be evaluated by reducing multiple neighborhood-school travel times simultaneously.

These results characterize demand responses rather than final assignments or the cost-minimizing transportation network. In neighborhoods where public transit makes top-rated schools difficult to reach, targeted reductions in travel time can substantially increase the probability that students rank those schools first. This provides one concrete mechanism through which hyper-local transportation policies could expand access to top-rated schools by improving the neighborhood-school connections where transit access is most binding.

8.1 Evaluating Human Capital Targets

The express bus counterfactuals that we consider above target schools with high public accountability ratings – those designated as 'Level 1+'. But our policy design tool enables us to consider a wide variety of schools that could be selected by education policymakers. In this section, we also estimate the impacts of targeting schools based on a human capital criterion – value-added. There is good reason to consider other schools as targets. Traditional accountability ratings may partly reflect the characteristics of enrolled students rather than the schools' causal contributions to learning, and therefore expanding access to highly-rated schools need not be equivalent to expanding access to schools that generate the largest achievement gains (Angrist, Hull, Pathak, and Walters [2024]).

We find that if shifting student demand towards higher value-added schools is an important policy objective, then targeted transportation is a potentially powerful policy lever to change demand. To see this, we first show that family preferences as captured by school mean utilities are only weakly correlated with value-added and more strongly correlated with accountability ratings. For Black students, estimated mean utility ($\widehat{\delta}_{j,\text{Black}}$) is strongly correlated with SQRP points, with $\rho = 0.48$, but much less correlated with estimated value-added: $\rho = 0.22$ for ELA and $\rho = 0.08$ for mathematics.[20] Thus, the schools that students appear to value most are not necessarily those with the largest estimated contributions to achievement growth, consistent with evidence in Abdulkadiroğlu, Angrist, and Pathak [2014] and Abdulkadiroglu, Pathak, Schellenberg, and Walters [2020].

Reducing transportation time to a high value-added school can increase demand even when that school's contribution to achievement is not fully reflected in family preferences. This stems from the fact that reducing the travel time to *any* school will increase the utility of the school and thus increase the probability that a student ranks that school first. More formally, when utility is specified as $U_{ij} = \delta_{jk} - \lambda_k d_{ij} + \epsilon_{ij}$, then reducing d_{ij} will increase the probability that $U_{ij} \geq U_{i\ell}\ \forall \ell \neq j, \ell \in J_i$, which will increase the demand for school j.

[20] We describe the construction of the value-added measures in Appendix A.3.

We find that targeted routes to high math value-added schools will substantially increase demand for those schools even when math value-added is only weakly related to family preferences. We continue to focus on the same South Side origin neighborhood but redirect the targeted routes toward schools in the top quintile of mathematics value-added.[21] For each targeted school, we estimate the increase in predicted demand among the 348 applicants when travel time is reduced from public transit time to driving time. Several schools in the top quintile of mathematics value-added experience sizable increases in predicted demand of 14.7, 10.9, 9.4, 8.7, 7.1, and 6.8 percentage points. Although these responses are generally smaller than the responses in the preceding exercise, they show that targeted reductions in travel time can meaningfully increase demand for schools selected on the basis of estimated contributions to achievement growth.

These results separate two policy choices: which schools the district wishes to make more accessible and how transportation can be redesigned to reduce the cost of reaching them. The initial exercise identifies transportation links to schools that families appear to value, whereas this exercise targets schools with high estimated contributions to achievement. In both cases, the model identifies highly local transportation changes that increase predicted demand for the targeted schools. More broadly, the framework provides a toolkit for identifying transportation interventions that can increase demand for policymaker-defined target schools from neighborhoods where transportation frictions are most binding.

