

NBER WORKING PAPER SERIES

WHO CHOOSES AND WHO BENEFITS? THE DESIGN OF PUBLIC SCHOOL CHOICE
SYSTEMS

Christopher Campos
Jesse Bruhn
Eric Chyn
Antonia Vazquez

Working Paper 34581
<http://www.nber.org/papers/w34581>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2025

Campos is listed first as lead author. The remaining authors (Bruhn, Chyn, and Vasquez) are listed alphabetically and contributed equally to the paper. We are grateful for insightful comments from Marianne Bertrand, John Friedman, Nathan Hendren, Matt Notowidigdo, Minseon Park, Canice Prendergast, Chris Walters, and seminar participants at Arizona State University, Opportunity Insights, Rochester University, University of Michigan, University of Chicago, and University of Maryland. Ryan Lee and Connor Fogal provided outstanding research assistance. Campos acknowledges support from the John E. Jeuck Faculty Fellowship. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Christopher Campos, Jesse Bruhn, Eric Chyn, and Antonia Vazquez. 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.

Who Chooses and Who Benefits? The Design of Public School Choice Systems
Christopher Campos, Jesse Bruhn, Eric Chyn, and Antonia Vazquez
NBER Working Paper No. 34581
December 2025
JEL No. I0, I20, I24

ABSTRACT

Public school choice has evolved rapidly in the past two decades, as districts roll out new magnet, dual-language, and themed programs to broaden educational opportunity. We use newly collected national data to document that opt-in (voluntary) systems: (i) are the modal design; (ii) are harder to navigate; and (iii) have participation that is concentrated among more advantaged students. These facts suggest a striking inconsistency: districts have largely adopted centralized assignment algorithms to broaden access, but most rely on optional participation that fragments public education. We study the implications of this design choice in the Los Angeles Unified School District, the largest opt-in system in the country, combining nearly two decades of administrative data, randomized lotteries, and quasi-experimental expansions in access. Participation is highly selective, consistent with national evidence, and lottery estimates suggest that the students with the lowest demand for choice schools are the ones who gain the most from attending. Opt-in participation therefore embeds a selection mechanism that screens out high-return students and leaves many effective programs with unused capacity. To evaluate system-level implications, we estimate a structural model linking applications, enrollment, and achievement. Choice schools are vertically differentiated and generate meaningful gains, but the opt-in participation rule—through high application costs and negative selection on gains—prevents these benefits from reaching the students who need them most. Counterfactual simulations make the design stakes clear: information and travel-cost reductions have limited effects, whereas reforms that change the participation architecture eliminate core inefficiencies and deliver the largest district-wide achievement gains. These results underscore that system design—not school effectiveness alone—shapes who benefits from public school choice and to what extent.

Christopher Campos
University of Chicago
Booth School of Business
and NBER
Christopher.Campos@chicagobooth.edu

Eric Chyn
University of Texas at Austin
Department of Economics
and NBER
eric.chyn@austin.utexas.edu

Jesse Bruhn
Brown University
Department of Economics
and NBER
jesse_bruhn@brown.edu

Antonia Vazquez
University of Texas at Austin
antonia.vazquez@utexas.edu

1 Introduction

Over the last two decades, U.S. public school districts have aggressively expanded school choice—opening new magnet schools, dual-language programs, theme-based academies, and selective-admissions exam schools that collectively provide new intra-district options outside of the traditional neighborhood system. These efforts were motivated by a desire to reduce persistent inequities in educational opportunity and by competitive pressures from charter schools (Betts et al. 2015; Hess et al. 2001; Kahlenberg et al. 2019; Lallinger 2024). The need for such reforms is stark: within the same district, students who attend the lowest-performing schools fall behind their peers at the highest-performing schools by roughly 0.20 standard deviations in achievement per year, a gap that education researchers equate to about two-thirds of a grade level of annual learning.¹ Yet, grafting choice onto a framework built on residential assignment creates an immediate market design challenge: how should these new education markets be organized so that access is equitable and seats are efficiently allocated? The potential benefits of school choice—through competition and better student-school matches (Bruhn et al. 2023; Campos and Kearns 2024; Hoxby 2003; Neilson 2021; Ovidi 2025)—will materialize only if families, especially those who stand to gain the most, are able to navigate these options and submit informed applications.

This paper provides new evidence on the structure of public school choice systems and how their design shapes educational opportunity. To motivate our analysis, we begin by documenting the national landscape of school choice—and how districts have tackled the market design problem—based on original data collected from the nation’s 150 largest school districts. Enrollment in non-neighborhood public schools of choice has more than doubled over the past two decades, with enrollment shares now exceeding those of the charter sector. We also uncover substantial heterogeneity in how families are incentivized to participate. Most large districts (64%) have adopted algorithms to assign students to schools, but a majority also require parents to proactively “opt-in” and submit applications. By contrast, only nine percent of districts require all students to submit preferences by mandate, as is done in cities like New York and Boston (Abdulkadiroğlu and Sönmez 2003; Abdulkadiroğlu et al. 2017). Districts with opt-in participation also provide less clear information and have more complex applications, creating frictions that raise the cost of entry. Consistent with this, we document marked under-representation of economically disadvantaged students in choice schools. These patterns highlight the concern that voluntary participation in school choice programs may segment the public schooling system by allowing the most advantaged students to leave behind their less advantaged peers.

Motivated by these findings, we study the consequences of design: how do the institutional features of opt-in systems shape achievement and inequality, and more generally, how would other designs perform? To frame our analysis, we borrow from Roth (2015)’s focus on questions of who, what, and the how of markets. *Who* opts-in? This question relates to students: are the characteristics of market participants different from those who opt-out? *What* do they choose? This question relates to whether choice schools are vertically differentiated and deliver real benefits for those who take them up. *How* do market design features shape outcomes? This

¹We calculate this statistic by averaging the 90-10 differential in within-district test score growth across all U.S. districts that have at least 10 schools reported in the Stanford Education Data Archive (SEDA). SEDA has harmonized test scores to be comparable across states and districts (Fahle et al. 2024).

question relates to the institutional details—such as the participation rules and the information environment—that can amplify or dilute the interaction between who participates and the options that are available.

To bring these questions into focus, we turn to the Los Angeles Unified School District (LAUSD), the largest voluntary public choice system in the U.S. and a uniquely rich environment for analyzing how design shapes participation and outcomes. Treated as its own district, LAUSD’s choice program would have the 23rd-highest enrollment in the country, serving roughly 110,000 students annually. Notably, the number of choice options has more than doubled in the past two decades, reflecting both parental demand and competitive pressures from the rapidly expanding charter sector. LAUSD’s choice school portfolio now includes hundreds of magnet programs, fifty affiliated charters, and an expanding set of dual-language and theme-based schools. Our analysis relies on nearly two decades worth of detailed application records, linked to district-wide administrative student data containing achievement and other key outcomes. Importantly, the data also include information on the randomized admission lotteries conducted at oversubscribed choice programs. We focus on middle-school applications, where achievement impacts can be observed soon after enrollment.

Our analysis of who participates in LAUSD’s opt-in choice system combines two exercises: a descriptive characterization of applicants and a quasi-experimental analysis of how participation responds to the expansion of programs. Choice schools in LAUSD’s opt-in system attract a student body that is more racially diverse while also being more economically and academically advantaged. Applicants have much stronger academic preparation—the baseline achievement is roughly 0.8σ higher than non-applicants—and are less likely to be low-income or English-learners. Proximity also plays a central role—students who live closer to a choice program are much more likely to apply. To better understand the importance of distance, we leverage the unprecedented expansion of LAUSD’s choice sector, which generated quasi-random changes in access. Event study estimates show that neighborhoods experiencing reductions in relative distance saw participation rates rise by 17 percent on average. These findings demonstrate that while academically advantaged students tend to sort into choice programs, supply-side policy—through program expansion or reductions in travel costs—can broaden access. Whether that expanded participation translates into achievement gains and reduced inequality, however, depends on which students stand to gain the most from attending a choice program.²

Next, we turn to nearly two decades of lottery data to examine which students gain the most from opting-in. Motivated by our analysis of application behavior, we estimate how treatment effects vary with distance to a choice program. We find that students who live close to choice programs benefit the most from attendance: The difference in reduced form impacts between students in the first and the fifth quintile of distance to choice school is roughly 0.1σ . This finding suggests that students with the lowest *unobserved* demand to “opt-in” also stand to gain the most from attendance.³ We further probe this interpretation with two supplemental

²Indeed, prior work studying selection into Charter schools has found that variation in distance predicts patterns of application behavior and treatment effect heterogeneity consistent with inverse-Roy style sorting (Walters 2018).

³More formally, suppose that distance serves as a barrier (i.e. a cost shifter) to application; in that case, our results would imply that the students who gain the most (i.e. those who live closest to a choice program) also face the lowest cost to apply; hence, they will submit an application even at low-levels of latent demand.

exercises. The first leverages the fact that application rates vary greatly by student subgroups (e.g., race or baseline ability) to ask whether students with low *observable* demand also benefit more. The second leverages plausible preferences for homophily (as in Corradini and Idoux 2025) to further probe the role of *unobservable* determinants of demand in a manner that is similar in spirit to the distance-based patterns. In both cases, we find similar results: low-demand predicts a larger benefit. Thus, taken together, our results to this point suggest that opt-in designs generate allocative inefficiency by screening out high-return students.

A key policy implication of this reduced form evidence is that the optional nature of participation can limit the benefits of public school choice. To evaluate whether system-level reforms could improve outcomes, we develop and estimate a structural model linking school applications, enrollment, and achievement. The model recovers the distribution of treatment effects for all students—not just the set of students who currently select into choice programs through the opt-in system. We use a generalized Roy model of application and enrollment demand that follows prior literature (Walters 2018) by linking idiosyncratic school preferences to potential outcomes. The link provides control functions for selection and parameterizes treatment effect heterogeneity. For identification, we leverage two distinct sources of variation: randomized admission offers at oversubscribed programs and policy-driven expansions of LAUSD choice schools that altered students’ relative distance. While lottery offers are standard for identification in this literature, there has been concern over the use of cross-sectional instruments based on distance (e.g. Carneiro and Heckman 2002). Rather than relying on cross-sectional differences in student proximity, we therefore exploit the quasi-experimental changes induced by the choice sector’s expansion, which we show balance well on baseline attributes. Together, these features provide plausibly exogenous variation that allows us to obtain credible estimates of causal effects, which feed into our final counterfactual analysis of how various alternative policies shape enrollment patterns and student achievement.

Our selection-corrected estimates show that choice schools are higher quality, generating average treatment effects of 0.29σ and 0.24σ for math and ELA achievement, respectively. These findings from LAUSD provide evidence that districts facing competitive pressures—such as from charter schools—respond by creating a vertically differentiated product and expanding it. At first glance, the combination of vertical differentiation and strong selection based on prior achievement into choice schools through an opt-in application design might exacerbate educational inequality. However, we find a negative association between preferences and causal effects, indicating that those most likely to opt in also gain the least—patterns consistent with our earlier reduced-form analysis. The difference in average effects between students with the weakest and strongest proclivity for choice school enrollment is on the order of magnitude of -0.30σ , indicating that selection into treatment based on gains dilutes the impacts of opt-in systems on inequality in achievement outcomes. Overall, the characterization of treatment effects for the entire population demonstrates that there is scope for improvements in district-level achievement.

With the structural estimates in hand, we turn to our counterfactual analysis exploring how alternative policies affect sorting and academic performance relative to a baseline simulation of LAUSD’s current opt-in system. Throughout the counterfactuals, we hold fixed the existing menu of programs and school capacities. Our analysis starts by examining participation-focused

policies that lower key access barriers. These include simulations of (i) information interventions that lower search frictions (Ainsworth et al. 2023; Campos 2024; Corcoran et al. 2018; Deming et al. 2014; Figlio and Rouse 2006; Hastings and Weinstein 2008; Neal and Root 2024), and (ii) busing-type policies that reduce travel costs (Angrist et al. 2022; Setren 2024; Trajkovski et al. 2021). We then turn to studying overall system design features by simulating outcomes under (iii) a decentralized market (in which families can apply to multiple schools and receive multiple offers) and (iv) centralized assignment with mandatory participation and deferred acceptance (Lincove and Valant 2023). Finally, we combine participation-enhancing policies with a mandatory assignment regime to evaluate how these reforms jointly shape sorting and achievement. In theory, these counterfactual policies can expand access by lowering information barriers or application costs and encourage participation among students who stand to gain the most. At the same time, each reform also introduces trade-offs as broader participation changes who enrolls and can create capacity pressures, motivating formal counterfactual analysis of how assignment design reshapes access, sorting, and achievement.

We find that general or targeted information interventions modestly broaden participation, while mandates mechanically increase participation. As participation widens, average achievement rises, with the largest gains under a fully centralized system (i.e., mandating participation and using a central clearinghouse for assignment). By contrast, a decentralized market results in more advantaged students occupying more seats in vertically differentiated schools and results in modest gains in district-level achievement relative to the opt-in status quo. Pairing reduced travel costs (via a busing-like policy) with centralization produces modest additional gains, highlighting that high application costs in an opt-in system are more of a barrier than travel costs. In terms of winners and losers of the fully centralized system, we find that all groups experience improvements in average outcomes. These results reflect the elimination of two distinct inefficiencies present in an opt-in system. The first is negative sorting on gains, whereby students with substantial gains to participating self-select out of participation. The second is substantial slack in the system, as many seats at high-quality schools go unfilled. Centralizing participation mostly eliminates both forms of inefficiency, producing achievement gains of 0.018σ and 0.016σ for math and ELA, respectively. The impacts on students who actually enrolled in choice schools (i.e., the treatment on the treated impacts) amount to $0.18 - 0.21\sigma$, compared to $0.08 - 0.11\sigma$ in the status quo opt-in allocation.

To further characterize the potential gains from our counterfactual policies, we conduct a simulation that removes all frictions and allocative inefficiencies by directly matching students to schools to maximize district-level achievement.⁴ Relative to this achievement-optimal benchmark, we find that centralization captures nearly 50 percent of the potential district-level gains in math and ELA, while combining centralization with busing captures over 55 percent. Although none of the policies we study reach the achievement-optimal allocation, the remaining gap points to a fundamental limitation: even when participation is universal and capacity is fully used, demand for effective schools remains imperfectly aligned with school quality. As we demonstrate quantitatively, this misalignment arises because families’ preferences reflect a mix

⁴In the language of Abdulkadiroğlu et al. (2025), this is the “Treatment Effect Maximizing Allocation” (TEMA) and solves a linear program that maximizes student achievement subject to capacity constraints as proposed by Shapley and Shubik (1971).

of academic, social, and informational factors, leading many to favor schools that are not the most effective. Consequently, while centralization substantially improves the efficiency of school assignment, the overall gains from reform are bounded by how demand for school effectiveness shapes equilibrium outcomes.

