

NBER WORKING PAPER SERIES

SCHOOL CHOICE AND SEGREGATION:
EVIDENCE FROM THE OAKLAND UNIFIED SCHOOL DISTRICT

Jesse Rothstein
Ini Umosen
Christopher R. Walters

Working Paper 34957
<http://www.nber.org/papers/w34957>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2026

The authors thank Joan Martinez for excellent research assistance and Oakland Unified School District (particularly Sonali Murarka, Jonathon Stewart, and Kilian Betlach) for supporting this work via data access and frequent conversations. We also thank Moonhawk Kim. This work was conducted under the auspices of the California Policy Lab. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Jesse Rothstein, Ini Umosen, and Christopher R. Walters. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

School Choice and Segregation: Evidence from the Oakland Unified School District
Jesse Rothstein, Ini Umosen, and Christopher R. Walters
NBER Working Paper No. 34957
March 2026
JEL No. I21

ABSTRACT

We study the prospects for changes in school priorities to reduce income segregation in a context of centralized school assignment, accounting for behavioral responses to school offers. Promoting integration is a central objective for large urban school districts in the US, and reforms to school assignment priorities are a prominent means of pursuing this goal. Such efforts may be constrained by students' decisions to exit the public school system in response to less-preferred school offers. Using data on kindergarten applicants to the Oakland Unified School District (OUSD), we show that offers of spots at first-choice schools boost the likelihood that applicants remain in OUSD. Nevertheless, simulations show that policy reforms giving priority for low-income students at high-income schools can substantially reduce segregation with minimal impacts on retention in the district.

Jesse Rothstein
University of California, Berkeley
Department of Economics
and NBER
rothstein@berkeley.edu

Christopher R. Walters
University of Chicago
Department of Economics
and NBER
crwalters@uchicago.edu

Ini Umosen
Ithaka S+R
ini.umosen@ithaka.org

Since the 1954 *Brown v. Board of Education* decision, a major priority of education policy in the United States has been to desegregate schools. Desegregation efforts are constrained by two important factors: Families may prefer to attend schools that are more racially isolated than desirable from the policymaker’s perspective, and extensive and ongoing residential segregation means that any assignment system based in whole or in part on proximity will tend to generate segregated school assignments. Both between-district and within-district residential segregation can be important. However, in *Milliken v. Bradley* (1974) the Supreme Court put severe constraints on the ability of policymakers to construct remedies to between-district segregation, and since then attention has focused on within-district school assignments.

Early desegregation efforts often used coercive assignment rules. For example, some districts paired a white neighborhood and a geographically distant black neighborhood, assigning all students from each to a single school that would thereby be integrated. But there was substantial parental resistance to school assignment systems that relied on busing ([Lukas, 1985](#); [Angrist et al., 2022](#)). Another type of rule, based on parental choice, has grown more common in recent years. Many large districts have adopted choice systems relying on centralized assignment algorithms to allocate students to schools according to their preferences while respecting school capacity constraints. Districts often attempt to adjust the rules of these algorithms to increase integration, by reserving some seats at each school for students from underrepresented groups or by other mechanisms for increasing the priority of applications from students who would diversify the school ([Dur et al., 2018](#)). However, such efforts may be ineffective if preferences or locations of students from different groups are sufficiently distinct ([Oosterbeek, Sóvágó and van der Klaauw, 2021](#)).

We assess the prospects for school assignment rules to reduce segregation using data on kindergarten applicants to the Oakland Unified School District (OUSD) in the period right before the COVID-19 pandemic. In OUSD, as in other urban

districts, neighborhoods are highly segregated by race and income. Moreover, the most oversubscribed schools tend to be in higher-income neighborhoods, so families' tendency to prefer nearby schools leads on its own to some degree of income segregation. This may be exacerbated by the use of neighborhood preferences, whereby students in a school's neighborhood are preferred over students outside it, in the assignment system.

We combine students' submitted preferences over schools with an implementation of the district's assignment algorithm to simulate the consequences of alternative priority rules that aim to reduce income segregation within the district.¹ This investigation employs a simple parameterization that allows us to continuously vary the preference given to low-income students who apply to over-subscribed high-income schools to which they do not otherwise have neighborhood preference. We simulate a sequence of hypothetical rules for generating school offers, starting from the status quo where all demand from neighborhood children is met before any non-neighborhood students are admitted, moving to alternative rules where low-income non-neighborhood students are admitted before non-low-income non-neighborhood students, and eventually winding up in an extreme alternative where low-income non-neighborhood students are treated as neighborhood residents. We show that feasible rules that give small but meaningful chances to low-income, non-neighborhood applicants would dramatically reduce income segregation in OUSD schools.

Importantly, our analysis accounts for a major concern that districts face in designing their school assignment rules: the potential for students to exit the district in response to less-preferred school offers. Rules that promote low-income students' access to preferred schools may reduce the share of higher-income students who are offered their most-preferred schools. This reduction imposes a cost on the district if students denied their first-choice schools respond by en-

¹We focus on income rather than racial segregation because, under current legal precedent, districts have more latitude to craft income-based school assignment policies than those based on race. Moreover, as we discuss below, only about 10% of OUSD students are white, limiting the scope for increasing integration by moving students within the district.

rolling in non-district alternatives like private or charter schools. Such attrition responses lead to losses in district funding, and may also reduce the impacts of policy changes on school integration.

We study this issue by using random lotteries embedded in the DA process to identify the causal effect of receiving a less preferred school on the probability that a student leaves the district. We show that students denied their first choices are more likely to exit OUSD. The magnitude of this effect is similar for low- and high-income students, and primarily reflects exit from the public school system rather than substitution to public charter schools. Because these attrition effects are similar for low-income and high-income students, and because our alternative offer rules only modestly reduce the overall share of students receiving their first choices, we find that student exit leads to only small changes in the simulated impacts of changes to school assignment policy. This suggests scope for changes in offer rules to reduce income segregation despite the behavioral responses induced by such changes.