9 Conclusions

School choice policies seek to weaken the link between residential location and educational opportunity. This paper shows that the link can persist due to the transportation network that creates a highly local geography of access. Using ranked-choice data and public transit times from students' homes to each school, we estimate a demand model that separately identifies school valuations from the disutility of travel time. We find that different demographic groups broadly agree on the ranking of schools, but large differences in travel time across neighborhoods yield substantial

[21] The target set includes multiple campuses in the Noble Charter Network. Prior lottery-based evidence finds substantial achievement gains from attending Noble schools, providing some external validation that their high value-added estimates do not simply reflect the characteristics of enrolled students [Davis and Heller, 2019, Cohodes et al., 2021].

differences in demand for top-rated schools. Counterfactuals show that travel time differences are an important determinant of the racial gap in preferences for top-rated schools: assigning Black students the travel times faced by white students increases predicted demand for top-rated schools by 23 percentage points. Well-targeted interventions such as express bus routes generate sizable increases in demand when they connect neighborhoods with poor transit access to schools that students value and to schools with large estimated contributions to achievement growth. This paper provides a toolkit to identify neighborhood-school connections where transportation improvements can increase demand for targeted schools.

This framework could be used to study education policy beyond the transportation counterfactuals in this paper. More broadly, it allows policymakers to quantify how access depends on the joint locations of programs and students. The same framework could be used to evaluate the distributional effects of siting new magnet or specialty programs, targeted capacity expansions, or school closures or consolidations. In each case, combining estimates of school demand with transportation access can help identify how policy changes affect different neighborhoods and student populations.

References

A. Abdulkadiroglu and T. Sonmez. School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3):729–747, 2003.

A. Abdulkadiroğlu, J. Angrist, and P. Pathak. The Elite Illusion: Achievement Effects at Boston and New York Exam Schools. *Econometrica*, 82(1):137–196, 2014.

A. Abdulkadiroğlu, N. Agarwal, and P. A. Pathak. The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match. *American Economic Review*, 107(12): 3635–3689, 2017.

A. Abdulkadiroğlu, P. A. Pathak, and C. R. Walters. Free to Choose: Can School Choice Reduce Student Achievement? *American Economic Journal: Applied Economics*, 10(1):175–206, 2018.

A. Abdulkadiroglu, P. A. Pathak, J. Schellenberg, and C. R. Walters. Do Parents Value School Effectiveness? *American Economic Review*, 110(5):1502–1539, 2020.

N. Agarwal and P. Somaini. Revealed Preference Analysis of School Choice Models. *Annual Review of Economics*, 12(1):471–501, 2020.

P. Agte, C. Allende, A. Kapor, C. Neilson, and C. Ochoa. School Choice Frictions: Information, Awareness, Search, and Beliefs. working paper, 2026.

R. Ainsworth, R. Dehejia, C. Pop-Eleches, and M. Urquiola. Why Do Households Leave School Value Added on the Table? The Roles of Information and Preferences. *American Economic Review*, 113(4):1049–1082, 2023.

C. Allende. Competition Under Social Interactions and the Design of Education Policies. working paper, 2021.

C. Allende, F. Gallego, and C. Neilson. Approximating the Equilibrium Effects of Informed School Choice. working paper, 2026.

J. Angrist, P. Hull, P. A. Pathak, and C. R. Walters. Race and the Mismeasure of School Quality. *American Economic Review: Insights*, 6(1):20–37, 2024.

F. Arteaga, A. J. Kapor, C. A. Neilson, and S. D. Zimmerman. Smart Matching Platforms and Heterogeneous Beliefs in Centralized School Choice. *The Quarterly Journal of Economics*, 137 (3):1791–1848, 2022.

L. Barrow, L. Sartain, and M. de la Torre. Increasing Access to Selective High Schools through Place-Based Affirmative Action: Unintended Consequences. *American Economic Journal: Applied Economics*, 12(4):135–163, 2020.

L. Barseghyan, D. Clark, and S. Coate. Peer Preferences, School Competition, and the Effects of Public School Choice. *American Economic Journal: Economic Policy*, 11(4):124–158, 2019.

D. Beuermann, C. K. Jackson, L. Navarro-Sola, and F. Pardo. What is a Good School, and Can Parents Tell? Evidence on the Multidimensionality of School Output. *Review of Economic Studies*, 90:65–101, 2023.

C. Calsamiglia, G. Haeringer, and F. Klijn. Constrained School Choice: An Experimental Study. *American Economic Review*, 100(4):1860–1874, 2010.

C. Calsamiglia, C. Fu, and M. Güell. Structural Estimation of a Model of School Choices: The Boston Mechanism versus Its Alternatives. *Journal of Political Economy*, 128(2):642–680, 2020.