Our analysis yields two policy lessons. First, rapidly expanding district-run choice programs—often introduced to stem enrollment losses and charter school competition—are vertically differentiated and, on average, raise student achievement, indicating that districts respond competitively when the K-12 landscape shifts. Second, layering voluntary choice on top of neighborhood assignment poses a market-design challenge, and the approach most districts now use appears to deepen inequality in access to opportunity: negative sorting on gains produces allocative inefficiencies, while complex applications and limited outreach leave many high-quality seats empty. Thus, expanding families’ options marks real progress, yet its promise will be met only when assignment mechanisms are tuned to distribute those opportunities equitably—the devil is in the details.

The main contribution of this paper is to the literature on education market design and student outcomes. Across education, health, nonprofit service delivery, and numerous other domains, market design research has shown that well-structured allocation mechanisms can substantially improve welfare (Abdulkadiroğlu et al. 2017; Prendergast 2017; Roth 2015). We begin by documenting how school districts across the U.S. have chosen to organize their education markets—ranging from decentralized systems and canonical mechanisms in the market design literature (Abdulkadiroğlu and Sönmez 2003) to the less-studied *opt-in* design. Prior work shows that uncoordinated school choice policies generate substantial mismatch, while reforms that centralize and coordinate assignment mechanisms yield large welfare gains (Abdulkadiroğlu et al. 2017; Avery et al. 2025). Yet, evidence on how transitions to centralized systems affect student outcomes—the dimension policymakers are most responsive to (Agarwal et al. 2025)—remains limited. We fill this gap by providing a comprehensive analysis of how outcomes vary under the most prevalent design, the opt-in system, and by assessing how alternative market structures and policy choices would influence aggregate achievement and inequality. This work advances our understanding of how education market design translates into market-level changes in student outcomes.

A related literature examines how to improve the performance of centralized assignment systems. This includes work on enhancing the quality of information about schools and admission probabilities (Ainsworth et al. 2023; Andrabi et al. 2017; Arteaga et al. 2022; Campos 2024; Corcoran et al. 2018; Corradini 2023; Corradini and Idoux 2025; Neal and Root 2024) and on the trade-offs between strategy-proof mechanisms and those that capture cardinal preferences (Agarwal and Somaini 2018; Calsamiglia et al. 2020; Kapor et al. 2020). Our data reveal wide variation across U.S. districts in both information environments and mechanism adoption, underscoring the importance of this existing work while also highlighting a largely overlooked design element—the opt-in structure. Drawing on canonical models of self-selection (Heckman and Vytlacil 2005; Walters 2018; Willis and Rosen 1979), we show that participation itself is a critical margin through which design influences both outcomes and allocative efficiency. In doing so, we extend the literature by demonstrating that participation design can be a powerful policy lever for improving student outcomes—one whose welfare implications may diverge from

its effects on achievement. While the existing body of work has focused on frictions, conditional on participation, we show that the first-mile problem of *opting in* is a first-order concern.

We also relate to the literature on how school choice intensifies sorting and stratification across schools. Urquiola (2005) shows that greater inter-district (Tiebout) choice in the United States is associated with sharper segregation of students across districts. In Chile, Hsieh and Urquiola (2006) find that a nationwide voucher program increased between-school stratification by socioeconomic status and ability, with little evidence of system-wide gains. In more recent work, Munteanu (2024) shows that expanded high school choice in Romania raises the variance of test scores and ability sorting without raising mean performance, while Machado and Szerman (2021) document that the introduction of a centralized college admissions platform in Brazil changes the academic and geographic composition of entrants to selective programs. Other work has found that preferences and travel costs also sharply limit school choice effectiveness and contribute to sorting (Idoux 2022; Laverde 2024), while reforms that target participation barriers can improve outcomes (Bergman 2018; Setren 2024). Our findings contribute to this literature by emphasizing that system design, and importantly, the participation rule, is a first-order determinant of student sorting and subsequently reshapes academic outcomes.

Finally, our findings connect to the literature on how public schools respond to competition (Bau 2022; Campos and Kearns 2024; Crema 2022; Figlio and Hart 2014; Gilraine et al. 2021; Hoxby 2000, 2003). Charter schools now educate roughly seven percent of all public school students and have been shown to be vertically differentiated across multiple contexts (Abdulkadiroğlu et al. 2011; Angrist et al. 2013; Dobbie and Fryer Jr 2011), contributing to declining enrollment in traditional districts nationwide (Mumma 2022). Our work departs from these facts and assesses public school district’s response to this competition. We show that public school choice has also grown substantially with enrollment now rivaling the charter sector. Drawing on original data and new analysis, we develop a taxonomy of how school districts structure their education markets and demonstrate that expanding intra-district choice options are also vertically differentiated. In doing so, we show that districts are responding competitively, and that the institutional details of market organization carry significant implications for school districts’ competitive response to outside competition.


2 Public School Choice Systems Across the U.S.

While charter schools have attracted substantial media coverage and policymaker attention in recent years, an equally important transformation has taken place within traditional public school districts themselves. Many large districts now operate extensive “intra-district” choice systems that allow families to select among magnet, dual-language, and themed programs outside their neighborhood zones. Figure 1 plots enrollment shares in district-run choice schools and local public charter schools in the 150 largest school districts.⁵ These statistics show that

⁵For enrollment, all data are from the Common Core of Data (U.S. Department of Education, National Center for Education Statistics 2024). For a given district, a local charter is defined to include all Charter schools that either share the district’s Local Education Agency ID or that operate in a zip-code served by a traditional public school with that district’s LEAID. To identify district-operated choice schools, we rely on the Common Core’s magnet flag, since other forms of choice are not reliably recorded. In a few states (e.g., Ohio), magnet status is not reported. For that reason, we also count a school as a within district choice option if it contains the word

the share of students enrolled in publicly operated choice schools has more than doubled since 1999 and consistently exceeded charter school enrollment. As a whole, more than 14% of all students are now educated in a district-run choice option—over 40% larger than the corresponding market share going to local charters. Appendix Figure A.1 uses data covering all U.S. districts and shows similarly large growth in choice school enrollment.

Figure 1: Public School Choice and Charter School Enrollment Trends (150 Largest Districts)



Notes: This figure reports enrollment shares for charter schools and district-operated choice schools (“District Option”) for the 150 largest school districts. Data are from the Common Core of Data (U.S. Department of Education, National Center for Education Statistics 2024). For a given district, a local Charter includes all Charters that have the district’s Local Education Agency ID or operated in a zip-code served by a traditional public school with that district’s LEAID. To identify district operated choice schools, we use the Common Core’s magnet flag, since other forms of choice are not reliably recorded. In a few states (e.g., Ohio), magnet status is not reported. For that reason, we also count a school as a within district choice option if the word “magnet” appears in its name. The district operated choice numbers in this figure are potentially conservative, since many districts have choice options (e.g., gifted and talented schools) which are not officially named as a magnet and otherwise not easily identified. The 150 largest school districts are those which have the highest enrollment in 2019 and are not comprised exclusively of Charter Schools.

Yet, the impacts of this rapid expansion depend on the details of the choice systems implemented in individual districts. Two dimensions are particularly important: whether participation is mandatory or voluntary, and how that design shapes who ultimately participates. These features vary widely across districts. For that reason, it is useful to highlight a few concrete examples. Boston Public Schools, for instance, requires that all families rank schools and submit them to a centralized clearinghouse, which then determines assignments using an algorithm

“magnet” in its name. Thus, the district-operated choice numbers in Figure 1 are almost surely an undercount, since many districts also have choice options (e.g., gifted and talented schools) which are not officially designated as a magnet and otherwise not easily identified in the Common Core.

(Abdulkadiroglu et al. 2006). In contrast, Miami-Dade County Public Schools defaults families to neighborhood schools unless they voluntarily opt into a centralized system that determines assignment. Districts also vary in the degree to which families are able or willing to participate. Parents in many districts find the system complex and often misunderstand key features.⁶ For example, Chicago Public Schools recently responded to concerns about complexity by undertaking large-scale efforts to simplify the choice system, with the stated goal of reducing inequities in application behavior across socio-economic groups (Sartain and Barrow 2022). Taken together, these examples not only illustrate the diversity of local policies but also underscore the limits of what we know: they come from just a handful of districts. To date, there has been no systematic effort to document how choice systems are actually structured across the U.S., despite the fact that these design features are central to understanding their welfare effects.

To learn more about the structure of public school choice systems nationally, we collected original data for the 150 largest U.S. school districts. Our goal was to understand how large districts vary in the practical, market design features that shape how efficiently students are allocated to schools. Data collection followed a hybrid approach: key information was obtained through direct phone interviews with district representatives and systematic reviews of district websites. We then combine this original data with publicly available data on enrollment shares from the Common Core, producing a dataset that covers districts educating roughly 27 percent of all U.S. public school students. Appendix B.1 provides additional details on the data collection and summary statistics for the sample.

Our analysis of the data yields three key findings. First, the textbook model of mandatory assignment is rare. The most common design instead features voluntary participation with a centralized assignment algorithm. Panel (a) of Figure 2 plots the share of districts in our sample that fall into four different categories of school choice systems: (i) no choice districts that do not offer intra-district choices beyond residential moves, such as Hawaii (1.4%); (ii) decentralized districts (e.g., Dallas ISD) that allow parents to submit applications and do not use a centralized algorithm, thereby allowing families to hold multiple offers (34.7%); opt-in districts (e.g., Clark County, Broward, Los Angeles) which use a centralized algorithm but without mandatory participation (55.1%); and, mandatory districts (e.g., NYC, Houston) that require participation and use a centralized algorithm (8.8%).⁷ Taken together, these results show that the majority of districts offer public school choice following an opt-in model. While prior research has focused heavily on the properties and effectiveness of mandatory centralized algorithms, our analysis reveals that such systems account for only a small share of districts nationwide. By contrast, we are not aware of existing work that systematically studies the design and performance of the opt-in systems with centralized assignment that characterize most real-world school choice markets.

Our second key finding is that choice systems are easier to navigate when school changes are more centralized. A district is classified as difficult to navigate if both the website and phone

⁶For example, recent reporting on San Francisco’s lottery system highlights how parents perceive the process as confusing and stressful, emphasizing the gap between the system’s intended design and how it is understood by families (Li 2025).

⁷In the case of no choice districts, we do not count hardship or capacity-contingent boundary waivers (such as Hawaii’s Geographic Exception permits) as intra-district school choice, since such procedures are limited to discretionary transfers rather than systematic options available to all families.

pathways required multiple steps to learn whether choice exists and how to apply. Panel (b) of Figure 2 plots the share of districts found to be difficult to navigate using the categorization of choice systems defined in the preceding paragraph. At one end of the spectrum, we find that 50% of districts that have no choice system, and hence require a residential move to change schools, were found difficult to navigate, as were 20% of the districts which have decentralized choice systems that require families to apply to schools one-by-one. At the other end of the spectrum, over 7% of districts with mandatory participation were difficult to navigate. Opt-in systems lie in between, with 12.3% of districts being difficult to navigate. These stark differences raise a natural follow-on question: Does “who chooses” vary by system type?

Our final finding is that districts with a more centralized process for changing schools tend to enroll a higher share of low-SES students in choice schools. Panel (c) of Figure 2 illustrates this fact by plotting the average share of the free and reduced price lunch population enrolled in a choice program.⁸ At one end of the spectrum, we find that the typical district in our data that does not offer any district choice option has only 3.3% of their low-SES student population enrolled in a choice program of any kind, with the small share that does exist resulting from non-district charters serving students in that area.⁹ Similarly, choice programs in fully decentralized systems also enroll a relatively small share of the low-SES population (11.6%). At the other end of the spectrum, a typical district with mandatory participation and a centralized algorithm enrolls nearly 30% of low-SES students in choice programs. Finally, we note that opt-in districts tend to fall between these extremes, with the low-SES student enrollment share being 21.5%. Interestingly, the results on low-SES enrollment exhibit an inverse pattern when compared to navigation difficulty, suggestive evidence that the two might be connected.

In sum, our survey evidence establishes three key facts: (i) opt-in systems are most common nationally, (ii) navigation is more difficult in less centralized systems, and (iii) disadvantaged students are less likely to enroll in choice schools in those systems. Taken together, these findings point to a central uncertainty: are the higher-SES students who tend to “opt-in” also those who benefit most from district choice offerings, or are the lower-SES students who are “screened out” by system complexity and other barriers to access the ones who would gain more? To investigate these questions, we now turn to a case study that allows us to examine participation and achievement effects in detail.


3 School Choice in LAUSD

Our national data collection demonstrated that the modal district-level system design is an opt-in system. We now turn our attention to the largest opt-in district in the country, the Los Angeles Unified School District (LAUSD). LAUSD provides a particularly rich setting for analysis because of its scale, diversity, and detailed administrative data. We now provide a brief overview of LAUSD’s choice landscape.

⁸A small number of districts in our sample no longer report free and reduced price lunch enrollment as a result of the movement towards universal meal programs. For that reason, we omit districts where less than 10% of students have free and reduced price lunch from Figure 2.

⁹Including charters in our count of low-SES students enrolled in choice schools is important in order to create an apples to apples comparison to districts like Washington D.C. and Denver where the unified enrollment systems also include non-district charter schools.

Figure 2: Different Systems of Choice Among the 150 Largest School Districts in the US




Notes: This figure reports summary statistics from an original data collection from the 150 largest school districts. Panel (a) plots the share of districts falling into four types of choice systems: “No district choice” refers to districts that do not offer an option to move across district-run programs other than through a residential move; “Decentralized” refers to districts that allow students to apply to non-neighborhood schools or participate in choice programs such as magnet schools, but they do not use a centralized clearinghouse. “Opt-in” districts allow families to submit applications through a choice-school clearinghouse, but participation in the algorithm is not mandatory. “Mandatory” refers to districts that require families to submit preferences over schools and use a centralized algorithm or clearinghouse to match students. Panel (b) plots the share of districts in each of these categories that had a choice system that was difficult to navigate. In the No District Choice category, difficult to navigate refers to learning whether the district offers intra-district choice at all. Panel (c) plots by district-type the average share of low-SES students (measured by Free-and-Reduced Price Lunch (FRPL) eligibility) who are enrolled in a choice program. A small number of districts no longer report FRPL enrollment as a result of the movement toward universal meal programs. For that reason, we omit the six districts where less than 10% of all students receive free and reduced price lunch from the analysis in Panel (c).