Our analysis adds to a growing literature studying the causes and consequences of segregation in US public schools. One prominent strand of research assesses impacts of court-ordered desegregation and busing efforts ([Angrist and Lang, 2004](#); [Reber, 2005a](#); [Rivkin and Welch, 2006](#); [Cascio et al., 2008](#); [Billings, Deming and Rockoff, 2013](#); [Reardon and Owens, 2014](#); [Angrist et al., 2022](#)). Other work looks at the roles of segregation in the housing market ([Epple and Sieg, 1999](#); [Bayer, Ferreira and McMillan, 2007](#)), school attendance boundaries ([Richards, 2014](#); [Monarrez, 2023](#)), and school choice programs ([Bifulco and Ladd, 2007](#); [Monarrez, Kisida and Chingos, 2022](#)) in exacerbating or ameliorating school segregation. The recent trend toward centralized school assignment in large urban districts has been motivated in part by a goal of increasing integration, but the impacts of these changes on segregation are controversial (see, e.g., [Quick 2016](#); [Marquez Martinez 2018](#); [Elsen-Rooney 2024](#)). Recent studies suggest that geographic sorting, family preferences, and school admission policies all continue to

play important roles in generating race and income segregation in districts with centralized assignment (Idoux, 2023; Laverde, 2024; Corradini and Idoux, 2025; Han and Idoux, 2025; Umosen, Puller and Moreira, 2025). We complement this literature by showing large potential integration effects for a class of admission reforms focused on income diversity, holding fixed geographic sorting and preferences while accounting for behavioral responses to changes in school offers.

The remainder of the paper is organized as follows. Section I gives background on the Oakland schooling landscape and describes the OUSD school assignment process, including the district’s implementation of the deferred acceptance (DA) algorithm. Section II leverages data from the assignment process to summarize student preferences and estimate causal effects of offers to preferred schools on retention in the district. Section III simulates impacts of counterfactual priority rules that privilege students who would increase income diversity. Section IV offers concluding thoughts and directions for future research.

I. Oakland Unified School District

A. Background on schooling in Oakland

Oakland is a large district in California, across San Francisco Bay from the city of San Francisco. It serves close to 50,000 students each year, of whom about 20% are Black or African American, 50% are Hispanic/Latino, 10% are white, and 10% are Asian. The district contains both impoverished neighborhoods and wealthy neighborhoods, distinguished in part by their locations – higher income families tend to live in the hills on the eastern edge of the city, while lower-income families live in the “flats.” Importantly, OUSD has both a large number of charter schools, enrolling about 20% of students living in the district, and a developed private school sector. The district has experienced declining enrollment in recent years, due in part to the expansion of the charter sector, and has faced repeated battles over closing schools. Maintaining district enrollment is thus a top district priority. Approximately 30% of students who participate in the OUSD school

assignment lottery ultimately do not enroll at OUSD - but this share is just 25% for students who are offered their first-choice schools and 44% for those who are not.

Schools in OUSD are highly segregated. Because the white share of district enrollment is so low, we focus on segregation between low-income and high-income students. Following OUSD guidelines, we refer to neighborhoods, and their resident students, as low-income if at least one of the following is true: (1) At least half of residents are Latino or African American and median income is at or below the free/reduced priced lunch (FRPL) threshold income for a family of 5, or (2) at least 20 students, and at least 70% of students, are eligible for FRPL. By these definitions, the average student from a low-income neighborhood in the district attends a school that is 68% low income, while the average student from a high-income neighborhood attends a school with only a 31% low-income share.

OUSD has been engaged in a multi-year process to identify ways to adjust its assignment system to promote more integration without driving students out of the district. In 2021, OUSD began testing giving some non-neighborhood applicants higher priority at a few schools with low shares of low-income students. Our analysis considers the expansion of this pilot, both in terms of the number of schools affected and the number of students eligible for adjusted priorities.

Our analysis uses information on all participants in the OUSD kindergarten public school choice lottery in 2019. These data include the full list of ranked preferences that each student submitted as well as the student's priority classification (sibling, neighborhood, etc.) at each school to which he or she applied. We also observe the student's eventual school offer, whether the student enrolled at OUSD and, if not, whether he or she enrolled at one of the charter schools in the district. Finally, we observe the census block group in which each student resides.

B. School Assignment in OUSD

OUSD uses a centralized matching process to assign students to schools. Students entering kindergarten, 6th, and 9th grades submit rank-ordered lists of up to six preferred schools to the district. Each school has a preestablished maximum enrollment capacity. Students are assigned priorities at each school based on location of residence, sibling status, and other criteria. The OUSD match applies a deferred acceptance (DA) algorithm that takes preferences, priorities, and school capacities as inputs, and generates a single school offer for each student.

To formalize this centralized assignment process, suppose there are N students indexed by i and J schools indexed by j within the OUSD district. Each student i submits a rank-ordered preference list $R_i = (R_{i1}, \dots, R_{i\ell(i)})$ to the district, where $R_{i1} \in \{1, \dots, J\}$ denotes the identity of i 's most-preferred school, R_{i2} represents her second choice, and so on, and $\ell(i)$ is the number of schools on i 's list. Each student i also has a priority ρ_{ij} at each school j . The vector $\rho_i = (\rho_{i1}, \dots, \rho_{iJ})$ collects student i 's priorities at all schools.

Student i 's *type* in the mechanism is the collection of her preferences and priorities, given by $\theta_i = (R_i, \rho_i)$, which captures all the non-random information about student i that feeds into the assignment system. Each student is also randomly assigned a tie-breaking number $\omega_{ij} \in [0, 1]$ at each school. The augmented priority $\tilde{\rho}_{ij} = (\rho_{ij}, \omega_{ij})$ combines priorities with random tie-breakers, resulting in a strict ordering of students at each school: Students are ranked lexicographically, first by ρ_{ij} and then, in the case of ties, by ω_{ij} . Finally, let c_j denote the capacity of school j . The DA algorithm uses student types and school capacities to compute offers as follows:

- 1) In round 1, each student i applies to his/her most-preferred school according to R_i . Each school j provisionally offers seats to applicants in order of priorities $\tilde{\rho}_{ij}$ up to capacity c_j , and rejects the rest.
- 2) In round $k > 1$, each student i who was not provisionally offered a seat in

round $k-1$ applies to his/her most-preferred school, according to R_i , among those that have not previously rejected him/her. Each school considers its currently-seated and new applicants, and provisionally offers seats to applicants in order of priorities $\tilde{\rho}_{ij}$ up to capacity c_j , rejecting the rest.

- 3) The algorithm terminates when all applicants are seated or every non-seated student i has been rejected by all schools strictly ranked on his/her preference list R_i .

This algorithm generates a school offer $Z_i \in \{1, \dots, J, \emptyset\}$ for each student, where $Z_i = \emptyset$ indicates no offer.²

Deferred acceptance has been shown to be strategy-proof - students cannot improve their offers by mis-representing their preferences when submitting applications to the system (Dubins and Freedman, 1981; Roth, 1982). This property follows from the fact that DA allows a student who is rejected by a preferred school in one round to displace a lower-priority student at a less-preferred school in the next round. Formal strategy-proofness requires that students be permitted to rank all available schools, but OUSD limits students' lists to six schools. In practice, over 90% of students receive one of their top three schools, so the limitation on the length of the list is generally not binding.³ Our analysis therefore treats students' submitted preferences as truthful and assumes these submissions would not change under counterfactual priority rules.