C. Campos and C. Kearns. The Impact of Neighborhood School Choice: Evidence from Los Angeles' Zones of Choice. working paper, 2020.

R. Chetty, J. Friedman, N. Hendren, M. R. Jones, and S. R. Porter. The opportunity atlas: Mapping the childhood roots of social mobility. *American Economic Review*, 116(1):1–51, 2026.

E. Chyn and L. F. Katz. Neighborhoods Matter: Assessing the Evidence for Place Effects. *Journal of Economic Perspectives*, 35(4):197–222, 2021.

S. R. Cohodes, E. M. Setren, and C. R. Walters. Can Successful Schools Replicate? Scaling Up Boston's Charter School Sector. *American Economic Journal: Economic Policy*, 13(1):138–67, February 2021.

V. Corradini. Information and Access in School Choice Systems: Evidence from New York City. working paper, 2024.

V. Corradini and C. Idoux. Overcoming Racial Gaps in School Preferences: The Effect of Peer Diversity on School Choice. working paper, 2025.

L. Costa, J. Naddeo, and M. J. Seiler. School Choice and Residential Commute Times. *Real Estate Economics*, 54:424–468, 2026.

M. Davis and B. Heller. No Excuses Charter Schools and College Enrollment: New Evidence from a High School Network in Chicago. *Education Finance and Policy*, 14(3):414–440, 2019.

N. Deutscher. Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia. *American Economic Journal: Applied Economics*, 12(2):220–249, 2020.

G. Fack, J. Grenet, and Y. He. Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review*, 109(4):1486–1529, Apr. 2019.

J. Hastings and J. Weinstein. Information, School Choice, and Academic Achievement: Evidence from Two Experiments. *The Quarterly Journal of Economics*, 123(4):1373–1414, 2008.

J. Hastings, T. Kane, and D. Staiger. Heterogeneous Preferences and the Efficacy of Public School Choice. NBER Working Paper 12145, 2008.

C. Hoxby and S. Turner. Expanding College Opportunities for High-Achieving, Low Income Students. SIEPR Discussion Paper no. 12-014, 2013.

C. Hoxby and S. Turner. What High-Achieving Low-Income Students Know About College. *American Economic Review*, 105(5):514–517, 2015.

A. Kapor, C. Neilson, and S. Zimmerman. Heterogeneous Beliefs and School Choice Mechanisms. *American Economic Review*, 110(5):1274–1315, 2020.

L. F. Katz. Presidential Address: Neighborhood Effects and Missing Markets for Opportunity. *American Economic Review*, 116(9):3197–3222, 2026.

J.-W. Laliberté. Long-Term Contextual Effects in Education: Schools and Neighborhoods. *American Economic Journal: Economic Policy*, 13(2):336–377, 2021.

M. Laverde. Distance to Schools and Equal Access in School Choice Systems. working paper, 2026.

E. Nakamura, J. Sigurdsson, and J. Steinsson. The Gift of Moving: Intergenerational Consequences of a Mobility Shock. *The Review of Economic Studies*, 89(3):1557–1592, 2022.

H. Oosterbeek, F. Sóvágó, and B. van der Klaauw. Preference Heterogeneity and School Segregation. *Journal of Public Economics*, 200:104465, 2021.

M. Park and D. W. Hahm. Location Choice, Commuting, and School Choice. working paper, 2022.

P. Sharkey. Neighborhoods, Cities, and Economic Mobility. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(2):159–177, 2016.