3.1 Background on LAUSD’s Choice Programs

The LAUSD has long relied on attendance-zone boundaries to assign students, but since 1982, it has layered on an evolving choice environment. That year, a court-ordered desegregation mandate spurred the creation of 42 magnet programs and schools. Choice broadened further in the 1990s with the advent of charter schools. Following California’s Charter Schools Act of 1992, LAUSD has served not only as an authorizer but also as the governing board for dozens of charter schools. In the 2000s, a variety of pilot initiatives emerged, ultimately leading to the launch of the Zones of Choice program for high-school students, a controlled choice environment in a subset of neighborhoods.

More recently, during the past fifteen years, LAUSD has pursued an aggressive expansion of its portfolio of choice schools. This growth reflects both competitive pressures from the charter sector and concerns about declining district enrollment, particularly in historically underserved neighborhoods. Beginning in the late 2000s, the district more than doubled the number of magnet programs, introduced dozens of affiliated charter conversions, and steadily scaled up dual-language and theme-based offerings. Figure 3 illustrates this trend, showing how the number expanded markedly between 2004 and 2023—a response to declining enrollment trends. Today, LAUSD is the largest district in the country to operate an opt-in system of choice, with roughly 110,000 students enrolled annually across hundreds of district-run choice programs.

Figure 3: LAUSD Enrollment and Choice Availability



Notes: This figure reports trends in LAUSD district-level enrollment and trends in the availability of LAUSD choice programs. The solid black line shows LAUSD enrollment spanning Grades 1-12 between 2001 and 2023. The dashed black line shows the total number of choice programs over the same period, where choice programs include magnet programs, dual-language programs, and affiliated charter schools.

Against this backdrop of expansion, demand for choice seats is strong, with many programs receiving more applications than available seats and, as further detailed below, conducting

randomized admission lotteries. Yet, despite oversubscription at particular campuses, there is nontrivial slack in the LAUSD choice system because capacity growth has outpaced demand in recent years. In the period most relevant for our analysis (2004-2013), aggregate capacity utilization across choice programs was around 60 percent. This coexistence of oversubscription and slack is a central empirical fact that we return to below, as our counterfactual analyses leverage available capacity to assess how changes in participation and application rules would reallocate students across programs.

We now turn to a detailed examination of the prominent alternative choice schooling options and how they relate to the education landscape in Los Angeles. Although this discussion is specific to LAUSD, the varied offerings mirror those of numerous other school districts that have embraced the portfolio management model (Hill and Campbell 2011; Hill 2013), with distinctions, among others, arising in how school districts organize assignment of students to schools (Marsh et al. 2021).

Magnet Schools LAUSD’s first foray into providing alternatives to neighborhood schools came with its magnet school program, launched in 1982 as part of a court-ordered desegregation plan. In the landmark *Crawford v. Board of Education* case, California courts found LAUSD had an obligation to alleviate school segregation. After mandatory busing was largely halted in 1981 due to a statewide ballot measure limiting state courts’ power to mandate busing, the district turned to voluntary integration programs.

LAUSD’s magnet program crept along modestly through the 1980s and 1990s—going from 42 schools at inception to roughly 150 by 2000—but the 2010s saw a rush to expand magnet offerings as charter schools began to chip away at district enrollment. Confronted with dozens of lean, new charters in underserved neighborhoods, the district pivoted to more than double its magnet offerings, from roughly 160 programs in 2008 to over 320 by 2018. Packaged into a new unified enrollment platform, this was a deliberate strategic countermove to reclaim families drawn to the rapidly expanding charter sector (Kohli et al. 2016; Los Angeles Unified School District 2022). Throughout this paper, we treat each magnet program—whether it occupies an entire campus or is co-located with other programs at the same school building—as a separate choice option.

Affiliated Charter Schools LAUSD distinguishes between independent charters and affiliated charters, a model unique to the district’s policy structure. Independent charters are fully autonomous non-profits with their own boards and that run their own admissions processes in accordance with state-mandated lottery provisions. Affiliated charters, in contrast, are semi-autonomous LAUSD schools. They are usually regular neighborhood schools that converted to charter status but remain under district governance, with teaching staff still maintaining ties to the local teacher union. Affiliated charters generally continue to serve their neighborhood attendance area first and then use a lottery for any open seats for out-of-area students, effectively behaving like district schools with expanded autonomy.

LAUSD’s affiliated charter sector evolved from a handful of early conversions in the 1990s and early 2000s into a core pillar of the district’s choice portfolio by the late 2010s. What began with isolated cases—such as Palisades Charter High School’s conversion in 1993—accelerated

around 2009, when district leaders embraced charter conversion as a tactic to stem student losses to competing independent charter schools. Over the ensuing decade, LAUSD authorized the conversion of dozens of neighborhood campuses into affiliated charters. The district now offers 52 affiliated charters, separate from the roughly 150 charter schools they serve as the authorizing body for.

Other Choice Options LAUSD offers additional choice options, including Dual Language programs and admission criteria schools. The former are primarily concentrated at the elementary level, and the latter consist of only a handful of programs. We provide further discussion of these other, less common, choice options in Appendix B.3.

3.2 LAUSD’s School Assignment Policies

As highlighted, LAUSD is the largest opt-in district in the U.S., with its choice offerings beginning with magnet programs created under the district’s court-ordered desegregation plan. From their inception, LAUSD magnet schools have used a centralized lottery system to allocate seats, with built-in design features to promote diversity. Unlike zoned schools, magnets accept students citywide, so the district created a central application and lottery process. Early on, magnets operated under racial enrollment quotas to fulfill their integration mission, often with 60 percent of seats reserved for predominantly Hispanic, Black, Asian, and Other Non-Anglo students. This use of race in magnet assignments continued even after court supervision ended and was later upheld as lawful despite California’s Proposition 209 ban on racial preferences, since the magnet program had been part of a pre-1996 court-ordered remedy (*American Civil Rights Foundation v. Los Angeles Unified School District* 2008).

Historically, LAUSD’s opt-in assignment system allowed families to rank at most one option—a single-application design that features prominently in our empirical analysis. In recent years, the system has expanded: families now submit rank-ordered lists for magnet and dual-language programs through a Unified Enrollment platform, while district-run charters continue to run their own applications systems. When students apply to oversubscribed programs, a single tie-breaking lottery number is assigned to each applicant. In the most recent period, the tie-breaking lottery numbers are still used and seats are allocated using a variant of Immediate Acceptance (First-Preference-First). Participation frictions commonly emphasized in the K-12 market design discourse—e.g., search and information frictions, coordination costs, and the implications of strategic complexity—are also present in LAUSD. The information environment in LAUSD is relatively weak (Campos 2024; Campos and Neilson 2025), but the use of a coordinated system likely helps (Abdulkadiroğlu et al. 2017), while strategic considerations likely disadvantage lower income families in the district (Abdulkadiroğlu et al. 2006). Although these frictions have been shown to be important and payoff relevant, the opt-in design is arguably a much larger participation friction, especially when that is coupled with relative weak information environments as Panel (b) of Figure 2 shows.

3.3 Data and Samples

Our analysis draws on administrative data from the LAUSD spanning 2004-2017, with the years chosen to avoid following students into the pandemic. We focus on cohorts of fifth grade students who are potentially applying to choice schools for the following year. These data include standard demographics, baseline standardized test scores in Math and English Language Arts (ELA), residential Census block, and application information. We link fifth grade cohorts with middle school outcomes in Grades 6-8 and construct two primary achievement measures: a Math score and an ELA score. Each measure is defined as the student-level average of all available standardized test scores in that subject across grades 6-8; when a subject-grade score is missing, the average is computed over the remaining available scores.

We rely on three primary samples for our analysis. The first is a baseline sample, which includes all fifth-grade cohorts observed in LAUSD between 2004 and 2017. This is the broadest sample and is used for our descriptive analysis of who applies to choice programs and how applicants differ from non-applicants. The second sample is restricted to students who applied to an oversubscribed choice school program and participated in an admissions offer lottery that we can observe over this same period. During our study period, roughly 90 percent of all choice programs held at least one lottery. Lotteries are run within strata defined by the cross of application year, student grade, program, race, and priority-points level, and our records contain all fields needed to reconstruct these strata. For each lottery, the files report each applicant’s randomly assigned number and whether an offer was ultimately extended.¹⁰ Since distance features prominently in both the reduced form and structural analysis, we further restrict this sample to students with valid geospatial data representing their home address, since this is necessary to construct measures of distance to a nearby choice school. There are a total of 1,033 lotteries contained in our sample of oversubscribed programs. The third sample is a restricted version of our baseline sample constructed to estimate a structural model of school applications and enrollments. The restrictions are as follows. Students must (i) have non-missing geo-spatial data; (ii) enroll in an LAUSD school in sixth grade; (iii) have at least one test score in Grades 6, 7, or 8; and (iv) be enrolled in the 2004-2013 academic years. This final restriction is convenient for estimating the structural model as it covers a period in which families could apply to at most one school, allowing us to sidestep common issues associated with strategic play or the curse of dimensionality (Agarwal and Somaini 2018).

Appendix Table A.1 provides summary statistics that compare the three samples. Relative to the baseline sample of fifth-grade cohorts, the lottery sample is an advantageously selected subset of applicants, with stronger prior achievement—roughly 0.65σ above district averages in both Math and ELA—and somewhat different demographic composition. This pattern of selection is similar to the differences that we describe in further detail in the next section, comparing all applicants to non-applicants. By contrast, the structural sample resembles the baseline population by construction: while it includes moderately more English learners and slightly lower fifth-grade achievement, in general, it closely tracks the broader LAUSD cohorts.

¹⁰ Appendix Section C details the lottery data and our procedure for inferring each lottery’s cutoff for extending offers. We construct an indicator equal to one if an applicant’s random number lies on the offer side of the inferred cutoff; this is our measure of receiving a lottery-based choice school offer.

4 Who Participates in Choice Schools?

We begin by examining patterns of selection into Los Angeles Unified School District’s (LAUSD) choice programs over the period 2004–2017. Table 1 presents descriptive statistics comparing non-applicants (column 1), applicants (column 2), and the estimated differences between the two groups (column 3). As shown in the top rows, participation in the school choice system is limited, as only 20 percent of fifth-grade students in LAUSD submit an application to a choice program.

Table 1: Summary Statistics: LAUSD Grade 5 Students (2004-2017)

	Non- Applicants (1)	Applicants (2)	Diff. (3)
Observations	477,068	116,482	593,550
Student Share	0.80	0.20	–
Student Demographics			
Hispanic	0.79	0.54	-0.25*** (0.02)
Black	0.09	0.13	0.04*** (0.01)
White	0.07	0.19	0.12*** (0.01)
Asian	0.03	0.11	0.08*** (0.01)
Female	0.49	0.51	0.02*** (0.00)
Poverty	0.75	0.60	-0.14*** (0.01)
English Learner	0.31	0.10	-0.21*** (0.01)
Standardized Test Scores			
ELA	-0.18	0.63	0.81*** (0.03)
Math	-0.16	0.60	0.77*** (0.03)
Distance to District Schools			
Nearest Choice (Rel. Dist.)	0.59	0.51	-0.08*** (0.02)

Notes: This table reports summary statistics that describe demographics, test scores, and proximity to school for a sample of LAUSD fifth-grade cohorts between 2004 and 2017. Columns 1 and 2 report summary statistics for students who did and did not apply to school choice programs in LAUSD, respectively. Column 3 reports mean differences with an associated standard error in parentheses below. The top two rows report statistics regarding their participation in the school choice process, the middle rows report student demographics and baseline test scores, and the bottom rows report proximity-related information. For proximity statistics, we report the relative distance to the nearest district choice school, where relative distance is defined as the difference between distance to the nearest choice school and distance to the nearest district non-choice schools.

Applicants differ systematically from the broader LAUSD student population along both demographic and socioeconomic lines. Hispanic students, who constitute nearly 80 percent of all non-applicant fifth graders, account for just 54 percent of applicants. In contrast, the shares of Black, White, and Asian students are significantly higher among applicants, by 4, 12, and 8 percentage points, respectively, yielding a more racially diverse pool of applicants relative to the overall student body. Applicants are also more socioeconomically advantaged. The share of students identified as living in poverty is 14 percentage points lower among applicants, and the share classified as English Learners is 21 percentage points lower, relative to non-applicants. Female students are also modestly overrepresented, comprising 51 percent of applicants compared to 49 percent of non-applicants.

Turning to academic characteristics, choice applicants are substantially higher achieving than their peers. Their baseline standardized test scores are nearly 0.8 standard deviations higher in both English Language Arts (ELA) and mathematics. This degree of positive selection on prior achievement is consistent with existing evidence from school choice settings. For example, Abdulkadiroğlu et al. (2011), Kline and Walters (2016), and Deming et al. (2014) document that students applying to charter and other selective schooling options tend to be drawn disproportionately from the upper tail of the prior achievement distribution.


Finally, the summary statistics demonstrate that student proximity to schooling options is closely connected to application behavior. Applicants tend to live relatively closer to a nearby choice school when compared to non-applicants. Figure 4 explores this relationship directly by plotting binned averages of the share of students applying to any choice program against relative distance to the nearest school. As illustrated in the figure, application rates decline sharply as distance grows, falling from a high just above 24 percentage points at the closest distance to a low of 18 percentage points at the furthest distance. These patterns are consistent with a large body of prior research documenting that distance is a critical factor in the demand for high-quality schooling options (Abdulkadiroğlu et al. 2020; Hastings et al. 2009; Laverde 2024; Walters 2018).

4.1 Does Distance Determine Participation? Evidence from Choice Expansion

The descriptive patterns show that an applicant’s proximity to nearby choice school programs is an important predictor of participating in LAUSD’s opt-in system. Of course, distance to schools may simply be correlated with other unobserved factors—such as parental motivation, information, or neighborhood characteristics—that also shape demand for choice schools. To more clearly test the importance of distance as a shifter of demand, we use LAUSD’s rapid expansion of choice programs as an empirical testing ground.