At OUSD, priority groups are defined lexicographically. The first three levels are unimportant for our purposes: Students who live in the district are prioritized over those who do not, children of OUSD employees are prioritized over others, and students with siblings already attending the school are prioritized over those without. We focus on the next priority grouping: Students who live

²If each student were to submit a complete preference ordering schools in the district (i.e., if $\ell(i) = J \forall i$) and if $\sum_j c_j \geq N$, all students would be offered seats. The first condition is not met at OUSD since students only rank up to six schools. The district handles this issue by augmenting a student's list with the student's neighborhood school, if it was not otherwise listed, and then appending other schools at random. This ensures that the algorithm offers every student a school.

³About 98% of students, and all students who listed six schools, are offered one of their listed schools.

within the designated geographic area for a school receive priority over those who live elsewhere. We refer to this as “neighborhood” priority.⁴ When a school is over-subscribed, priority groupings (defined by the interaction of each of the four levels) are admitted as long as there is space, with random tie-breaking used only for the marginal priority group that includes more students selecting a school than there is remaining capacity.

The deferred acceptance algorithm induces a probability of an offer for each student at each school, also known as an *offer propensity score* (Abdulkadiroğlu et al., 2017). Let $Z_{ij} = 1\{Z_i = j\}$ denote an indicator for an offer at school j . The offer propensity score for student i at this school is given by

$$(1) \quad p_{ij} = \Pr[Z_{ij} = 1|\theta_i],$$

where the randomness that defines this probability is determined by the assignment of random numbers ω_{ij} for all students and schools. The offer propensity score equals zero for a student who did not rank a school, and equals one for a student who ranked school j first and is guaranteed an offer due to membership in an inframarginal priority group (e.g., the child of an OUSD employee) or to the school being undersubscribed. A student with $p_{ij} \in (0, 1)$ faces non-degenerate offer risk at school j as a result of randomized tie-breaking.⁵

Our analysis focuses on choice patterns and retention effects for students’ first choices (R_{i1}). The likelihood that student i is offered her first choice is given by the first-choice offer propensity score:

$$(2) \quad p_i^F = \sum_j p_{ij} 1\{R_{i1} = j\}.$$

⁴School neighborhoods sometimes overlap, in which case students will receive neighborhood preferences at multiple schools.

⁵Most obviously, a student in a marginal priority group will receive an offer with a very good random number ω_{ij} and not receive an offer with a very bad one. But the student’s offer may also depend on *other* students’ random numbers - for example, a student in a higher priority group might not take a spot at this school if she draws a favorable random number at a school that she prefers more.

We compute first-choice propensity scores by repeatedly simulating the OUSD offer mechanism, holding fixed students’ preferences and priorities, and calculate p_i^F as the share of simulations in which student i is assigned to her first choice. We next use offers and propensity scores to describe preferences and lottery outcomes.

II. Preferences and retention effects

A. Describing choices and propensity scores

Low-income students are more likely to choose schools in low-income neighborhoods. Moreover, when they do not do so, they are unlikely to get their first choices. Figure 1 divides schools by their share low-income (0-20%, 20-40%, and so on). Bars show the number of students listing a school in each category as their first choice, for low-income students in the left panel and high-income students in the right panel. Each bar is separated into a blue segment indicating the expected number of students offered their first choice schools (calculated as the sum of first-choice propensity scores p_i^F) and a red segment indicating the expected number of students not assigned their first choice (the sum of $(1 - p_i^F)$).

A clear majority (58%) of low-income students list as their first choices schools with low-income shares above 80%. Only 26% choose schools with less than 40% low-income. Nearly all of the students in the former group get their first-choice school, while only 33% of students in the latter group do.

Among non-low-income students, 77% choose first-choice schools with low-income shares under 20%, and vanishingly few choose schools with shares above 60%. High-income students who list low-poverty schools are more likely to be offered them than are low-income students who choose those schools, largely reflecting the importance of neighborhood priorities. This suggests that a weakening of those priorities has the potential to increase the share of low-income students choosing high-income schools who are able to get those schools. However, there are limits to this potential: Even among high-income students choosing high-




FIGURE 1. FIRST CHOICE SCHOOLS AND OFFERS, BY SCHOOL LOW-INCOME SHARE AND STUDENT TYPE - KINDERGARTEN

Note: This figure is constructed by first binning schools by their low-income shares. The bins range from 0-20% on the left to 80%-100% on the right. The left panel shows low-income students and the right panel high-income students. Blue areas represent the share of students who receive offers at their first-choice school, while the red areas reflect the likelihood that a student was not offered their first choice. The height of the bar represents the number of students in each bin. These are computed as sums of the relevant propensity scores, $\sum_{R_{i1} \in b} p_i^F$ and $\sum_{R_{i1} \in b} 1 - p_i^F$, where b indexes bins, separately for low-income and high-income students.

income schools, many of whom have neighborhood priority, a sizable share do not receive their first choice, as these schools tend to be severely oversubscribed.

B. Preferences and the role of distance

Figure 1 indicates that low- and high-income students make sharply different choices among schools. This does not primarily reflect differences in assessments of the inherent valuation of schools, but rather the location of these schools relative to where students live: Students prefer nearby schools, and high-income schools tend to be located far from low-income neighborhoods.

We explore this by fitting random utility models for student preferences separately for high- and low-income students, allowing the two groups to have independent assessments of the different schools along with group-specific distance costs. Specifically, we assume that student i 's rank-ordered list R_i ranks schools in decreasing order of latent utilities generated by the model:

$$(3) \quad U_{ij} = \delta_j(L_i) - \tau(L_i)Dist_{ij} + \eta_{ij},$$

where $L_i \in \{0, 1\}$ indicates whether a student is low-income, $\delta_j(\ell)$ summarizes the overall popularity of school j for income group ℓ , $Dist_{ij}$ represents the distance from student i 's home to school j , and $\tau(\ell)$ is a distance cost for income group ℓ . The unobserved components of utility, η_{ij} , are assumed to be independently and identically distributed according an extreme value type I distribution, which makes equation (3) a rank-ordered logit model (Hausman and Ruud, 1987). We estimate the unknown parameters of this choice model, which consist of the $\delta_j(\ell)$'s along with the distance costs $\tau(\ell)$, by maximum likelihood. We also compare estimates with and without distance included in the model to assess the role of spatial location in generating differences in school popularity by income.

Figure 2 shows scatterplots of school utility parameters against the schools' low-income enrollment share. The left panel presents results from a model that does not allow distance to enter a family's evaluation of a school (i.e., $\tau(\ell)$ is constrained to zero), while the right model does control for distance.