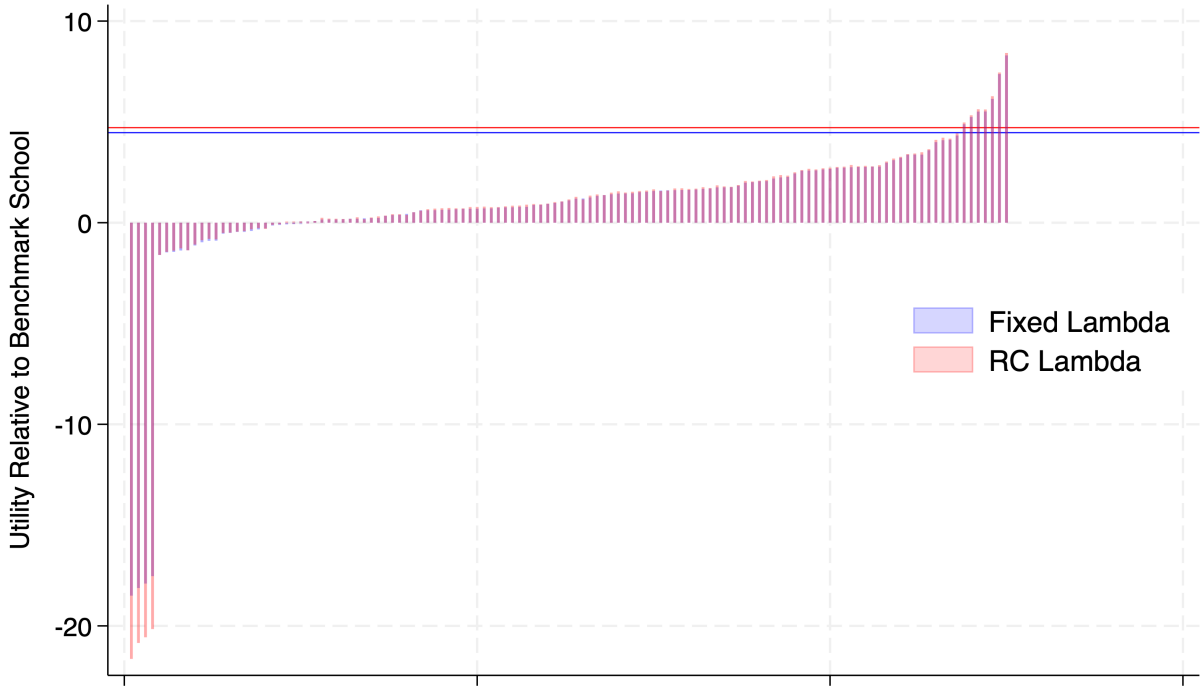
S. Trajkovski, J. Zabel, and A. E. Schwartz. Do School Buses Make School Choice Work? *Regional Science and Urban Economics*, 86:103607, 2021.