As highlighted in Section 3, this expansion had a substantial impact on the availability of programs: the district more than doubled the number of choice programs over the past two decades, creating meaningful variation in access over time. Importantly, as shown in Appendix Figure A.2, the expansion also reshaped geographic access across neighborhoods in Los Angeles. Panels (a) and (b) map the distance to the nearest choice school in 2004 and 2023, respectively, with the darker shading in the latter year reflecting the expansion of choice schools throughout

Figure 4: School Distance Predicts Participation in the Choice System



Notes: This figure reports a binscatter plot assessing the relationship between students’ likelihood of applying to choice programs and relative distance to nearest choice programs within LAUSD sub-districts. There are a total of six LAUSD sub-districts that segment the district into large areas. Choice programs include magnet programs, dual-language programs, and district-run affiliated charter schools. The sample comprises student-level data of all cohorts of fifth-grade students from 2004 to 2017.

the city. Panel (c) summarizes these changes by computing each area’s change in relative distance, showing that neighborhoods in the northeastern Valley foothills (e.g., Sun Valley and Shadow Hills), the western San Fernando Valley (e.g., Woodland Hills and West Hills), and the Harbor Area (e.g., San Pedro and Wilmington) saw the largest increases in choice-program access.

To test whether participation responds to access, we use an event study approach that exploits the policy shock of choice program expansion by comparing changes in application rates across neighborhoods that do experience entry to those that do not. Specifically, we construct our sample as follows. Treated tracts are those with non-zero LAUSD enrollment that experience a reduction in distance to the nearest choice school of more than half a mile in a given year. Control tracts are those with non-zero enrollment that saw no change in distance during the previous three academic years and that are sufficiently far away from treated tracts. More specifically, to mitigate concerns about spatial spillovers, we exclude potential control tracts located within 2.5 miles of a treated tract’s centroid. We then organize these treated and control tracts into a stacked event-study sample, where each “event” corresponds to the first year a tract experiences a qualifying distance reduction. In the stacked sample, we re-index time relative to that event year, pooling across all cohorts to trace out dynamic responses to

improved access.¹¹

Our formal analysis relies on the stacked sample, where we estimate the following event-study specification:

$$Y_{nts} = \alpha_{ns} + \alpha_{ts} + \sum_{k \neq -1} \beta_k D_{ntk} + u_{nts}, \quad (1)$$

where Y_{nts} is an outcome for neighborhood n in event-time t and stack s , α_{ns} are neighborhood-by-stack fixed effects and α_{ts} are event-time-by-stack fixed effects, and $D_{ntk} \equiv \mathbf{1}\{\text{event-time} = k\} \times \mathbf{1}\{\text{Neighborhood } n \text{ is treated}\}$, and u_{nts} is an error term. The inclusion of stack-specific fixed effects restricts to within-stack comparisons of changes in outcomes between treated and control units. The coefficients β_k represent average differences in changes in outcomes between treated and control units between time k and the year before treatment ($k = -1$).

Figure 5 reports results from our event study analysis of choice program demand. We observe a sudden jump in applications that is precisely timed with the expansion. Neighborhoods that experience entry see a roughly 20% increase in neighborhood application rates. Neighborhoods that experience entry were also trending similarly to control neighborhoods in the years leading up to the event, providing reassurance for the parallel-trends assumption underlying our design. Taken together, these results confirm that improvements in geographic access meaningfully raise participation in LAUSD’s opt-in system. In line with the cross-sectional patterns documented in Table 1 and Figure 4, the quasi-experimental evidence reinforces that distance is a key determinant of choice-school demand rather than a mere correlate of unobserved neighborhood or family characteristics.

In sum, our choice program entry event study results show that reductions in distance meaningfully raise participation. This analysis highlights access costs as a central margin in LAUSD’s opt-in system. This pattern aligns with prior studies highlighted above that treat distance as a cost shifter in school choice (e.g., Walters 2018), and therefore naturally raise the question: if distance shifts who applies, does it also shape who benefits when an offer arrives? We provide evidence on this question in Section 5 by testing for distance-graded heterogeneity in lottery-based impacts.


5 Impacts of Choice Schools: Lottery-Based Evidence

Having established who participates in LAUSD’s opt-in choice system, our next goal is to examine which students gain the most from attending these programs. We use a flexible approach to estimating causal effects of enrollment in choice schools that exploits the district’s randomized admission lotteries. Oversubscribed programs allocate seats using random priority number lotteries within strata defined by program, year, grade, priority points, and race combinations.¹²

¹¹The data construction begins with a block-level sample where we calculate changes in distance in order to assign blocks to treatment and control status based on the block’s distance and whether the areas have non-zero LAUSD enrollment. We aggregate the blocks that meet the treatment and control definitions to the tract level. This process discards partially treated blocks (i.e., those that experience smaller reductions in distance). Although the block-level provides the most granular measure, enrollment counts at the block level fluctuate substantially across cohorts, often cycling from a few students to none. Aggregating to the tract level smooths out these fluctuations and yields a more stable unit of analysis.

¹²LAUSD has yet to reach unitary status so still operates with a court-mandated desegregation order and many programs have race-specific quotas.

Figure 5: Changes in Application Rates in Response to Choice School Entry



Notes: This figure reports event-study estimates based on Equation (1) where the dependent variable is the log of the neighborhood applications for choice schools. The event study specification is based on identifying the first academic year in which a neighborhood (census tract) experiences a reduction in its distance to the nearest choice school of more than 0.5 miles. Control tracts exhibit no change in distance over the prior three academic years, and are located at least 2.5 miles from the centroid of any treated tract to mitigate spatial spillovers. Only tracts with non-zero LAUSD enrollment are included in the sample. We report 95-percent confidence intervals where standard errors are robust and clustered at the neighborhood level.

Motivated by the evidence on application behavior and proximity in the preceding section, we estimate how the reduced form impacts of winning an offer vary with baseline distance to the nearest choice school. Specifically, we estimate the ITT effect within distance bins using the following specification:

$$Y_i = \alpha_{\ell(i)} + \sum_{q=1}^5 \beta_q (Z_i \times \mathbf{1}(Q_i = q)) + \sum_{q \neq 1}^5 \kappa_q \mathbf{1}(Q_i = q) + u_i, \quad (2)$$

where Y_i is a post-lottery achievement outcome for student i , Q_i identifies their quintile of baseline distance to the nearest choice school, $\alpha_{\ell(i)}$ are lottery fixed effects, and Z_i is an indicator denoting whether applicant i received an offer in the lottery.¹³ Because Z_i is randomized within each lottery stratum, the β_q identify ITT effects *within* distance bins. Appendix Table A.2 provides evidence of lottery offer balance across baseline student co-variates for each baseline distance bin. The last row of the table reports results from a joint test that all of the coefficients are identically zero, and we consistently fail to reject the null. For inference, we cluster standard

¹³We use distance to the nearest choice school rather than distance to the specific program each student applied to in order to avoid endogeneity concerns: families may strategically apply to distant programs if they particularly value those schools, creating spurious correlations between distance and treatment effects. The nearest-choice-school measure provides a more exogenous proxy for baseline access costs.

errors at the lottery level. We report estimates of β_q and test the hypothesis $H_0 : \beta_1 = \beta_5$.

Figure 6 reports our main reduced-form estimates.¹⁴ Panel (a) of Figure 6 shows that Math ITT effects vary sharply with baseline distance to the nearest choice school: students in the nearest quintile gain roughly 0.05σ when they receive an offer, whereas those in the most distant quintile lose about 0.05σ , and we reject equality across quintiles. Panel (b) displays qualitatively similar patterns. First-stage offer effects are comparable across distance quintiles (0.38-0.45; Appendix Table A.3), implying that the heterogeneity is not driven by take-up differences. Moreover, because the first and fifth quintile effects are opposite in sign, scaling to IV will mechanically magnify the dispersion in impacts evident in Figure 6.

How should we interpret this pattern of results? The distance-based gradient in treatment effects is consistent with findings from Walters (2018), who documents similar heterogeneity in charter-school lotteries: students who live farther from charter schools—those facing higher application costs and thus revealing stronger demand—tend to experience smaller achievement gains once enrolled. This pattern implies a form of reverse Roy sorting, in which students with the greatest willingness to participate are not those who benefit the most. In our setting, reductions in distance appear to draw in additional applicants who are observably similar but differ in their unobserved inclination to apply, revealing that geographic frictions screen out many students with potentially large gains.


This reverse-Roy interpretation based on latent demand raises two main concerns. First, distance is potentially correlated with observable student characteristics, such that the distance achievement gradient may reflect other forms of observable treatment-effect heterogeneity.¹⁵ To explore this possibility, we proceed in two steps. We define subgroups based on baseline characteristics and measure their demand using subgroup application rates; Appendix Figure A.4 plots the resulting treatment-effect estimates against each subgroup’s application rate. Because each point corresponds to a subgroup and its complement (e.g., low- versus high-achiever), we can compare within each pair; in nearly all cases, the group with the higher demand has the smaller estimated treatment effect, further evidence of reverse-Roy sorting. Next, we ask whether this observable heterogeneity can account for the distance gradient itself. To do so, we augment Equation (2) with interactions between offers and these same observable characteristics. Figure 7 reports the results: controlling for effect heterogeneity with respect to baseline achievement, poverty status, and other demographics leaves the distance gradient virtually unchanged, indicating that the pattern in Figure 6 is not driven solely by observable heterogeneity and instead points to an important role for unobserved demand.

Second, one might worry that distance is not the only plausible proxy for unobserved demand. To probe this concern, we construct an alternative index measuring how similar each applicant is to the typical applicant in their lottery. Intuitively, prior research such as Hastings et al. (2009) and Corradini and Idoux (2025) documents that parents and students value


¹⁴The heterogeneity revealed in our main results is obscured when we estimate a pooled specification that ignores distance; the average ITT of winning an offer is negative at -0.027σ ($p < 0.01$) and -0.036σ ($p < 0.01$) for Math and ELA, respectively. This is driven by the disproportionate representation of students with negative achievement benefits at farther distances from nearest choice school in the lottery sample. As will become clear in our structural exercise in Section 7.2, we find that a significant share of students who observably apply to choice schools are “Type 3” students that experience negative gains (see Figure 11).

¹⁵For example, suppose high-ability students both live closer and benefit the most from choice schools. Then the distance gradient in Figure 6 would simply reflect subgroup treatment effect heterogeneity.

Figure 6: Oversubscribed Choice School Lottery Effects by Distance Quintile




(a) Math



(b) ELA

Notes: This figure reports reduced form effects of receiving an offer on Math and ELA achievement in Panel (a) and Panel (b), respectively. We report estimates for different quintiles of distance to choice school at baseline. The estimates come from a regression of achievement—Math or ELA—on lottery strata fixed effects, main effects for distance quintiles, and interactions of offers and distance quintiles. We report estimates of the interactions in each sub-figure. We report 95-percent confidence intervals that use robust standard errors clustered at the lottery strata level. We also report p -values from hypothesis tests that test the null hypothesis that the distance-specific effects are equal and another test that tests the null hypothesis that the distance effects are jointly equal to zero.

Figure 7: Robustness of Distance Gradient in Achievement Impacts



Notes: This figure reports results that assess the robustness of the estimated distance gradient, the difference in estimated impacts of winning a lottery for students who live closest and furthest away from their nearest choice school. All estimates are based on Equation (2) where the coefficients of interest are interactions between a lottery offer indicator and indicators for distance quintiles based on the students' residential address and their distance to choice schools. The left-most results correspond to results from our baseline specification for ELA (gray) and math (black) achievement, respectively. Each additional pair of estimates augments the baseline model with a main effect for the corresponding covariate—e.g., Achievement or Suspensions—and an interaction between a focal covariate and offer indicators. We then calculate the estimated quintile 1 (Q1) and quintile 5 (Q5) difference in reduced form effects and report that with its associated 95 percent confidence interval. We use robust standard errors that are clustered at the lottery strata level.

homophily—preferences for peers who are observably similar to themselves. Thus, students who are “outliers” in the distribution of applicant characteristics must have stronger idiosyncratic, unobserved tastes for the program.¹⁶ Lottery effects by quintiles of this index (Appendix Figure A.3) display a nearly identical negative gradient: students least similar to their peers—those likely to have the strongest unobserved demand for the program—experience the smallest gains. These results reinforce the interpretation that unobserved demand is an important source of treatment effect heterogeneity that we will need to incorporate into our structural analysis in the next section.

6 Beyond Lottery Effects

The lottery-based analysis provides valuable evidence on the effects of choice programs for families who apply. However, as shown in Table 1, less than a quarter of all families submit an application, and an even smaller subset appears in the highly selected lottery sample. To understand how the opt-in design shapes aggregate outcomes and inequality, we must look beyond these applicants. Doing so requires a framework that connects application and enrollment decisions to potential outcomes, allowing us to generalize treatment effects to the broader pop-


¹⁶We quantify this by computing the Euclidean distance between each applicant's observable characteristics and those of the average applicant to their program, producing a simple peer-similarity index.

ulation not currently engaged with the system. Equally important, such a framework provides the foundation for counterfactual policy analysis—enabling us to evaluate how alternative designs, such as reducing application costs or mandating participation, would change district-wide outcomes. In this section, we develop such a model, linking choice school demand to potential outcomes in a way that allows us to address the paper’s core questions.

6.1 Setup and Timing

Our framework models the optimal choice problem faced by forward-looking students and their families when deciding whether to apply to a LAUSD choice program and which program to select. The figure below illustrates the sequence of decisions and events in our model. In Stage 1, consistent with our interest in LAUSD and its policies during the 2004–2013 academic years, each rising sixth-grade student i chooses whether to apply to a single choice school program $j \in \{1, \dots, J\}$, with the binary variable A_{ij} indicating an application to school j . Students are always guaranteed access to their default neighborhood school, denoted $j = 0$. In Stage 2, each choice school $j > 0$ makes admission offers to applicants, with Z_{ij} denoting whether the student receives an offer; when schools are oversubscribed, admission is determined by randomized lotteries, generating exogenous variation in offers. In Stage 3, students make enrollment decisions over programs, denoted by $E_i \in \{0, 1, \dots, J\}$ where $E_i = 0$ indicates neighborhood school attendance. Finally, students attend their chosen schools, take achievement exams, and realize academic outcomes Y_i in Stage 4.

Figure 8: Stages of Decisions and Outcomes in the Model



Notes: This figure summarizes key decisions and endogenous variables in our model.

6.2 Student Preferences

Let $i \in \mathcal{I}$ index rising sixth-grade students in LAUSD and $j \in \mathcal{J}$ index the set of choice schools. Each student may apply to any choice school and, if admitted, enroll there; otherwise, they default to their neighborhood school ($j = 0$) which varies across neighborhoods. We define the indirect utility of student i from enrolling in choice school j relative to their neighborhood school as:

$$U_{ij}^* = U_{ij} - U_{i0},$$

so that $U_{ij}^* > 0$ indicates a strict preference for choice school j over the default.