FIGURE 2. STUDENT PREFERENCES AND SCHOOL LOW-INCOME ENROLLMENT SHARE - KINDERGARTEN

Note: This figure plots estimated mean utility parameters (δ 's) derived from a rank-ordered logit model against a school's low-income student enrollment share. In the left panel, the utility parameters come from a model that does not control for distance, while the right panel uses a model that does control for distance. For both panels, the models are estimated separately for low-income and high-income students. The figure excludes high-income δ_j s for two schools with estimated δ_j s below -30 that were selected as the first choice by 0 and 1 students, respectively.

In the left panel, high income students exhibit a strong preference for schools with lower low-income shares. Low income students do not - there is essentially no relationship between a school's low-income share and low-income demand. In

the right panel, however, the two preference curves are much more similar - both high- and low-income students tend to prefer schools with low low-income shares. The estimated δ_j 's for the high-income and low-income groups have a correlation of 0.45 (vs. -0.14 in the left panel).⁶ The contrast is explained by the distance parameter $\tau(\ell)$, which is positive and substantial for both groups of students. Low-income students tend to live farther from high-income schools, and thus are unlikely to rank them highly even though they would, distance aside, prefer them to the low-income schools that they do rank highly.

C. Causal effects of offers to preferred schools

Students' distaste for distance represents a constraint on the district's ability to achieve socioeconomic integration via changes in priority rules. Insofar as students of a particular group avoid a school because it is too far away, there is little to be done about this beyond changing either school or student locations. A second constraint on the district's ability to achieve integration is that students may exit the district in response to offers to less-preferred schools. We explore this possibility by leveraging methods developed by [Abdulkadiroğlu et al. \(2017\)](#) and [Abdulkadiroğlu et al. \(2022\)](#) for extracting quasi-experiments from centralized offer systems with partially-randomized tie-breaking.

For students with the same type in the mechanism, θ_i , differences in school offers are determined solely by random tie-breaking numbers ω_{ij} . Comparisons of outcomes for students offered different schools conditional on θ_i therefore reveal causal effects of school offers. Conditioning on individual values of θ_i is not feasible since few students share the exact same preferences and priorities at all schools. However, the propensity score theorem of [Rosenbaum and Rubin \(1983\)](#) implies that it is only necessary to control for the offer propensity score to isolate random offers and eliminate selection bias. Students with different θ_i 's may share a first-choice propensity score p_i^F , e.g. if they have equal admission chances at a

⁶We have also estimated the correlations between the δ_j s, net of sampling error. These corrected correlations are -0.15 in the left panel and 0.47 in the right panel.

shared first-choice school but different odds at lower-ranked schools.

We estimate causal effects of offers to preferred schools with ordinary least squares (OLS) regressions of the form:

$$(4) \quad Y_i = \alpha + \beta F_i + \gamma p_i^F + \epsilon_i,$$

where Y_i is an indicator for enrolling in OUSD and $F_i = 1\{Z_i = R_{i1}\}$ is an indicator equal to one if student i is offered his or her first-choice school. Variation in F_i conditional on p_i^F solely derives from student i 's random number draw, so is unconfounded. As a result, coefficient β measures the causal impact of a first-choice offer on the likelihood of remaining in the district.⁷

Offer of a seat at a first-choice school increases the likelihood that a student enrolls in an OUSD school. This can be seen in Table 1, which reports estimates of equation (4). The first three columns present results where the outcome is an indicator for enrolling at an OUSD school. We estimate the model first using the full sample, then separately for high- and low-income students. Receiving an offer at a first choice school increases the likelihood of enrolling at OUSD by 16 percentage points on a base of 63 percentage points (for a student with $p_i^F = 50\%$). This effect is larger for high-income students (18 p.p.) and smaller for low-income students (14 p.p.), but the differences are minor.

The remainder of the table considers two additional outcomes: Enrolling at an Oakland charter school, in columns 4-6, or exiting the district (i.e., not enrolling at either an OUSD or a charter school), in columns 7-9. We estimate small, statistically insignificant effects of a first-choice offer on charter enrollment.⁸ This suggests the effects of preferred-school offers on OUSD enrollment operate through the margin of drawing students into OUSD who would otherwise exit the district

⁷Since $E[F_i | p_i^F]$ is linear in p_i^F , it is only necessary to control for a linear term in p_i^F to isolate the causal effect of a first choice offer (Abdulkadiroğlu et al., 2017). If first choice effects vary across values of p_i^F , OLS estimation of equation (4) recovers a convex weighted average of first-choice effects, with weights proportional to sample size times $p_i^F(1 - p_i^F)$ (Angrist, 1998).

⁸When we estimate similar models with 6th grade applicants, we do find negative effects of first-choice offers on charter enrollment.

	OUSD			Charter			Not OUSD or Charter		
	All	High Inc.	Low Inc.	All	High Inc.	Low Inc.	All	High Inc.	Low Inc.
Offered first	0.158 (0.045)	0.179 (0.060)	0.136 (0.068)	-0.017 (0.032)	-0.029 (0.029)	-0.009 (0.057)	-0.141 (0.037)	-0.150 (0.056)	-0.127 (0.048)
Propensity score	0.020 (0.057)	-0.005 (0.072)	0.054 (0.092)	-0.082 (0.039)	-0.068 (0.030)	-0.110 (0.078)	0.062 (0.048)	0.073 (0.069)	0.056 (0.065)
Constant	0.621 (0.025)	0.634 (0.031)	0.601 (0.041)	0.159 (0.018)	0.119 (0.020)	0.222 (0.034)	0.220 (0.021)	0.248 (0.028)	0.177 (0.031)
Observations	901	516	385	901	516	385	901	516	385
R-squared	0.035	0.037	0.033	0.015	0.023	0.015	0.020	0.017	0.021

TABLE 1—FIRST-CHOICE TREATMENT EFFECT BY INCOME STATUS

Note: Data only estimated on 2019-20 data for which census block group-based income is available. Sample is the subset of in-district Kindergarten applicants with a propensity to receive their first choice school strictly between 0 and 1.

entirely - either to another district, or to a private school.

It’s possible that offers at more vs. less-preferred schools affect retention even below a student’s first choice. In additional specifications not reported here, we found no evidence of effects of offers to second vs. lower-ranked choices, though with 76% of students receiving one of their first two choices the data offer limited statistical power for this purpose. Another question is whether the impact of a first-choice offer differs according to school income level or overall popularity. Again, we have found no evidence of this in models interacting first-choice offers with these school characteristics.