A Additional Figures and Tables

Figure A.1: Screenshot of GoCPS Website

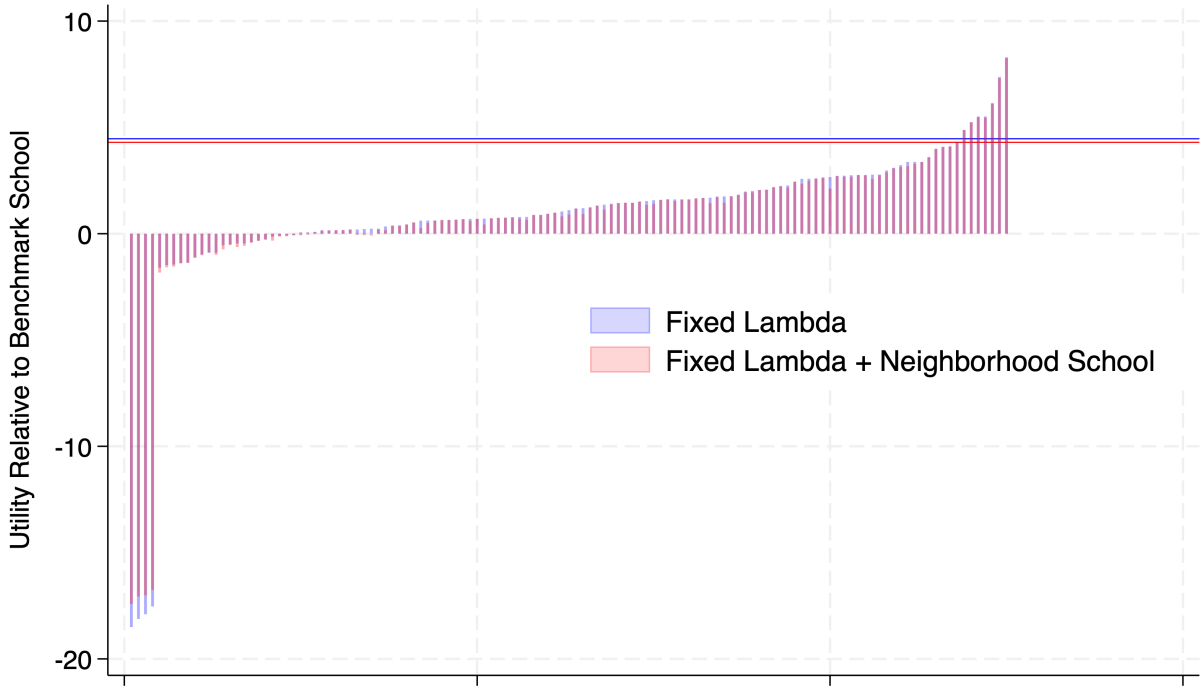


Figure A.2: Benchmark Estimates with Random Coefficients on Travel Time



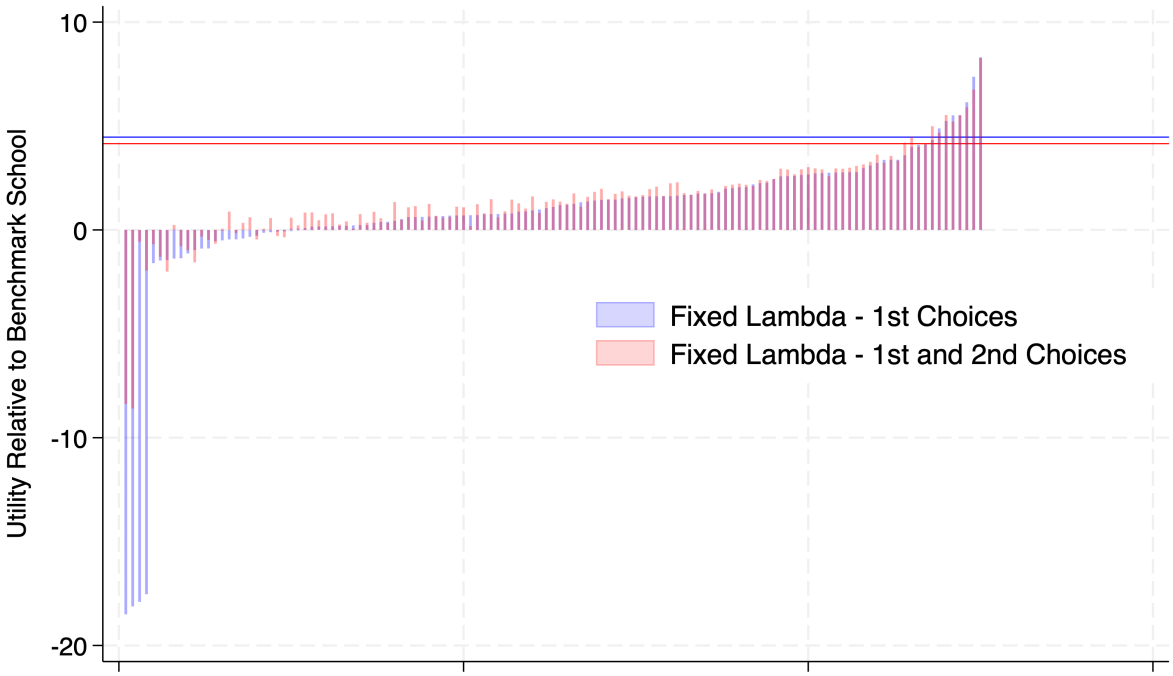
This figure compares estimates from the pooled benchmark model to estimates from a mixed-logit specification that allows the disutility of travel time to vary across students. The horizontal line represents the estimated disutility of one hour of one-way public transit travel time. The vertical bars show each school's mean utility relative to the excluded benchmark school, with schools ordered by estimated mean utility. Blue estimates assume a common fixed disutility of travel time for each student. Red estimates come from the random-coefficients model, where school mean utilities are fixed but the disutility of travel time varies across students, following a Normal distribution with mean μ_λ and standard deviation σ_λ. In the random-coefficients model, the red horizontal line depicts the estimate of μ_λ. In this benchmark model, parameters do not vary by race.