This utility difference is specified as follows:

$$U_{ij}^* = \underbrace{\theta_i + \beta_c X_i + \delta_j - (\psi_D D_{ij} + \psi_{D2} D_{ij}^2 + \psi_{Dx} D_{ij} \times X_i)}_{v_j(\theta_i, X_i, D_i)} + \epsilon_{ij}, \quad (3)$$

where X_i is a vector of observable student-level characteristics, and D_{ij} measures the residential distance between student i 's address and school j . In our framework, the term δ_j measures the mean popularity of choice school j and serves as the component that absorbs average differences in school quality and preferences over achievement. This model allows for two main sources of choice school preference heterogeneity for each student i . The first is observable preference heterogeneity embedded in the vector X_i , which includes race, socioeconomic status, baseline achievement, neighborhood quality, and baseline peer achievement scores among other common demographic characteristics. The preference vector β_c does not depend on j , implying that observable preference heterogeneity applies uniformly across all choice schools. Second, the term θ_i allows for unobservable preference heterogeneity. We model the travel costs of distance flexibly by specifying a quadratic function for D_{ij} that also features an interaction between distance and the observable student characteristics X_i . Finally, ϵ_{ij} represents idiosyncratic heterogeneity in student tastes for specific school-student pairings, which also encompasses school-specific match effects in realized achievement.

The individual taste parameter θ_i plays an important role in our model of student preferences. A few points about θ_i are worth highlighting for our subsequent analysis. First, because θ_i enters identically for every $j > 0$, it shifts all choice-school utilities in the same direction, implying symmetric substitution patterns between choice schools. Second, we must make assumptions on θ_i 's distribution $F(\theta_i)$ to make the model tractable. We flexibly approximate $F(\theta_i)$ via a finite mixture of K -many normals. Third, as we detail further below in our model of achievement, we allow student achievement outcomes to depend on θ_i , where it captures the relationship between unobserved choice school demand and potential achievement outcomes.

Application Decisions: The application decision depends on both preferences for schools outlined in Equation (3), the probability of receiving offers, and application costs. Consistent with the LAUSD application rules in place from 2001 to 2013, we model each family as being permitted to apply to a single program.¹⁷ The application vector is defined as $A_i = (A_{i1}, \dots, A_{iJ})$ where A_{ij} is the indicator for whether student i applies to school j . Similarly, the offer vector is $Z_i = (Z_{i1}, \dots, Z_{iJ})$ where the indicator Z_{ij} denotes receipt of an offer. As noted above, offers are realized from a lottery embedded in the centralized assignment system conditional on student i applying and school j being oversubscribed. Students face different probabilities, denoted π_{ij} , of receiving a lottery-based offer to school j based on priority groups defined based on sibling status, race, and neighborhood proximity. Given the uncertainty in admissions, the expected utility from a choice set that includes an offer from j , $Z_{ij} = 1$, is:

$$u(Z_{ij} = 1 \mid X_i, D_i, \theta_i) = \log \left(1 + \exp((\delta_j + \theta_i) + \beta_c X_i - (\psi_D D_{ij} + \psi_{D2} D_{ij}^2 + \psi_{Dx} D_{ij} \times X_i)) \right) + \gamma, \quad (4)$$

where γ is Euler's constant. Therefore, the ex-ante expected utility from applying to j is:

$$e(A_{ij} = 1 \mid X_i, D_i, \theta_i) = \pi_{ij} u(Z_{ij} = 1 \mid X_i, D_i, \theta_i) + (1 - \pi_{ij}) \gamma. \quad (5)$$

¹⁷While our baseline model reflects LAUSD's historical one-application regime, we relax this restriction in the "decentralized" counterfactual scenario (Section 8), where families are permitted to submit multiple applications.

Students weigh expected utility from different choices of application vectors against their application costs. Let $a = (a_1, \dots, a_J)$ denote a possible choice school application vector. The utility cost of submitting this set of application choices is given by $c(a, X_i, \eta_i)$ and η_i governs unobserved application cost heterogeneity. We assume $\eta_i \sim N(0, \sigma_\eta^2)$, and the realizations of η_i are known to each student. Application costs are assumed to have the following form:

$$c(a, X_i, \eta_i) = \mathbf{1}\{\sum_j a_j > 0\} \times (\exp(\mu_c + X_i \mu_{cx}) + \eta_i), \quad (6)$$

where the parameters μ_c and μ_{cx} are the mean application cost common to all applications and the heterogeneous application costs with respect to observables, respectively. The unobserved application cost η_i is only incurred when a student chooses to apply to the given school. The functional form of application costs allows us to separately identify preference heterogeneity from cost heterogeneity.

Putting this all together, students select an application vector $A_i \in \mathcal{J} \cup \{0\}$ that maximizes their utility:

$$A_i = \arg \max_a \sum_j \left(e(a_j = 1 \mid X_i, D_i, \theta_i) \right) - c(a, X_i, \eta_i). \quad (7)$$

Equation (7) does not have a closed form expression, so we follow standard approaches in approximating the choice probabilities using a logit kernel smoother (Train 2009; Walters 2018). For a small λ , we can approximate the application choice probabilities as:

$$P_A(A_i = a \mid X_i, D_i, \theta_i, \eta_i) \approx \frac{\exp\left(\{\sum_j e(a_j = 1 \mid X_i, D_i, \theta_i) - c(a, X_i, \eta_i)\}/\lambda\right)}{\sum_{a'} \exp\left(\{\sum_j e(a'_j = 1 \mid X_i, D_i, \theta_i) - c(a', X_i, \eta_i)\}/\lambda\right)}, \quad (8)$$

where a' is the index over the set of possible application vectors.

Enrollment Decisions: When students submit applications, they subsequently receive an offer set $\mathcal{O}_i$ that includes the single choice school to which they applied and their default neighborhood school; students who did not apply have only their neighborhood school in $\mathcal{O}_i$. At the time of enrollment decisions, we introduce post-lottery preference shocks ξ_{ij} , which follow a Type I Extreme Value distribution. These shocks capture the possibility that some students may apply to a choice school, receive an offer, yet ultimately decline enrollment. Conditional on the realized offer set and the vector of covariates $(Z_i, X_i, D_{ij}, \theta_i)$, the probability that student i enrolls in school j —conditional on receiving an offer—is therefore given by a standard logit over the available options:

$$P_E(E_i = j \mid Z_{ij} = 1, X_i, D_i, \theta_i) = \frac{\exp(v_j(X_i, D_i, \theta_i))}{\sum_{j' \in \mathcal{O}_i} \exp(v_{j'}(X_i, D_i, \theta_i))}, \quad (9)$$

where $v_j(\cdot)$ denotes systematic utility. Equations (8) and (9) demonstrate that both observable and latent tastes for choice schools, θ_i , influence application and enrollment decisions. In the achievement outcome model specified below, it is this model-implied variation that we use to account for selection into both application and eventual enrollment. We return to this below,

but first discuss estimation of the model.

Estimating Preference and Cost Parameters: Conditional on the unobservable terms θ_i and η_i , Equations (8) and (9) imply that the likelihood contribution of student i who submits application A_i , receives offer vector Z_i , and enrolls in school E_i is:

$$\mathcal{L}(A_i, Z_i, E_i \mid X_i, D_i, \theta_i, \eta) = P_A(A_i \mid X_i, D_i, \theta_i, \eta) \times P(Z_i \mid A_i, X_i) \times P_E(E_i \mid Z_i, X_i, D_i, \theta_i),$$

where $P(Z_i \mid A_i, X_i)$ is the probability mass function of offers, conditional on applications. In practice, we integrate out the random coefficients θ_i and η_i by simulating draws from the postulated distributions:

$$\begin{aligned} \mathcal{L}(A_i, Z_i, E_i \mid X_i, D_i) &= \int P_A(A_i \mid X_i, D_i, \theta, \eta) \times P(Z_i \mid A_i, X_i) \\ &\quad \times P_E(E_i \mid Z_i, X_i, D_i, \theta) dF(\theta, \eta \mid X_i, D_i). \end{aligned} \tag{10}$$

In our main analysis, we assume $K > 1$, implying that the density is:

$$dF(\theta, \eta \mid X_i, D_i) = \left(\sum_{k=1}^K p_k \phi(\theta; \mu_k, \sigma_k^2) \right) \phi(\eta; 0, \sigma_\eta^2) d\theta d\eta.$$

Therefore, the individual likelihood contribution in the general case is:

$$\begin{aligned} \mathcal{L}(A_i, Z_i, E_i \mid X_i, D_i) &= \sum_{k=1}^K p_k \left(\int P_A(A_i \mid X_i, D_i, \theta, \eta) \times P(Z_i \mid A_i, X_i) \right. \\ &\quad \left. \times P_E(E_i \mid Z_i, X_i, D_i, \theta) \phi(\theta; \mu_k, \sigma_k^2) \phi(\eta; 0, \sigma_\eta^2) d\theta d\eta \right), \end{aligned} \tag{11}$$

where p_k correspond to the type- k share, a quantity that is also estimated in the $K > 1$ case. We maximize the simulated log likelihood function in Equation (11) using all Grade 5 students in the sample. We recover a vector of parameters, $\Omega = (\{\delta_j\}, \beta_x, \psi_D, \psi_{D2}, \mu_c, \mu_{cx}, \{\mu_k, \sigma_k, p_k\}, \sigma_\eta)$. See Appendix Section D for further details on the demand model and estimation.

Identification of Demand Parameters: Before turning to the model of student achievement in the next section, we briefly discuss the identification of our choice school demand model. We organize our discussion around three ingredients: (i) enrollment preferences where identification is aided by randomized offers Z_i ; (ii) application costs identified from extensive margin behavior; (iii) the idiosyncratic preferences for choice schools θ_i recovered from the joint behavior of application and enrollment decisions.

First, consider the enrollment decision within our model. Here, we highlight that randomized choice-school offers Z_i are realized after applications are submitted, ensuring that they are exogenous to unobserved tastes and application costs. This timing implies that Z_i is excluded from the application decision and helps identify the mean utilities δ_j from observed enrollment patterns. Variation in Z_i within lottery strata provides random changes in the feasible enrollment set, allowing the model to trace how take-up varies across schools and thereby recover mean utilities and preference loadings. Cross-sectional differences in enrollment behavior identify β_c

by comparing how choice school demand varies with demographics and baseline achievement, while the degree to which otherwise similar students vary enrollment with respect to changes in distance identifies the travel-cost parameters (e.g., ψ_D and ψ_{D2}).

Next, our second point of discussion focuses on the application side of the model given its central importance for analyzing counterfactual policies that shift participation in the choice system. Given enrollment preference parameter estimates, the application decision depends on the expected utility of enrolling at each school j and on the known offer probabilities π_{ij} implied by the lottery rules and capacities. Because offers are realized after applications, Z_i is excluded from this decision, and the enrollment-side preference estimates ensure that we are separating participation frictions from tastes rather than conflating them. In the baseline model, where each student can submit at most one choice-school application, identification of the application cost parameters comes from variation in the propensity to apply at all—that is, from overall application rates vary across cohorts, neighborhoods and demographic groups as expected utilities and π_{ij} change. This extensive-margin identification of application frictions parallels approaches that separate preferences from participation/search costs in education markets (e.g., Fu 2014).¹⁸

Finally, the third point that we highlight is that the idiosyncratic parameter θ_i acts as a common “taste for choice schools” that shifts both application and enrollment utilities. It is disciplined by the joint behavior of application and enrollment decisions: stronger substitution toward choice schools at enrollment and higher propensities to apply when expected option values rise both load onto θ_i . This is a standard finite-mixture/random-coefficients identification argument: multiple observed margins that share a latent factor help recover its location dispersion (Train 2009). In our $K > 1$ specification, the mixture parameters $\{\theta; \mu_k, \sigma_k^2\}$ are recovered from heterogeneous responsiveness on the enrollment and application margins, while distance identifies travel disutility, so θ_i is not correlated with geography.

6.3 Student Achievement Model

We are interested in a model of the causal effects of choice schools on the average student in the district. Our framework posits that students choose schools optimally implying that students enrolled at any given school are not a random sample of the population. We address selection by allowing mean potential outcomes to depend on the unobserved preferences that shape school enrollment using two primary approaches. Specifically, we begin by assuming:

$$E[Y_{ij} | X_i, D_i, \tau_i] = \alpha_j + \gamma'_0 X_i + \gamma'_c X_i \times \mathbf{1}\{j > 0\} + h_j(\tau_i) \quad j = 0, 1, \dots, J, \quad (12)$$

¹⁸ Andrews et al. (2020) recommend transparent, low-burden diagnostics that clarify which descriptive moments discipline particular parameter estimates in structural models, while emphasizing that such exercises speak to transparency and local informativeness rather than formal identification. In this spirit, we shed light on the key features of our data that drive the fixed application-cost estimate. We focus on this parameter and scale it from 0.5 to 1.5 of our main estimate; we then recompute the model-predicted overall application rate, holding enrollment-side preference parameters at their estimates and using the known offer probabilities implied by the lottery rules, and compare the prediction to the observed rate. Appendix Figure A.5 shows our results and demonstrates that predicted and empirical application rates coincide at the estimated value, while small deviations produce sharp over- or under-prediction, mirroring the single-crossing diagnostic in Autor et al. (2019).

where the indicator $1(j > 0)$ denotes whether school j is a choice program, $h_j(\cdot)$ is a function that satisfies $E[h_j(\tau_i)|X_i]=0$, and $\tau_i \in \{1, \dots, K\}$ represents the latent type indicating which component distribution of $F(\theta)$ generates θ_i .

Our approach imposes the following restrictions. First, we exclude the lottery offer vector Z_i from the achievement potential outcomes, a standard approach in the school choice literature. This assumption embodies the idea that lottery offers have no direct impacts on achievement. Second, we exclude D_i from the achievement equation, so that residential proximity influences achievement only through its role in shaping the costs of application and enrollment. This restriction parallels the use of distance instruments in the charter school literature and rules out direct effects of commuting distance on test scores once school choice is held fixed. Third, we assume application costs $c(a, X_i, \eta_i)$ are unrelated to potential outcomes, so that heterogeneity in costs influences selection into the applicant pool but not academic achievement directly. Fourth, we exclude the post-lottery preference shocks ξ_{ij} from potential outcomes, reflecting the view that these shocks capture idiosyncratic realizations of taste at the enrollment stage rather than achievement-relevant information.