The powerful effect of first-choice offers shown in Table 1 indicates that changes in lottery rules that impact the share of students offered their first choice may have important effects on student exit from the district. Because maintaining enrollment is a high priority for OUSD, the need to keep the first-choice share high may act as a constraint on its willingness to consider alternative assignment rules. However, the fact that the “offered first” effect is similar for high- and low-income students suggests that policies that raise the share of one group getting its first choice at the expense of the other may not have first-order effects on overall enrollment. In the simulations presented below, we assess this by using the first-choice effects in columns 2 and 3 of Table 1 to model the effect of changing offer priorities on OUSD enrollment.

III. Priorities and counterfactual assignment rules

A. Defining counterfactual policies

Let i index students and j schools. If student i ranks school j , her priority for the school is defined by whether she resides in OUSD at the time of application, O_i ; by whether she has a sibling already attending school j , S_{ij} ; and by whether she lives in the school’s neighborhood, N_{ij} : $\rho_{ij} = \rho(O_i, S_{ij}, N_{ij})$.⁹ These characteristics are considered lexicographically: OUSD residents are always preferred over non-residents; among students with the same O_i , siblings are preferred over non-siblings; and among students with the same O_i and S_{ij} , neighborhood students are preferred over out-of-neighborhood applicants. As discussed in Section I, ties between students with the same priority are broken with random numbers ω_{ij} drawn from uniform distributions on the unit interval. We can represent the resulting strict ordering as:

$$\tilde{\rho}_{ij} = 100O_i + 6S_{ij} + 2N_{ij} + \omega_{ij}.$$

A student with higher $\tilde{\rho}_{ij}$ will always beat out a student with lower $\tilde{\rho}_{ij}$ for a spot at school j .

Our analysis of counterfactual assignment policies focuses on a class of alternative priorities that privilege students who would increase income diversity relative to the surrounding neighborhood. We consider a family of rules indexed by a parameter $\pi \in [0, 2]$ that governs the probability that a low-income, non-neighborhood student can compete with a neighborhood student for a spot at a high-income school (i.e., a school with a status quo low-income share below the district average of 39%). Let D_{ij} denote an indicator equal to one if student i is low-income, school j is high-income, and i does not reside in the neighborhood of

⁹This neglects some complexities - e.g., that there are priorities for children of OUSD employees. These affect less than 1% of kindergarten applicants. We neglect them in our simulations, both of the status quo policy and counterfactuals. We do account for them in generating the propensity scores used in Table 1.

j (i.e., $N_{ij} = 0$).¹⁰ Our counterfactual rules form new strict priority orderings as:

$$\bar{\rho}_{ij}(\pi) = 100O_i + 6S_{ij} + 2N_{ij} + \pi D_{ij} + \omega_{ij}.$$

The status quo rule corresponds to $\pi = 0$. In this case, within any $\{O_i, S_{ij}\}$ cell, neighborhood students are admitted before any non-neighborhood students are considered, and if these non-neighborhood students are considered they all have equal probabilities of admission. As π rises above zero, non-neighborhood students with $D_{ij} = 1$ are increasingly likely to beat out other non-neighborhood students with $D_{ij} = 0$. When $\pi = 1$, all $D_{ij} = 1$ students are admitted before any non-neighborhood $D_{ij} = 0$ students (again within $\{O_i, S_{ij}\}$ cells), though after all neighborhood students. As π rises above 1, non-neighborhood $D_{ij} = 1$ students with very high ω_{ij} begin to be admitted before neighborhood students with low ω_{ij} . At the maximum value of $\pi = 2$, non-neighborhood $D_{ij} = 1$ students compete on equal footing with neighborhood students, with random ordering of students from the two groups. This, in effect, allows all low-income students to be awarded neighborhood priority at all high-income schools.

In principle, one could consider $\pi > 2$, in which case non-neighborhood diversifying students would have *higher* probabilities of being admitted than neighborhood students. Such rules seem unlikely to be politically feasible, so our analysis focuses on cases with $\pi \leq 2$.

B. Alternative rules lead to more integration

We simulate alternative rules (i.e., different values of π) by applying the deferred acceptance algorithm to students' stated preferences R_i and the modified ordering $\bar{\rho}_{ij}(\pi)$. For each value of π , we conduct 1000 simulations, drawing new random numbers ω_{ij} for each, and average results across them. Figures 3 and

¹⁰In the appendix, we consider an alternative definition in which non-low-income students also get boosts when applying to out-of-neighborhood low-income schools. This makes little difference - as Figure 1, below, indicates, few non-low-income students list low-income schools as their first choices.

4 display the effects of changing π on first choice offers, travel distances, and school segregation. In each panel of the figures, the horizontal axis indexes π , which governs the magnitude of the change to the priority rule. Recall that $\pi = 0$ corresponds to the status quo rule, $\pi = 1$ to a rule where all low-income non-neighborhood students are admitted to high-income schools before any high-income non-neighborhood students (with the same district and sibling priorities), and $\pi = 2$ to a rule where non-neighborhood low-income students compete on equal footing with neighborhood students for admission to high-income schools.

Increasing the diversity preference parameter π has little impact on the overall share of students offered their first-choice schools, but increases travel distances modestly. The left-hand panel of Figure 3 shows that under the status quo rule ($\pi = 0$), about 74% of low-income and 63% of high-income students are offered their first-choice schools. This largely reflects the fact that high-income students' first-choice schools are much more likely to be oversubscribed (see Figure 1). The right-hand panel shows that low-income students travel about 1.1 miles to their offered school, on average, while high-income students travel about 25% further. As π increases above zero, a few more low-income students get their first choices, and they travel further on average. Effects are in the opposite direction, but more muted, for high-income students. As a result, the total share of students offered their first choice schools remains roughly constant, while the average distance traveled increases slightly. Once π reaches about 0.4, further increases have little additional effect until π exceeds 1. (Recall that it is only when $\pi > 1$ that any non-neighborhood students are able to come out ahead of neighborhood students; lower values merely adjust probabilities among non-neighborhood applicants.) As π grows above 1, we see increasing effects, with rising first-choice shares for low-income students and falling first-choice shares for high-income students. As before, the pooled first-choice share is not much affected. Both high- and low-income students travel farther, on average, as the share offered their neighborhood schools falls.




FIGURE 3. EFFECT OF ALTERNATIVE OFFER RULES ON FIRST-CHOICE SHARE (LEFT) AND AVERAGE DISTANCE TRAVELED (RIGHT)

Note: The horizontal axis in each panel represents the diversity priority parameter, π . Each point comes from a separate simulation of the deferred acceptance assignment rule, using students' stated preference rankings but varying the school priorities. Panel A shows the share of students who are offered their first choices, by student income. Panel B shows the mean distance between a student's home and their offered school, again by student income.