Figure A.3: Benchmark Estimates with Assigned Neighborhood-School Preference



This figure compares estimates from the pooled benchmark model to estimates from a specification that adds a separate preference for the student's assigned neighborhood school. The model includes an indicator for whether school j is student i's assigned neighborhood school: $U_{ij} = \delta_j + \beta \cdot 1(j = \text{i's zoned school}) - \lambda d_{ij} + \epsilon_{ij}$. This specification allows the assigned neighborhood school to have value beyond the school's estimated mean utility net of travel time (e.g. friends from 8th grade are more likely to attend neighborhood school). The vertical bars show estimated school mean utilities relative to the excluded benchmark school, and the horizontal line represents the estimated disutility of one hour of one-way public transit travel time.

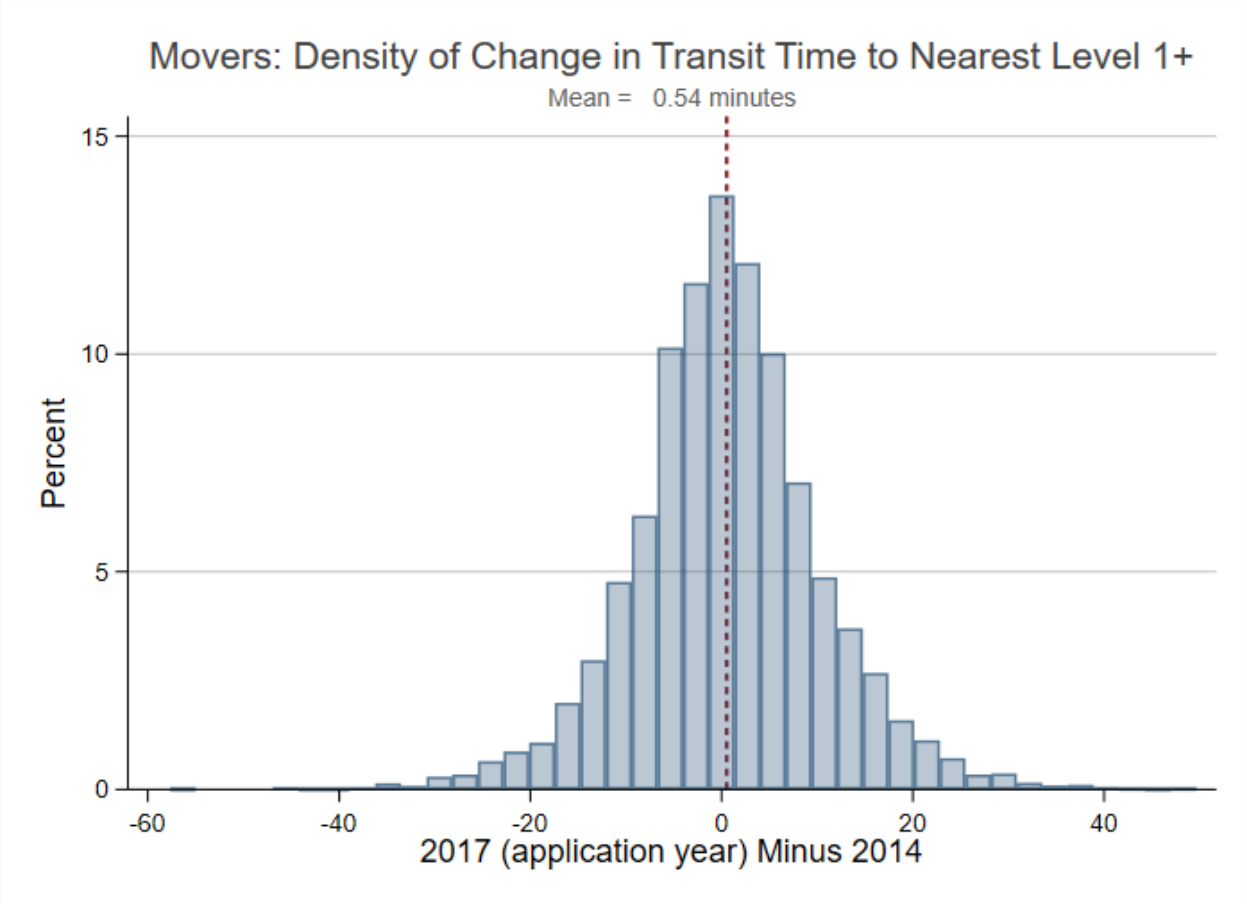
Figure A.4: Benchmark Estimates Using First and Second Choices



Excluding 4 schools with very low $\hat{\delta}_j$, the correlation = 0.98.