Taken together, these restrictions imply that any bias plaguing ordinary least squares estimates of the choice school effect is due to preference heterogeneity governing application and enrollment decisions captured by an individual's type τ_i . In practice, we impose the following parametric restrictions on h_j . Let $T_{ik} = 1(\tau_i = k)$ with $\sum_k T_{ik} = 1$. Our parameterization of h_j assumes:

$$h_j(\tau_i) = \sum_{k \neq 1} \gamma_k T_{ik} + \sum_{k \neq 1} \gamma_{ck} T_{ik} \times \mathbf{1}\{j > 0\}. \quad (13)$$

The parameters γ_k and γ_{ck} allow for flexible heterogeneity in terms of selection on levels and selection on gains into the choice sector. That is, each type's achievement in the untreated state (attending a neighborhood school) is allowed to differ and the causal impacts of choice school enrollment are also allowed to freely vary. A finding that $\gamma_{k'} > \gamma_k > 0$ for $k' > k$ is indicative of positive selection into the choice sector on achievement levels. Similarly, finding that $\gamma_{ck'} > \gamma_{ck} > 0$ is indicative of positive selection on achievement gains. By allowing for type-specific selection parameters, the model allows for rich sorting patterns that do *not* need to be symmetric across types.

As an alternative approach, we can model selection by allowing potential outcomes to depend directly on the unobserved preferences θ_i . Specifically, we assume:

$$E[Y_{ij} | X_i, D_i, \theta_i] = \alpha_j + \gamma'_0 X_i + \gamma'_c X_i \times \mathbf{1}\{j > 0\} + g_j(\theta_i), \quad j = 0, 1, \dots, J \quad (14)$$

where $g_j(\cdot)$ is a function that satisfies $E[g_j(\theta_i)|X_i]=0$. We follow prior work by assuming a linear functional form for g_j (Abdulkadiroğlu et al. 2020; Bruhn et al. 2023; Otero et al. 2023; Walters 2018). That is, we assume:

$$g_j(\theta_i) = \gamma_\theta \theta_i + \gamma_{c\theta} \theta_i \times \mathbf{1}\{j > 0\}. \quad (15)$$

In this model, the parameters γ_θ and $\gamma_{c\theta}$ serve the same interpretive roles as before, capturing selection on achievement levels and on treatment gains, respectively. If $\gamma_\theta > 0$, then students with stronger idiosyncratic tastes for choice schools tend to have higher achievement, regardless

of their eventual enrolled school. Evidence that $\gamma_{c\theta} > 0$ implies that families select schools based on achievement-based comparative advantage.

Our assumptions imply that we can write control functions for both models as the following quantities:

$$E[h_j(\tau_i) \mid X_i, D_i, A_i, Z_i, E_i] = \sum_{k \neq 1} \gamma_k p_{ik}^* + \sum_{k \neq 1} \gamma_{ck} p_{ik}^* \times \mathbf{1}\{E_i > 0\} \quad (\text{M1})$$

$$E[g_j(\theta_i) \mid X_i, D_i, A_i, Z_i, E_i] = \gamma_\theta \theta_i^* + \gamma_{c\theta} \theta_i^* \times \mathbf{1}\{E_i > 0\}. \quad (\text{M2})$$


In the first expression, $p_{ik}^* = E[T_{ik} \mid X_i, D_i, A_i, Z_i, E_i]$ is the posterior probability of belonging to type k , conditional on observed choices which are functions of student attributes, relative distance, application decisions, lottery offers, and enrollment decisions. Similarly, $\theta_i^* = E[\theta_i \mid X_i, D_i, A_i, Z_i, E_i]$ is the posterior mean of θ_i . All posterior means are calculated by simulations after estimating the demand model. As in Heckman (1979), the p_{ik}^* or θ_i^* terms serve as a control function that accounts for the non-random sorting of students across schools. By conditioning on the simulated posterior of θ_i or type probability p_{ik} , the model accounts for unobserved preference heterogeneity that jointly influences school enrollment and achievement, allowing us to obtain unbiased estimates. Having the control functions in hand, we estimate the empirical analogs of Equations (12) and (14) in a specification that is augmented with neighborhood and year effects. Importantly, the outcome model is not restricted to lottery applicants and accounts for selection into application and enrollment, allowing us to characterize treatment effect heterogeneity for the entire population.

Identification of the Student Achievement Parameters: Identification of the student achievement model relies on variation in lottery offers and relative distance. Intuitively, lottery offers provide exogenous variation in terms of who enrolls, conditional on applying, while variation in distance changes the composition of the applicant pool. The advantage of including neighborhood and year effects is that they allow us to leverage policy-induced changes in distance for identification, rather than relying solely on cross-sectional variation. This feature of our approach provides a source of variation that is plausibly exogenous with respect to achievement outcomes.

Distinct from related prior work on charter school demand and achievement in Walters (2018), our design exploits policy-driven changes in proximity generated by LAUSD’s rapid expansion of choice programs. As established in Section 4.1, new program entry substantially increases local application rates (Figure 5). Here, we extend that evidence to the key identification question for achievement: do these proximity shocks alter who applies on observables, or do they primarily shift latent demand?

To assess composition, we construct a tract-level composition index that is constructed by first calculating the fitted values from a regression of student-level achievement on a vector of baseline covariates and then averaging the fitted values among applicants within a candidate tract. We then estimate the same stacked event-study used for applications but replace the outcome with this index (and its components). Figure 9 shows flat pre-trends and no discernible post-entry shifts in the index. In short, expansions draw in more applicants without

Figure 9: Changes in Student Composition in Response to Choice School Entry




Notes: This figure reports event-study estimates based on Equation (1) where the dependent variable is a neighborhood-level composition index among applicants to choice schools. The composition index is constructed by first calculating the fitted values from a regression of student-level achievement on a vector of baseline covariates and then averaging the fitted values among applicants within a candidate neighborhood. We report 95-percent confidence intervals where standard errors are robust and clustered at the neighborhood level.

systematically changing the observable composition of applicants.

We also verify that changes in proximity to choice schools are not systematically related to student characteristics. Figure 10 summarizes results from student-level regressions that examine whether relative distance to the nearest choice school is correlated with baseline characteristics. When the specification includes only academic year fixed effects (circle markers), distance is moderately correlated with several key student observables (e.g., race and poverty measures). However, once both year and neighborhood fixed effects are included (triangle markers)—so that identification comes from within-neighborhood changes in proximity—the estimated correlations largely shrink and are uniformly small. This pattern confirms that variation in access induced by new school openings is not predicted by observable student characteristics.

Taken together, these facts line up with the achievement model introduced above: proximity shocks raise applications among students who are observably similar but differ in their latent preference for choice schools, θ_i . We therefore use two sources of quasi-experimental variation to identify achievement effects and their heterogeneity: (i) randomized lottery offers, which shift enrollment conditional on applying, and (ii) expansion-induced changes in relative distance, which shift the composition of applicants on unobservables (i.e., θ_i) rather than average observed traits. The inclusion of neighborhood and time effects ensures that identification is driven by within-neighborhood timing variation, not cross-sectional differences in baseline access.

Figure 10: Correlations Between Choice School Distance and Student Characteristics



Notes: This figure reports coefficients from student-level regressions testing whether proximity to choice schools is correlated with baseline characteristics. Each coefficient corresponds to a regression of a given characteristic on relative distance to the nearest choice school, defined as the difference between a block's distance to the nearest choice program and the nearest non-choice neighborhood school. The sample includes all fifth-grade students from 2004-2017. Circle markers correspond to specifications with only academic year fixed effects, while triangle markers include both neighborhood and year fixed effects, isolating within-neighborhood changes in distance. We report 95-percent confidence intervals where standard errors are robust and clustered at the neighborhood level.

7 Model Estimates and Causal Effects

In this section, we first report estimates of the structural model of application and enrollment demand. We then summarize causal effects for the entire population of LAUSD students, uncovering some of the allocative inefficiencies of opt-in systems. Last, we conclude with an analysis of the alignment of school-specific demand and school-specific causal effects.

7.1 Model Fit

We estimate three variants of our model that vary in the number of components K in the mixture of normals used to approximate the distribution of unobserved preference heterogeneity. Appendix Table A.5 reports model fit statistics and shows that the log-likelihood is maximized with a three-type model ($K = 3$), although the incremental gain relative to the two-type specification is modest. Any further improvement in fit beyond the three-type model is also modest, so we adopt the three-type model as our preferred specification. The model with two-types nonetheless represents an important improvement in fit relative to the single-type case, leading us to reject the simple specification. We report estimates for the two-type mixture model in Appendix Tables A.6 and A.7.

We next conduct an out-of-sample validation to assess the overall fit of our preferred model

that uses a three-type mixture distribution. Specifically, we estimate the model in a 25% random training sample and assess the fit using the remaining 75%. We run 200 simulations in which we use parameter estimates from the training sample to simulate applications, offers, and enrollment in the holdout sample. We then compare simulated averages to actual data across a range of moments that capture school-level and group-by-school-level behavior.

Appendix Figure A.6 compares school-level aggregates (maroon triangle markers) and group-by-school level aggregates (gray circle markers). Panel (a) shows visually that the model is able to match school-specific application demand closely. For applications, the model forecasts school-specific demand accurately on average with a forecast coefficient equal to 0.99 (std. err. = 0.043). Panel (b) shows similarly strong performance for enrollment, with a forecast coefficient equal to 0.97 (std. err. = 0.034). For both stages, we find that the model predictions of application and enrollment behavior are forecast unbiased. Overall, the model is able to reproduce aggregate and group-specific moments in the holdout sample, providing reassuring evidence for its ability to capture important determinants governing selection into application and enrollment.

7.2 Demand Model Estimates

Table 2 summarizes our demand model estimates. Panel (a) reports estimates of the utility and application cost parameters from our preferred mixture model. Column 1 shows estimates of the utility parameters δ_j and β , while Column 2 presents estimates of the distance cost parameters ψ_D , ψ_{D2} , and ψ_{Dx} . In the top row of Column 1, the constant (main effect) is the average of school utility intercepts (i.e., δ_j). The estimated constant is negative and indicates that the average student prefers their neighborhood school to a typical choice school, even in the absence of application or distance costs. Consistent with prior evidence on the importance of travel distance for school choice (e.g., Hastings et al. 2009; Laverde 2024), the estimated distance cost parameter is statistically significant, confirming that distance exerts a strong deterrent effect on demand. Based on these estimates, the implied mean utility distance for a typical choice school is equivalent to traveling approximately 2.75 miles ($0.897/0.325$). Column 3 reports estimates of the natural logarithm of fixed costs (i.e., μ_c and μ_{cx}). The estimate for the main effect indicates that application costs are substantially larger than travel costs.

Although the average utility of a typical choice school is negative, the standard deviation is about 3.2 miles, indicating that popularity varies widely. Some schools are so attractive that, absent travel or application costs, virtually all students would apply. In terms of observable heterogeneity, the estimates in Column 1 indicate that Black and Hispanic students express weaker preferences for choice schools than Asian students. Limited English proficiency students also show lower demand. Consistent with Table 1, students with higher incoming achievement both value choice schools more and face lower application costs, helping explain the strong sorting of higher-achieving students into these schools.

Heterogeneity in the distance cost parameters is modest, but application costs vary more. Relative to Asian students, the results in Column 3 indicate that Hispanic students face larger application costs, while Black and White students face lower application costs. Application costs are lower for students with higher baseline achievement. Students enrolled in choice schools at

Table 2: Demand Estimates: Mixture Model

Panel (a): Estimates for Observable Parameters

	Utility (1)	Distance Cost (2)	Log Cost (3)
Main Effects	-0.897 [1.037]	0.325 (0.005)	0.209 (0.014)
Female	0.002 (0.023)	0.002 (0.003)	0.002 (0.003)
Black	-0.106 (0.057)	-0.016 (0.007)	-0.035 (0.008)
Hispanic	-0.106 (0.039)	0.012 (0.005)	0.065 (0.006)
White	-1.098 (0.045)	-0.131 (0.007)	-0.083 (0.007)
Poverty	0.070 (0.026)	0.015 (0.004)	0.014 (0.004)
LEP	-0.276 (0.040)	-0.016 (0.005)	0.010 (0.005)
Speaks English at Home	-0.137 (0.031)	-0.024 (0.004)	-0.010 (0.004)
Baseline ELA	0.175 (0.018)	0.000 (0.002)	-0.033 (0.003)
Baseline Math	0.202 (0.018)	0.006 (0.002)	-0.018 (0.003)
Neighborhood Median Income	-0.059 (0.005)	0.001 (0.001)	0.005 (0.001)
Baseline Choice Enrollment	1.275 (0.038)	-0.076 (0.004)	-0.097 (0.007)
Baseline Peer Quality	0.123 (0.028)	-0.026 (0.004)	-0.026 (0.005)

Panel (b): Estimates for Unobservable Parameters

	μ (1)	σ (2)	$Pr(K_i = k)$ (3)
Type 1	-0.876 (0.046)	0.189 (0.062)	0.446
Type 2	0.430 (0.042)	0.125 (0.043)	0.488
Type 3	2.749 (1.051)	0.541 (0.039)	0.066
Cost Heterogeneity		0.363 (0.006)	

Notes: This table reports demand model results (Section 6.2) estimated via simulated maximum likelihood with 300 draws for taste heterogeneity (θ_i) and cost heterogeneity (η_i). Panel (a) reports observable heterogeneity: the first row gives main effects for distance and log cost; school mean utilities are shown as averages with noise-adjusted SDs in brackets; remaining rows are heterogeneity terms. Panel (b) reports unobservables: a mixture-of-normals for tastes (means in col. 1, SDs in col. 2, type probabilities in col. 3) and mean-zero normal costs with the SD in col. 2. Standard errors are in parentheses.

the time of application also face lower application costs.

Finally, Panel (b) reports summary statistics for the parameters that govern the distribution

of unobserved preferences. The results show that preference heterogeneity is substantial, while cost heterogeneity is more limited. The estimated means from the three-type mixture model in Column 1 suggest that there is a type that dislikes choice schools, with mean $\mu_1 = -0.876$, another that moderately likes choice schools with $\mu_2 = 0.430$, and a third type that strongly prefers choice schools with $\mu_3 = 2.75$. The type-specific standard deviations in Column 2 range from 0.19 to 0.54, underscoring meaningful within-type heterogeneity in choice school preference. Column 3 indicates that the population splits roughly evenly between Type 1, which tends to dislike choice schools (45 percent), and Type 2, which moderately likes them (49 percent), while Type 3-students with strong preferences for choice schools represent only 7 percent of the population but are over-represented among applicants and enrollees. As for cost heterogeneity, we estimate the standard deviation is 0.36, roughly one-third the size of the observable costs for the average student in the population.