Figure 4 shows simulated impacts of the alternative policies on an exposure index measuring the share non-low-income in the average low-income student's school. Higher values of this exposure index correspond to more integration of the two groups of students. We show two series - one that is based on all offers, and another that weights each student's enrollment (and contribution to the school low-income share) by their estimated probability of remaining in the district. This latter series is computed from the results in Table 1 given the student's expressed preferences and simulated school offer.

Beginning with the first, unweighted series, the leftmost point shows that under the status quo, the average low-income student is offered a seat at a school where just 29% of offers go to non-low-income students. As π increases, the exposure index rises, to 0.32 when $\pi = 1$ and then even more sharply to 0.39 when $\pi = 2$. That is, moving from the status quo to a rule that gives all low-income students neighborhood priority at non-low-income schools increases exposure by 0.10. For comparison, in a similar exercise to ours, [Idoux \(2023\)](#) finds that eliminating all geographic and academic preferences in New York City would increase exposure of

low-income to non-low-income students there by only about 0.03. [Reber \(2005b\)](#) finds that the implementation of a court-ordered desegregation plan in the 1960s, 1970s, and 1980s raised a similar measure of Black students' exposure to white schoolmates by between 0.10 and 0.13.¹¹ Our counterfactual policies are thus approximately as effective as much more disruptive school desegregation plans, which often involved forced busing and other dramatic changes to districts.

The initial estimates are based on students' school offers. They may not reflect realized segregation, as not all students comply with their assignments. In particular, students may exit the district in response to an undesired offer. The second series in Figure 4, plotted in blue dots, uses the weighted student counts to simulate each school's composition of actually enrolling students. This somewhat reduces exposure - our estimates in Table 1 indicate that high-income students are more likely to leave the district if offered schools with high low-income shares, which are less likely to be their first choice, and thus that schools are more segregated than initial offers. However, it has little effect on our estimates of the impact of alternative assignment rules. Our alternative rules only raise or lower first-choice shares by ten percentage points or less, which by Table 1 means that weights change by only about 0.015, too little to meaningfully change school composition. Thus, even accounting for student exit, we find large effects of alternative rules on exposure.¹²

¹¹[Guryan \(2004\)](#) finds somewhat larger effects, around 0.15. [Lutz \(2011\)](#), however, finds that removal of a desegregation order reduced exposure by only about 0.03 over ten years. All three studies also report effects on the dissimilarity index, another measure of segregation. We report effects of simulated policies on dissimilarity in Appendix Figure B4.

¹²[Reber \(2005b\)](#) finds that many white families left desegregating school districts in the years following desegregation, an analog to the exit that we study here, and that in the long run desegregation plans raised exposure by only 0.07 to 0.09.




FIGURE 4. EFFECT OF ALTERNATIVE OFFER RULES ON SEGREGATION

Note: This figure shows the exposure index of low-income to non-low-income students in OUSD kindergartens, computed first over all students offered spaces (yellow circles) and then adjusting for enrollment responses (blue triangles). The latter series weights each student by $p(X_i)$, the fitted value from the regressions in columns 2 and 3 of Table 1, before computing first the school enrollment shares and then the exposure index.

IV. Conclusion

Our analysis indicates that neighborhood preferences are an important contributing factor to school segregation in Oakland. Although low-income students are more likely to prefer low-income schools to high-income schools, largely because the former are much closer to home, enough low-income students would prefer to attend high-income schools to create the possibility of meaningful integration. We show that relatively modest changes to the way that priorities are assigned, somewhat weakening high-income students' priority at the schools in their neighborhoods, can have large effects on school segregation. However, these effects come with two costs. First, the share of high-income students who get their first choice schools falls. This drives some students out of OUSD, to private or out-of-district schools. Second, average distance traveled to school rises, for both low-income and high-income students alike. Nevertheless, despite these costs, our results indicate that neighborhood preferences serve as an important contributor to school segregation and that they could feasibly be weakened without driving large numbers of students out of the district.

REFERENCES

- Abdulkadiroğlu, Atila, Joshua D. Angrist, Yusuke Narita, and Parag A. Pathak.** 2017. "Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation." *Econometrica*, 85(5): 1373–1432.
- Abdulkadiroğlu, Atila, Joshua D. Angrist, Yusuke Narita, and Parag A. Pathak.** 2022. "Breaking Ties: Regression Discontinuity Design Meets Market Design." *Econometrica*, 90(1): 117–151.
- Angrist, Joshua D.** 1998. "Estimating the Labor Market Impact of Voluntary Military Service Using Social Security Data on Military Applicants." *Econometrica*, 66(2): 249–288.

- Angrist, Joshua D., and Kevin Lang.** 2004. “Does School Integration Generate Peer Effects? Evidence from Boston’s Metco Program.” *American Economic Review*, 94(5): 1613–1634.
- Angrist, Joshua, Guthrie Gray-Lobe, Clemence M. Idoux, and Parag A Pathak.** 2022. “Still Worth the Trip? School Busing Effects in Boston and New York.” National Bureau of Economic Research Working Paper 30308.
- Bayer, Patrick, Fernando Ferreira, and Robert McMillan.** 2007. “A Unified Framework for Measuring Preferences for Schools and Neighborhoods.” *Journal of Political Economy*, 115(4): 588–638.
- Bifulco, Robert, and Helen F. Ladd.** 2007. “School choice, racial segregation, and test-score gaps: Evidence from North Carolina’s charter school program*.” *Journal of Policy Analysis and Management*, 26(1): 31–56.
- Billings, Stephen B., David J. Deming, and Jonah Rockoff.** 2013. “School Segregation, Educational Attainment, and Crime: Evidence from the End of Busing in Charlotte-Mecklenburg *.” *The Quarterly Journal of Economics*, 129(1): 435–476.
- Cascio, Elizabeth, Nora Gordon, Ethan Lewis, and Sarah Reber.** 2008. “From Brown to busing.” *Journal of Urban Economics*, 64(2): 296–325.
- Corradini, Viola, and Clemence Idoux.** 2025. “Overcoming Racial Gaps in School Preferences: The Effect of Peer Diversity on School Choice.” Working Paper.
- Dubins, Lester E., and David A. Freedman.** 1981. “Machiavelli and the Gale-Shapley algorithm.” 88: 485–494.
- Dur, Umut, Scott Duke Kominers, Parag A Pathak, and Tayfun Sönmez.** 2018. “Reserve design: Unintended consequences and the demise of Boston’s walk zones.” *Journal of Political Economy*, 126(6): 2457–2479.