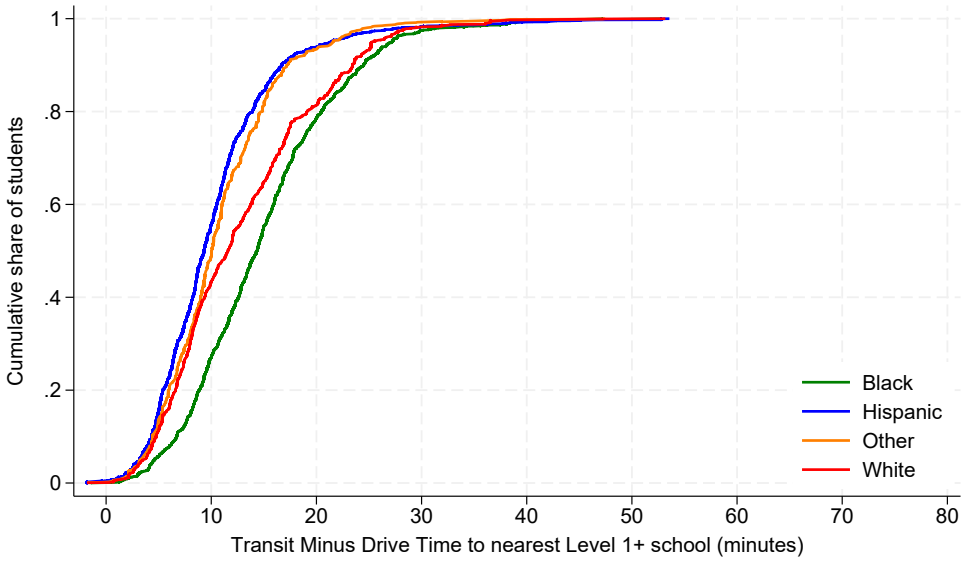
This figure compares estimates from the pooled benchmark model to estimates from a specification that uses information from both students' first-ranked and second-ranked schools in the likelihood. The benchmark model uses first choices. This robustness check assesses whether the estimated school mean utilities and travel-time disutility are sensitive to incorporating additional information from submitted rank-order lists. The vertical bars show estimated school mean utilities relative to the excluded benchmark school, and the horizontal line represents the estimated disutility of one hour of one-way public transit travel time.

Figure A.5: Movers Only: Pre-Application Moves and Transit Time to the Nearest Level 1+ School



This figure focuses on a subsample of students who change residential census block groups between 2017 and 2014. The figure depicts the distribution of changes in transit time to the nearest Level 1+ school between a student's residential location in the year of high school application, 2017, and the student's residential location three years prior, in 2014. A negative value indicates having a shorter public transit travel time to a Level 1+ school in 2017 than in 2014. 54% of applicants moved to a different block group between 2014 and 2017. Because block groups are small geographic units, this definition captures both substantial residential moves and short-distance moves within the same broader neighborhood.

Figure A.6: Difference Between Public Transit and Driving Time to the Nearest Level 1+ School



Notes: This figure plots the distribution, by race, of the difference between public transit travel time and driving time to the nearest Level 1+ school. Positive values indicate that public transit takes longer than driving.