In sum, the estimates show that choice schools are generally less attractive than neighborhood schools, with travel and application costs further discouraging participation. Unobservable preference heterogeneity—likely reflecting differences in information access, parental engagement, and other frictions—also weakens demand, since most students fall into Type 1 (strong distaste for choice schools) or Type 2 (moderate idiosyncratic tastes). Together, these patterns clarify the main barriers shaping participation in LAUSD’s opt-in system: sizable application costs that deter families from pursuing more effective schools, and wide dispersion in tastes that likely arises from information frictions and uneven parental motivation. Regardless of the underlying micro-foundation, the evidence indicates that high barriers to participation limit the allocative efficiency of opt-in systems. The details of how families are allowed to participate matter: without adequate information provision or institutional support to reduce application costs, demand for choice schools—and the system’s ability to deliver their benefits to the students who need them most—will remain constrained.

7.3 Causal Effects on Student Achievement

Table 3 reports the selection-corrected parameter estimates based on the student achievement model specified by Equations (12) and (14).¹⁹ Columns 1-4 report results for the parameterization that allows for type-specific heterogeneity, while Columns 5-8 correspond to the linear control function results. We report estimates for math and ELA achievement separately.

¹⁹Appendix Table A.8 reports bootstrapped estimates that take into account estimation error in the posterior means and type probabilities that we use as control functions. Results based on our bootstrap exercise are quantitatively similar as our main results.

Table 3: Student Achievement Model Estimates

	Model 1: Type Probabilities				Model 2: Linear Control Function			
	Math		ELA		Math		ELA	
	Neighborhood School	Choice School Diff.	Neighborhood School	Choice School Diff.	Neighborhood School	Choice School Diff.	Neighborhood School	Choice School Diff.
	(γ_0) (1)	(γ_c) (2)	(γ_0) (3)	(γ_c) (4)	(γ_0) (5)	(γ_c) (6)	(γ_0) (7)	(γ_c) (8)
Main Effects		0.290 [0.160]		0.243 [0.113]		0.252 [0.158]		0.226 [0.111]
Female	-0.022 (0.002)	-0.009 (0.007)	0.111 (0.002)	0.002 (0.006)	-0.022 (0.002)	-0.010 (0.007)	0.111 (0.002)	0.001 (0.006)
Black	-0.196 (0.011)	0.029 (0.022)	-0.161 (0.010)	0.027 (0.020)	-0.197 (0.011)	0.028 (0.022)	-0.161 (0.010)	0.027 (0.020)
White	-0.004 (0.012)	-0.027 (0.021)	0.001 (0.009)	-0.027 (0.021)	-0.004 (0.012)	-0.033 (0.021)	0.001 (0.009)	-0.030 (0.021)
Hispanic	-0.149 (0.010)	-0.004 (0.020)	-0.119 (0.008)	0.012 (0.019)	-0.150 (0.010)	-0.009 (0.020)	-0.119 (0.008)	0.010 (0.019)
Asian	0.211 (0.014)	-0.062 (0.025)	0.075 (0.010)	-0.042 (0.022)	0.210 (0.014)	-0.059 (0.025)	0.075 (0.009)	-0.041 (0.022)
Poverty	0.010 (0.003)	-0.011 (0.010)	0.003 (0.002)	-0.006 (0.007)	0.010 (0.003)	-0.012 (0.010)	0.003 (0.002)	-0.007 (0.007)
LEP	-0.087 (0.004)	0.037 (0.013)	-0.171 (0.003)	0.010 (0.014)	-0.087 (0.004)	0.041 (0.013)	-0.171 (0.003)	0.012 (0.014)
Median Income		-0.005 (0.002)		-0.006 (0.002)		-0.006 (0.002)		-0.006 (0.002)
English Home	-0.070 (0.004)	0.009 (0.012)	-0.059 (0.004)	0.044 (0.010)	-0.070 (0.004)	0.009 (0.012)	-0.059 (0.004)	0.044 (0.010)
Baseline Math	0.526 (0.004)	-0.038 (0.009)	0.182 (0.002)	-0.023 (0.006)	0.526 (0.004)	-0.037 (0.009)	0.182 (0.002)	-0.023 (0.006)
Baseline ELA	0.210 (0.003)	0.012 (0.007)	0.591 (0.003)	-0.016 (0.008)	0.210 (0.003)	0.014 (0.007)	0.591 (0.003)	-0.016 (0.008)
Baseline Peer Quality	-0.149 (0.012)	0.032 (0.018)	-0.095 (0.009)	0.024 (0.014)	-0.147 (0.012)	0.036 (0.019)	-0.094 (0.009)	0.025 (0.014)
Pr(Type 2)	-0.053 (0.032)	-0.167 (0.047)	-0.006 (0.023)	-0.118 (0.040)				
Pr(Type 3)	0.076 (0.015)	-0.385 (0.039)	0.102 (0.012)	-0.309 (0.032)				
Choice School Preference θ_i					0.023 (0.003)	-0.076 (0.009)	0.029 (0.003)	-0.071 (0.007)
Nghd., Year, Sub-district FE	✓		✓		✓		✓	
Observations	334,166		334,166		334,166		334,166	

Notes: This table reports estimates of Equation 12 (Cols 1-4) and 14 (Cols 5-8). The odd columns (1, 3, 5, and 7) report baseline coefficients (γ_0); the even columns (2, 4, 6, and 8) report choice-enrollee interactions (γ_c). The total effect for choice students is $\gamma_0 + \gamma_c$. The “Main effects” row gives the enrollment-weighted average choice effect; noise-adjusted SDs are in brackets. All specs include Census-block, year, and sub-district fixed effects. No main effect for Median Income is shown because it is collinear with block FEs; its interaction is identified from within-block variation in the choice indicator.

Both models paint similar pictures of the average effects of choice schools. The constant (“main effects” row) reports the enrollment-weighted average school-level causal effects, with the estimated standard deviation below in brackets. In the type-specific model, these main effects represent the impact of the average choice schools for Type 1 students—those with the lowest idiosyncratic demand for choice schools. Columns 2 and 4 indicate that, for these students, attending an average choice school boosts test scores by 0.29σ in Math and 0.24σ in ELA, with standard deviations of 0.16σ and 0.11σ across schools, respectively. In the linear control function, the first row shows that choice schools typically boost test scores for the average LAUSD students by 0.25σ in math and 0.23σ in ELA (Columns 6 and 8), with a similar estimated heterogeneity across schools (e.g., standard deviations of 0.160 and 0.158). Taken together, the estimates imply that some choice schools boost scores substantially, while others are roughly comparable to neighborhood schools, and a subset exhibit negative causal effects. These findings indicate that the question of what students access through the choice sector is a vertically differentiated set of schools. They are generally more effective than neighborhood schools, and the evidence suggests that the school district is responding to charter competition by creating vertically differentiated options that families can sort into.

The subsequent rows of Column 2 and Column 4 (Column 6 and 8) report treatment effect heterogeneity estimates in terms of observables and unobservables. All else equal, Black students tend to experience larger gains from enrolling in choice schools, while Asian students experience relatively negative gains. Treatment effects tend to decline with baseline achievement, indicating that higher-achieving students benefit less from enrollment. Overall, these patterns in observable effect heterogeneity, which were not specifically targeted in the estimation, line up well with the reduced form results in Figure A.4. The totality of the evidence reinforces the broader narrative of negative selection on gains: students with characteristics associated with lower application rates (e.g., those with low baseline achievement) tend to experience larger treatment effects.

In terms of idiosyncratic tastes, Column 1 and Column 3 of Table 3 shows that students with larger estimated tastes for choice schools are positively selected—Type 3 students with the largest estimated tastes for choice schools perform 0.08 - 0.10σ better on standardized exams than Type 1 students regardless of the school they enroll in—a finding that is consistent with descriptive evidence in Table 1. Columns 5 and 7 paint a similar picture via the linear dependence of θ_i and achievement. Columns 2 and 4 report a negative association between preferences and causal effects. Type 2 student’s causal effects are 0.12 - 0.17σ lower than Type 1 students, and Type 3 students have causal effects that are $0.31 - 0.38\sigma$ lower than Type 1 students, patterns that are consistent with negative selection on achievement gains. The linear control function approach reveals similar patterns: a one-unit increase in θ_i is associated with a roughly $0.07 - 0.08\sigma$ decrease in choice school causal effects, about one third of the average. The negative selection on achievement gains we find adds to the growing body of evidence of this kind of sorting in education and other settings (Chyn 2018; Cornelissen et al. 2018; Kline and Walters 2016; Walters 2018).

To aid interpretation of these selection patterns, Appendix Figure A.7 summarizes choice-school preferences for the subset of students who appear in the lottery sample. We begin by taking the posterior estimates of θ_i for the full population, ranking these values, and constructing percentile measures. We repeat this procedure for an overall preference index, which is defined

as $I_i \equiv \beta_c X_i + \theta_i$. We then attach these percentile ranks to the lottery applicants to show where they lie in each distribution. Panel (a) plots percentiles of θ_i itself, and Panel (b) plots percentiles of the combined preference index. Across both measures, lottery applicants are sharply concentrated in the far right tail: very few fall below the 80th percentile, and a substantial mass lies in the top decile. This pattern places the lottery sample almost entirely among students with the strongest idiosyncratic tastes for choice schools—those most likely to be Type 3 and, as the achievement estimates above show, the group with the smallest (and sometimes negative) causal effects. This figure, therefore, highlights a central allocative inefficiency of the opt-in system: the students most inclined to apply are precisely those who stand to benefit the least from attending choice schools. We next use a more direct visualization of type-specific causal effects to summarize these reverse-Roy patterns.

Finally, Figure 11 provides a complementary, visual summary of the relationship between enrollment preferences and causal effects by plotting Type-specific average causal effects. Recall that demand for choice schools is increasing with the type index. We see that Type 1 and 2 students, who have weak or moderate tastes for choice schools, have positive and relatively large-in-magnitude treatment effects from choice school enrollment. In contrast, Type 3 students—those with the strongest choice school tastes—have negative effects. Taken together with Appendix Figure A.7, which shows that lottery applicants are drawn disproportionately from the high-preference (likely Type 3) tail of the distribution, this figure makes the reverse-Roy pattern transparent. In summary, our analysis of causal effects reveals two takeaways: (i) choice schools are vertically differentiated, likely reflecting a competitive response of school districts in response to growing competition, and (ii) there is stark heterogeneity in treatment effects—a finding that alludes to allocative inefficiency in terms of district-level achievement maximization.


7.4 Demand for School Effectiveness

The preceding analyses show that those students who exhibit strongest demand for choice schools also tend to realize relatively worse achievement gains. This pattern points to an allocative inefficiency of the opt-in design, where willingness to participate is not tightly aligned with school effectiveness. We next examine this relationship more directly by exploring how demand varies with school-level causal effects.

Are more effective schools relatively more popular? Figure 12 provides a visualization that suggests that demand is not higher for high-achieving schools. Specifically, the figure reports the bivariate relationship between school-specific average causal effects (α_j) and demand scaled in miles (based on δ_j and the distance cost parameters). If anything, the relationship between school effectiveness and school popularity is negative. The misalignment of demand and school effectiveness adds to the mounting evidence finding weak relationships on this margin (Abdulkadiroğlu et al. 2020; Ainsworth et al. 2023; Rothstein 2006).

The misalignment of demand with school effectiveness could be due to a host of reasons. A preference for peers is often advanced as an explanation for these kinds of empirical findings (Rothstein 2006). At the same time, it is unclear if peer quality acts as a proxy for school quality in environments with substantial information frictions. In the same spirit that lack of

Figure 11: Choice School Causal Effects by Preference Type



Notes: This figure plots causal achievement effects for the three types implied by our demand model outlined in Section 6.2. Table 3 reports the parameter estimates we used to construct these type-specific estimates. We report 95 percent confidence intervals where standard errors are robust and clustered at the school level, and use the delta method where appropriate. We report effects for Math and ELA in black and maroon, respectively, using 334,166 students.


information inhibits applying to choice schools, lack of information about school effectiveness dampens potential demand for effective schools by either a lack of awareness or inducing families to select schools based on other attributes that are weakly correlated with school effectiveness. Information interventions that target this information friction have been successful at boosting demand (Ainsworth et al. 2023; Campos 2024), but the multifaceted nature of demand will always limit the extent to which these interventions can completely align demand with school quality (Beuermann et al. 2023).

7.5 Robustness and Additional Validation

The causal effects we estimate rely on our measures of latent demand adequately summarizing selection into application and enrollment. Appendix Figure A.6 demonstrated that our demand model reproduces a variety of moments in a holdout sample. This provides reassuring evidence that the model replicates key selection patterns we aim to account for with our counterfactual analysis. We now discuss two additional exercises that bolster the strength of our claim.

First, we address potential identification concerns. As outlined in Section 6.3, we rely on lottery-based offers and policy-induced variation in access to estimate our model of enrollment and achievement. Both instruments are excluded in the student achievement outcome equation, but one may worry that our effect heterogeneity is conflated by potential distance-based heterogeneity. Appendix Table A.9 reports estimates of a model that accounts for distance-

Figure 12: Demand for Effective Choice Schools



Notes: This figure reports the relationship of estimated school effectiveness and demand, where demand is in miles of willingness to travel. School effectiveness is estimated in a model that adjusts for selection into application and enrollment using the posterior mean control function estimates from the estimated demand model.

based heterogeneity in treatment effects and results are essentially unchanged. Note that this model still leverages the same identifying variation as results in Table 3, but assesses whether distance-based treatment effect heterogeneity explains away the θ -based heterogeneity in causal effects. This robustness check tests for direct effects of distance and finds them negligible, with our core findings robust to a distance-based gradient. The relatively weak distance gradient also suggests that the lottery-based distance heterogeneity reported in Figure 6 was indeed reflective of unobservable preference-based heterogeneity.

Second, we conduct a more formal validation exercise for our outcome model. One assessment of the model's reliability is to implement a lottery-based "forecast" test following Angrist et al. (2017). We do this by using randomized offer variation to check whether changes in *predicted* student outcomes, denoted $\tilde{Y}_i$, move one-for-one with lottery-induced changes in observed outcomes Y_i . Concretely, we construct a model-based prediction $\tilde{Y}_i$ for each student's achievement using our model from Section 6.3 and the school that each student attends. Because oversubscribed lotteries shift enrollment (thereby shifting $\tilde{Y}_i$), we can use the randomized offers Z_i as an instrument for the model-based predicted outcomes. We can then estimate a 2SLS model where the second-stage focuses on the relationship between Y_i and $\tilde{Y}_i$ after conditioning on the appropriate lottery-strata fixed effects. As our test, we focus on whether the 2SLS estimate for the coefficient on $\tilde{Y}_i$ is equal one, since a value of one indicates the model's predictions are unbiased for the causal effects revealed by the lotteries. We also summarize

this with a visual IV plot: for each lottery, we plot the first-stage offer effect on $\tilde{Y}_i$ against the associated reduced form offer effect on Y_i .