- Elsen-Rooney, Michael.** 2024. “NYC Reformed High School Admissions 20 Years Ago. Did It Make Things Better?” Accessed January 21, 2026.
- Epple, Dennis, and Holger Sieg.** 1999. “Estimating Equilibrium Models of Local Jurisdictions.” *Journal of Political Economy*, 107(4): 645–681.
- Guryan, Jonathan.** 2004. “Desegregation and black dropout rates.” *American Economic Review*, 94(4): 919–943.
- Han, Raymond, and Clemence Idoux.** 2025. “Equilibrium Effects of Neighborhood Schools.” Working Paper.
- Hausman, Jerry A., and Paul A. Ruud.** 1987. “Specifying and testing econometric models for rank-ordered data.” *Journal of Econometrics*, 34(1): 83–104.
- Idoux, Clemence.** 2023. “Integrating new york city schools: The role of admission criteria and family preferences.” MIT unpublished manuscript.
- Laverde, Mariana.** 2024. “Distance to School and Equal Access in School Choice Systems.” Working Paper.
- Lukas, J. Anthony Lukas.** 1985. “Common Ground: A Turbulent Decade in the Lives of Three American Families.”
- Lutz, Byron.** 2011. “The end of court-ordered desegregation.” *American Economic Journal: Economic Policy*, 3(2): 130–168.
- Marquez Martinez, Nuria.** 2018. “Should Oakland Schools Finally Try to Integrate?” Alameda Magazine (Education). Accessed January 21, 2026.
- Monarrez, Tomás, Brian Kisida, and Matthew Chingos.** 2022. “The Effect of Charter Schools on School Segregation.” *American Economic Journal: Economic Policy*, 14(1): 301–40.
- Monarrez, Tomás E.** 2023. “School Attendance Boundaries and the Segregation of Public Schools in the United States.” *American Economic Journal: Applied Economics*, 15(3).

- Oosterbeek, Hessel, Sándor Sóvágó, and Bas van der Klaauw.** 2021. "Preference heterogeneity and school segregation." *Journal of Public Economics*, 197: 104400.
- Quick, Kimberly.** 2016. "Chicago Public Schools: Ensuring Diversity in Selective Enrollment and Magnet Schools." Accessed January 21, 2026.
- Reardon, S. F., and A. Owens.** 2014. "60 years after Brown: Trends and consequences of school segregation." *Annual Review of Sociology*, 40: 199–218.
- Reber, Sarah J.** 2005*a*. "Court-Ordered Desegregation: Successes and Failures Integrating American Schools since Brown versus Board of Education." *The Journal of Human Resources*, 40(3): 559–590.
- Reber, Sarah J.** 2005*b*. "Court-ordered desegregation: Successes and failures integrating American schools since Brown versus Board of Education." *Journal of Human Resources*, 40(3): 559–590.
- Richards, Meredith P.** 2014. "The Gerrymandering of School Attendance Zones and the Segregation of Public Schools: A Geospatial Analysis." *American Educational Research Journal*, 51(6): 1119–1157.
- Rivkin, Steven, and Finis Welch.** 2006. "Chapter 17 Has School Desegregation Improved Academic and Economic Outcomes for Blacks?" In . Vol. 2 of *Handbook of the Economics of Education*, , ed. E. Hanushek and F. Welch, 1019–1049. Elsevier.
- Rosenbaum, Paul, and Donald Rubin.** 1983. "The central role of the propensity score in observational studies for causal effects." *Biometrika*, 70(1): 41–55.
- Roth, Alvin E.** 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research*, 7: 617–628.

Umosen, Ini, Steven L. Puller, and Cecilia Moreira. 2025. “Transportation as a Barrier to Education Access: Evidence from Chicago Public Schools.” Working Paper.

DATA APPENDIX

The data for this study are provided by the Oakland Unified School District (OUSD). OUSD supplies application and enrollment data from its centralized school assignment system. Data on enrollment at charter schools are provided by Oakland Enrolls, a nonprofit organization which manages the common application process for 98% of Oakland’s charter schools.

A1. School Choice Data

Under OUSD’s open enrollment system, students submit ranked lists of up to six schools. The software for the school choice system automatically generates an application to any neighborhood school that is not manually ranked by the student, ranking these after all submitted choices. As a result, some students have applications for more than six schools.

We receive records for every unique application to a unique OUSD school, along with the priority categories (e.g., neighborhood, sibling, employee) that apply to that application. Students receive neighborhood priority at a school if their official, permanent residence is within the school’s neighborhood boundary. School boundaries sometimes overlap, so it is possible for students to receive neighborhood priority at more than one school.

A2. Enrollment Data

Data on enrollment in OUSD schools come from the OUSD Aeries student information system. Since the application and enrollment datasets do not share a common unique identifier, we link these datasets using a fuzzy matching process that relies on a student’s full name, birth date, and grade level.

A3. Student income

We classify students as low income or not low income based on their home addresses. We geocode these addresses to the Census block group level and match

to data on block group median income and racial and ethnic composition. Following OUSD guidelines, block groups are classified as low income if either (1) at least half of residents are Latino or African American and median income is at or below the free/reduced-price lunch (FRPL) threshold for a family of 5 or (2) at least 20 students in the block group, and at least 70% of students in the block group, are eligible for FRPL.

A4. Sample Restrictions

Oakland processes applications during an “on-time” application window, which generally runs from December to February, and then a late application window that opens in February. We focus on students applying during the 2019-20 on-time application window to enroll in kindergarten in the 2020-21 school year.

The final sample contains 2,296 kindergarten applicants. Of these, 80 have home addresses outside of the district boundaries. Under OUSD’s choice system, in-district students always receive higher priority than out-of-district students. We include out-of-district students when estimating the causal effect of a first-choice offer and for simulations of alternative assignment rules, but exclude them from analyses involving distance (e.g., in the analysis of preferences in Section II.B).

A5. Calculating the First-Choice Propensity Score

In our analyses of the causal estimates of the effect of being offered one’s first-choice school, we rely on a propensity score that measures the likelihood that a student is offered his or her first choice. We construct this by simulating the student-school matching process 1,000 times, drawing new random numbers ω_{ij} for each. A student’s first-choice propensity score is the share of these simulations in which the student is offered his or her first choice school. For robustness exercises, we also computed other propensity scores - for a second-choice offer, for an offer at a low-income school, etc.

ADDITIONAL RESULTS

B1. Data descriptives

Table B1 presents summary statistics for our sample of kindergarteners. Two-thirds of applicants are offered a spot in their first-choice school, and 83% ultimately wind up enrolling at OUSD - 80% in their offered schools and 3% in other schools. Both the OUSD enrollment rate and the share of students who enroll in their offered school decline for students not offered their first choices.