A.1 Quadratic Travel-Time Disutility

The benchmark model assumes that the marginal disutility of travel time is constant across commute lengths. A natural concern is that longer commutes could become disproportionately costly, for example because they require earlier departures from home or impose more daily time in transit. To assess this possibility, we estimate a pooled specification that includes both travel time and travel time squared:

$$U_{ij} = \delta_j + \lambda_1 d_{ij} + \lambda_2 d_{ij}^2 + \epsilon_{ij}.$$

The estimates provide no evidence that longer commutes become disproportionately costly. Instead, the marginal disutility of travel time declines slightly as commute length increases. We estimate $\hat{\lambda}_1 = -5.66$ and $\hat{\lambda}_2 = 0.85$ in the pooled sample. These estimates imply that the marginal utility of an additional hour of travel time is -5.27 utils at the 25th percentile of the travel-time distribution to the nearest Level 1+ school and -4.68 utils at the 75th percentile. Thus, while there is some curvature, it is economically modest and goes in the opposite direction of the concern that very long commutes become sharply more costly. The benchmark linear model therefore does not appear to mask a strong nonlinear pattern over the commute lengths relevant for the counterfactuals.

A.2 Testing Whether Transit Time Proxies for Physical Distance

A potential concern with the benchmark specification is that the estimated disutility of public transit time reflects not only features of the transit network, but also a broader preference for schools located closer to home. Several features of the data and earlier specifications already make this interpretation less likely. Students routinely include schools located well beyond their immediate neighborhoods on their submitted rank lists, suggesting that limited awareness of distant schools is not the primary explanation for lower demand. In addition, allowing students to place a separate value on their assigned neighborhood school yields estimates similar to the benchmark model, indicating that $\hat{\lambda}$ is not simply capturing attachment to the local school.

We conduct an additional specification check by including both public transit travel time and roadway distance:

$$U_{ij} = \delta_j - \beta \, \text{Distance}_{ij} - \lambda d_{ij} + \epsilon_{ij},$$

where Distance_{ij} is the roadway distance, in miles, between student i's residential census block group and school j, and d_{ij} is one-way public transit travel time, measured in hours, used in the benchmark specification. The inclusion of roadway distance allows demand to vary with broad physical separation while using differences in transit time among similarly distant student–school pairs to assess the role of transit connectivity.

When both measures are included, we estimate $\widehat{\lambda} = 2.81$ utils per hour of public transit travel and $\widehat{\beta} = 0.16$ utils per mile of roadway distance. Thus, even after allowing physically more distant schools to be less attractive, longer public transit trips remain strongly associated with lower demand. In particular, among schools located a similar roadway distance from a student's home, those requiring more time on public transit are substantially less likely to be preferred.

This evidence indicates that the strong relationship between public transit time and school demand is not driven solely by unfamiliarity with distant schools, attachment to the neighborhood school, or physical distance itself. Instead, the organization of the transit network – including routes, transfers, waiting times, and network connections – constitutes an important friction in school choice.

A.3 Construction of School Value-Added Measures

We construct value-added estimates by regressing standardized ninth-grade PSAT scores on standardized eighth-grade NWEA scores and student characteristics, then averaging residuals by school.

The sample links eighth-grade NWEA Mathematics and Reading RIT scores (spring 2016) to ninth-grade PSAT Math and Evidence-Based Reading and Writing scores (2016–17) via student ID, restricted to students matched across the NWEA, PSAT, and district student-characteristics files. The 2016–17 outcome year aligns the value-added period with the SQRP school-performance period. All prior-achievement and outcome measures are standardized to mean zero, SD one within the analytic sample.

For each subject $d \in \{ELA, Math\}$, we estimate by OLS:

$$Y_{isd} = \alpha_d + \beta_d \,\text{PriorAchievement}_{id} + \gamma_d' X_i + \varepsilon_{isd} \tag{1}$$

where Y_{isd} is student i's standardized PSAT score, $PriorAchievement_{id}$ is the same-subject standardized eighth-grade NWEA score, and X_i is a vector of controls: race/ethnicity, gender, FRL status, special education status, has a Section 504 plan, in an ESL program, has limited English proficiency (LEP), homelessness status, and whether or not the student is repeating a grade.

We compute residuals $\hat{\varepsilon}_{isd} = Y_{isd} - \hat{Y}_{isd}$ and average them (unweighted) within school:

$$VA_{sd} = \frac{1}{n_s} \sum_{i \in s} \hat{\varepsilon}_{isd} \tag{2}$$

Because outcomes are standardized, VA_{sd} is in SD units of ninth-grade PSAT achievement, e.g., $VA_{sd} = 0.10$ means students at school s scored 0.10 SD above prediction on average.