TABLE B1—SUMMARY STATISTICS, 2019-20 KINDERGARTEN APPLICANTS

	By offer status		By rank of offered school				
	Initial	No	(among students with initial offers)				
	offers	initial	1st	2nd	3rd	4th–6th	Unranked
	(1)	offers	choice	choice	choice	choice	school
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
No. of students	2,296	626	1,503	232	113	248	200
% of initial offers			65%	10%	5%	11%	9%
Enrollment outcome (column shares)							
Enroll in Charter	0.08	0.05	0.06	0.10	0.20	0.12	0.11
Enroll in OUSD	0.74	0.80	0.83	0.67	0.55	0.51	0.51
Offered school	0.59	-	0.80	0.34	0.20	0.14	0.08
Different school	0.15	-	0.03	0.33	0.35	0.37	0.44
No information on enrollment school	0.18	0.16	0.11	0.23	0.25	0.37	0.38

Note: Students who received initial offers applied in the first round of the assignment process. Students without initial offers submitted applications after the initial deadline, and are not included in analyses.

Table B2 presents additional information about the schools that students choose, by the school’s rank in the student’s list. A bit more than half of high-income students and a bit less than half of low-income students have neighborhood priority at their first-choice school. In each group, about one-quarter have sibling preference at their first choice. Both shares decline with lower ranks, though the

7th ranked school has a higher neighborhood priority share. (Recall that students rank only six schools, but their lists are augmented with a seventh choice if the neighborhood school is not otherwise ranked.) High-income students list schools with much lower low-income shares and much higher oversubscription ratios than do low-income students.

TABLE B2—STUDENT CHOICE DESCRIPTIVE STATISTICS - KINDERGARTEN

High-Income Students					
	Neighborhood School Priority	Sibling / Continuing School Priority	Low-income Student Enrollment Share	Over- subscription	N
1	0.522	0.25	0.147	1.489	1327
2	0.238	0.006	0.164	1.373	1144
3	0.149	0.002	0.136	1.341	974
4	0.111	0	0.146	1.334	820
5	0.117	0.001	0.146	1.269	668
6	0.168	0	0.166	1.23	549
7	0.485	0	0.284	0.878	340
Low-Income Students					
	Neighborhood School Priority	Sibling / Continuing School Priority	Low-income Student Enrollment Share	Over- subscription Ratio	N
1	0.448	0.262	0.679	1.044	889
2	0.481	0.01	0.675	0.954	770
3	0.377	0.002	0.607	0.972	546
4	0.353	0.005	0.544	1.024	382
5	0.202	0	0.482	1.119	263
6	0.285	0	0.473	1.112	186
7	0.505	0	0.551	0.915	107

Note: Oversubscription ratio is the number of students who ranked a school first divided by the school's capacity.

B2. Inferring student preferences

Figure B1 plots the estimated school utility parameters (δ_j) for each school on the X-axis and the share of students who list the school as their first choice on the Y-axis. Here, the δ_j s come from our model that controls for distance (i.e., from the model plotted in the right panel of Figure 2).

For low-income students, these are basically uncorrelated, reflecting the fact that the highest δ_j schools tend to be far from low-income neighborhoods so are not often listed as first choices. For high-income students, many schools with low δ_j s receive zero first choices¹³, while high- δ_j schools receive large shares of first choices - over 12% for the highest- δ_j school.

Figure B2 plots the school δ_j s from the low-income-student model against the corresponding δ_j s from the high-income-student model. As above, we use estimates from our model that controls for distance. The two groups' preferences are strongly though not perfectly correlated.

B3. Simulations of changing priority rules

Figure B.B3 shows how two policies compare in the share of low-income students offered each school. The X-axis shows the status quo policy, while the Y-axis shows our most aggressive counterfactual policy, with $\pi = 2$. The counterfactual policy increases the low-income share at the schools with the lowest low-income share under the status quo policy, while reducing it at all others.

Figure 4 of the main paper reports the effects of changing neighborhood priority rules on the exposure index, one measure of segregation of low-income and non-low-income students across schools. Figure B4 reports estimates for two alternative segregation measures: The isolation index in the left panel and the dissimilarity index in the right panel. Both indices range from 0 to 1, with 0 corresponding to no segregation and 1 corresponding to complete segregation. In

¹³Recall that we use students' full ranked lists, not just the first choice, to estimate the δ_j s, so it is possible for us to estimate a δ_j for a school that is never ranked first.




FIGURE B1. SCHOOL-LEVEL FIRST-CHOICE SHARE VS MEAN UTILITIES - KINDERGARTEN

Note: The figure describes the relationship between a school's mean utility (δ) and the share of students who ranked a particular school first. School mean utilities are estimated from a discrete choice logit model. The model is estimated separately between low-income and non-low-income applicants to kindergarten. Two schools with estimated δ_j s below -30 (selected as the first choice by 0 and 1 students, respectively) are excluded from the non-low-income series.

each case, as the rules increase the priority for low-income students applying to high-income schools, segregation declines.

Finally, Figure B5 shows the effects of alternative policies. In our primary policy simulations, only low-income students applying to non-low-income schools were eligible for priority boosts. Here, we also provide priority boosts to non-low-income students applying to low-income schools.




FIGURE B2. LOW-INCOME VS HIGH-INCOME MEAN UTILITIES - KINDERGARTEN

Note: Points are mean school utilities estimated from the discrete choice logit model (3). The model is estimated separately for between low-income and high-income applicants to kindergarten.




FIGURE B3. LOW-INCOME SHARE OF ASSIGNED STUDENTS

Note: This figure plots each school's assigned low-income student share when $\pi = 0$ (the status quo rule, with no priority boost for low-income students applying to high-income schools) and when $\pi = 2$ (non-neighborhood low-income students receive full neighborhood priority at non-low-income schools).




FIGURE B4. EFFECT OF PRIORITY RULES ON ALTERNATIVE MEASURES OF SEGREGATION: THE ISOLATION INDEX (LEFT) AND THE DISSIMILARITY INDEX (RIGHT)

Note: See notes to Figure 4.




FIGURE B5. ALTERNATIVE POLICIES THAT PROVIDE BOOSTS TO BOTH LOW- AND NON-LOW-INCOME STUDENTS: POLICY EFFECTS ON FIRST-CHOICE SHARE (LEFT) AND SEGREGATION (RIGHT)

Note: X-axis in each panel represents the policy parameter, π . Panel A shows the share of students who receive their first choices, by student income. Panel B shows the exposure index for OUSD kindergartens, computed first over all students offered spaces (orange) and then adjusting for enrollment responses (blue).