Appendix Figure A.8 reports the visual IV plot. The plot shows a tight, approximately linear relationship between each lottery’s first stage and the reduced form impact. The estimated forecast coefficient is 0.91, and we are unable to statistically reject it equals one. Notably, the results also show relatively modest dispersion around the estimated slope. This evidence suggests that our outcome model predicts treatment effects that are commensurate to observed lottery-based treatment effects. Equipped with validated causal effects, we now turn to an assessment of district-level policies and their impacts on aggregate student achievement.

8 Counterfactual Policy Analysis

Our findings to this point suggest that participation in LAUSD’s choice system—an opt-in structure—remains limited and uneven, with substantial heterogeneity in both who applies and how much schools contribute to student achievement. These patterns raise the possibility that alternative policies and system design could reshape the distribution of achievement by altering how families access and engage with choice programs. To explore these possibilities, we use the estimated structural model to simulate policy counterfactuals by adjusting key parameters that capture how different features of system design affect application behavior, enrollment behavior and achievement outcomes.

Our counterfactual analysis evaluates participation-targeted policies and how application and enrollment system design influences shape sorting and achievement outcomes. LAUSD’s current opt-in structure serves as the baseline simulation, providing a benchmark for comparisons with scenarios that incorporate either targeted policy interventions or system-level reforms. Our analysis begins by focusing on two participation-focused policies that can operate within any public choice framework and address key barriers to entry: (i) information interventions that reduce search frictions and improve families’ understanding of available options, modeled as increases in θ_i , the parameter capturing access to information; and (ii) reductions in travel costs—for example, through universal busing programs—implemented as downward adjustments to the distance cost parameters. Next, we examine two system-level reforms: (iii) mandatory participation as implemented in New York City and Boston, modeled by setting application cost parameters to zero; and (iv) decentralized choice markets, in which families may apply to multiple schools and, in the absence of a centralized clearinghouse, can receive multiple offers. Collectively, our simulations capture the dominant system designs adopted by nearly all large districts—roughly 99 percent, as shown in Figure 2—and incorporate participation-enhancing policies, such as busing and information interventions, that many districts have implemented.²⁰

²⁰Given that our counterfactuals involve realigning students to schools, it is natural to wonder whether incorporating peer effects into the model could change key conclusions. To explore this possibility, we plot reduced form school-specific achievement effects against their lottery induced impacts on peer quality (Appendix Figure A.9). If peer quality is an important driver of student outcomes, we would expect a strong correlation between these two variables. Instead, we find essentially no relationship between changes in peer quality and own-achievement effects. This suggests that incorporating peer effects is unlikely to undermine the key conclusions of the counterfactual analysis.

On the supply side, we hold fixed the set of choice schools, their locations, and their school-level utilities (δ_j) and achievement effects (α_j). We also fix each program’s race-specific seat capacities and run admissions within program-year-grade-priority points-race strata, mirroring LAUSD’s lottery practice and the district’s court-ordered desegregation regime.²¹ Maintaining these race-specific seat targets in all counterfactuals ensures that any changes we report come from family decision-making (who applies and enrolls) rather than from relaxing binding integration mandates. Admissions continue to follow the same lottery rules and capacity constraints as in the baseline LAUSD system, but each policy counterfactual induces changes in demand that we must account for when students decide to apply to schools. For each relevant policy counterfactual, we calculate equilibrium admission probabilities that schools set under the demand structure implied by the policy change. Our assumptions ensure that the counterfactuals isolate how changes in family decision-making—rather than school expansion or reorganization—affect participation and achievement. As previewed, we compare simulated outcomes under each policy to the baseline opt-in allocation to assess how alternative designs would reshape access, participation, and aggregate achievement. Appendix Section E provides further details on the counterfactual exercises.

Before discussing policies that substantially alter system design, we hold constant the status-quo opt-in design and focus on policies that explicitly target access with informational interventions. These exercises examine how improving families’ access to information affects school choice decisions and, in turn, achievement outcomes. The findings outlined in Table 3 suggest that inducing students with lower θ_i to apply will lead to increases in system-wide learning. Specifically, we simulate two experiments in our model designed to approximate the types of information provision strategies commonly studied in the literature (Agte et al. 2024; Ainsworth et al. 2023; Andrabi et al. 2017; Hastings and Weinstein 2008). Both simulations use an adjusted preference parameter $\theta'_i = \theta_i + \Delta_\theta$ in the application and enrollment decisions of the model, while preserving the original θ_i when calculating student achievement.


In the initial general information experiment, we set Δ_θ to a type-specific one standard deviation (reported in Column 2 of Table 2) for a randomly selected subset of students, capturing imperfect compliance and heterogeneity in take-up. In the next, school-targeted experiment, we focus only on students enrolled in low-achieving schools, defined as those with average achievement below the district average, and apply the same Δ_θ increase to θ'_i for a random subset within this group to mimic targeted outreach. Together, these scenarios illustrate how information provision could shift participation and sorting in LAUSD’s choice system by narrowing informational gaps across families.

Figure 13 summarizes the gains from the information interventions relative to LAUSD’s existing opt-in policy, showing two bars for Math and ELA outcomes, respectively. The first two sets of bars on the left show that both information interventions produce extremely modest achievement gains, not exceeding half of a percent of a standard deviation in district-level achievement. Table 4 provides additional insights, demonstrating application rates increase from 16 percent to roughly 20 percent, and the number of occupied seats increases by 11-

²¹As noted in Section 3, choice programs have used racial enrollment quotas that have historically reserved a large share of seats for students from PHBAO (Predominantly Hispanic, Black, Asian, and Other non-Anglo) neighborhoods; these requirements persisted after court supervision ended and were upheld as lawful given their pre-1996 remedial origin.

12 percentage points from a base of 55.3 percent. By demographic group, Asian and White students have the largest enrollment responses in percentage-point terms to both information campaigns, as has been shown in prior reduced form work (Corcoran et al. 2018). These skewed take-up rates produce minimal impacts on district-level achievement, given that the affected students induced to participate (particularly Asian) have relatively weaker achievement gains (see Table 3 and Appendix Figure A.10). Overall, information campaigns have scope to produce reallocations of students between schools as has been shown in prior work (Agte et al. 2024; Campos 2024; Hastings and Weinstein 2008), but are not enough to overcome substantial barriers to participation that are present in opt-in systems.

Figure 13: Estimated Effects of Counterfactual Policies on District-Level Achievement



Notes: This figure reports mean changes in district-level standardized achievement for Math (grey) and ELA (maroon) relative to a baseline simulation of LAUSD’s current opt-in system. We consider participation-targeted policies: General Information, which raises the information parameter θ_i for a randomly selected 50% of students; Targeted Information, which raises θ_i for a randomly selected 50% of students enrolled in schools whose average achievement is below median; and Busing, which sets distance costs to zero. We also consider two system-level designs: Decentralized choice, which allows families to submit multiple applications and hold simultaneous offers while schools run independent lotteries, and Mandatory application, which sets application costs to zero, requires a rank-ordered list (which may include the neighborhood school), and assigns seats via deferred acceptance. We also report combinations of these components (Mandatory + Busing; Mandatory + Information; Mandatory + Information + Busing) and an Achievement-Optimal benchmark that assigns students in order of modeled match quality. Each scenario is simulated 100 times and we report average effects; for cases with school-run lotteries, we recompute equilibrium best-response admission probabilities under policy-induced demand (see Appendix Section E).

Next, we evaluate the impact of creating expanded busing policies as a means of addressing travel barriers limiting participation. Busing policies have a long history of broadening access

to schools, particularly in connection with court-ordered desegregation efforts.²² While existing evidence on busing is mixed—Angrist et al. (2022) find minimal achievement impacts from traveling longer distances to schools with higher achieving peers, Setren (2024) reports substantially improved long-run outcomes in the Boston’s METCO program, and Bergman (2018) detect gains (and some negative externalities) in a busing program for the Bay Area—we bring these ideas into our structural framework. Because we estimate sizable travel costs in the context of our model of LAUSD’s opt-in system, a policy that would improve busing availability or its salience may induce more participation from disadvantaged groups that are least likely to participate.

The results presented third from the left in Figure 13 show that implementing a busing program that removed travel costs would raise district-level achievement by roughly one percent of a standard deviation, with more pronounced effects for white students.²³ Under this regime, Table 4 shows that 33 percent of families participate and 86 percent of seats are occupied, a substantial improvement from the baseline opt-in system. Consistent with reductions in travel costs drawing in relatively disadvantaged students, the effects on the set of enrolled students (i.e., the treatment on the treated) increase by 4 percent of a standard deviation for both Math and ELA. Although these estimates provide an encouraging benchmark for the achievement gains, a key caveat for policy is that our simulation is silent about the costs imposed by a regime that would reduce travel costs by such a substantial margin.

Because access-targeting in the opt-in context still leaves substantial non-participation, we now assess two system-level alternatives. We begin by considering a decentralized system that lets families submit multiple applications and allows schools to run independent lotteries, yielding multiple simultaneous offers. To implement this policy, we introduce a marginal per-application cost to discipline portfolio size in the model. The ability to submit multiple applications has the potential to thicken the market and reduce slack, but portfolio expansion creates congestion that can discourage marginal entrants. Next, we consider a mandatory system that requires every student to submit a rank-ordered list and coordinates offers via a deferred acceptance clearinghouse with uncapped list lengths. In this counterfactual, we set application costs to zero, while families still weigh preferences and travel costs when ranking. By removing application frictions and eliminating multi-offer congestion, the system increases participation, fills capacity, and reallocates students toward higher-gain seats. Prior work finds that transitions to centralized assignment can improve welfare (Abdulkadiroğlu et al. 2017). Together, these counterfactuals highlight the core trade-off: decentralization buys thickness at the cost of congestion and selective participation, whereas mandatory centralization trades administrative reach for broad access and coordinated assignment.

²²In Los Angeles, busing was once part of court-ordered integration efforts, and today the district operates a voluntary busing program that facilitates access to magnet programs (Blume 2019).

²³See Appendix Figure A.10 for race-specific results for each counterfactual considered in this section.

Table 4: Analysis of Participation Reforms and and Alternative Choice Systems

Panel (a): Baseline Simulation

	Participation (%)		All Student Achievement (σ)		Choice Enrollee TOT (σ)		Demographic Group Choice Enrollment (%)				
	Apply (1)	Choice Seats Filled (2)	Math (3)	ELA (4)	Math (5)	ELA (6)	Asian (7)	Black (8)	Hispanic (9)	White (10)	Low SES (11)
Opt-in (Voluntary)	15.72	55.33	.	.	0.111	0.078	17.47	7.68	5.01	14.36	5.78

Panel (b): Alternative Policies/System Impacts (Relative to Baseline)

	Participation Gains (%)		All Student Achievement Gains (σ)		Choice Enrollee TOT Gains (σ)		Demographic Group Choice Enrollment Gains (%)				
	Apply (1)	Choice Seats Filled (2)	Math (3)	ELA (4)	Math (5)	ELA (6)	Asian (7)	Black (8)	Hispanic (9)	White (10)	Low SES (11)
Broad Information	+ 3.84	+ 11.43	+ 0.003	+ 0.002	+ 0.015	+ 0.014	+ 2.21	+ 2.19	+ 1.15	+ 2.05	+ 1.27
Targeted Information	+ 3.90	+ 11.85	+ 0.003	+ 0.002	+ 0.012	+ 0.012	+ 1.54	+ 2.31	+ 1.29	+ 1.67	+ 1.40
Busing Program	+ 17.45	+ 30.33	+ 0.009	+ 0.007	+ 0.042	+ 0.040	+ 1.56	+ 3.35	+ 3.26	+ 9.35	+ 3.83
Decentralized	-5.28	+ 18.53	+ 0.003	+ 0.002	+ 0.001	-0.001	+ 7.09	+ 4.01	+ 1.45	+ 4.72	+ 1.87
Mandatory App.	+ 69.12	+ 42.49	+ 0.018	+ 0.016	+ 0.105	+ 0.103	-14.53	+ 3.06	+ 4.94	+ 15.74	+ 5.33
Mandatory App. + Busing	+ 83.07	+ 44.39	+ 0.021	+ 0.019	+ 0.123	+ 0.121	-15.00	+ 2.29	+ 5.21	+ 18.41	+ 5.61
Mandatory App. + Broad Info	+ 70.00	+ 42.76	+ 0.018	+ 0.016	+ 0.106	+ 0.104	-14.55	+ 2.99	+ 4.97	+ 16.05	+ 5.35
Mandatory App. + Broad Info + Busing	+ 83.17	+ 44.41	+ 0.021	+ 0.019	+ 0.124	+ 0.121	-15.01	+ 2.27	+ 5.22	+ 18.43	+ 5.61
Maximized Achievement	-	+ 44.67	+ 0.038	+ 0.034	+ 0.258	+ 0.245	-16.67	+ 2.99	+ 5.88	+ 18.83	+ 5.41

Notes: This table reports simulated results from our model of application, enrollment, and achievement. Panel (a) reports results for the baseline opt-in system. Panel (b) reports counterfactual impacts, expressed as the difference relative to the baseline scenario. We report the percent of students that apply to choice schools as their most-preferred option (Column 1), and the percent of choice school seats that are filled (Column 2). For achievement, we report district-wide achievement (Columns 3 and 4) and treatment on the treated estimates (Columns 5 and 6). Demographic outcomes include the percent of students within each racial or income-group enrolling in a choice school (Columns 7-11). The General Information scenario increases the θ_i of a randomly selected 50% of students in the district, while the Targeted Information scenario increases the θ_i of a randomly selected 50% of students who are enrolled in schools whose average achievement is below median. The No Travel Costs scenario mirrors a generous busing policy that sets travel costs to zero. The Decentralized scenario allows families to submit multiple independent applications to schools, which then run their own lotteries, potentially generating multiple simultaneous offers. The Mandatory scenario eliminates application costs, allowing families to rank their neighborhood school as their most-preferred option and assignments are done via a deferred acceptance mechanism. For each counterfactual, we simulate the economy 100 times. For scenarios that retain school-run lotteries (information, busing, and decentralized choice), we calculate equilibrium best response admissions probabilities that schools would set under the changed demand implied by the policy (see Appendix Section E for details).

The results in Figure 13 and Table 4 show that the decentralized system yields only modest gains in district achievement—comparable in size to information interventions—despite having larger impacts on choice school seat utilization. The mechanism is straightforward: as families expand their application portfolios, equilibrium admission probabilities fall, which induces some marginal families to stop applying and overall participation drops by about 5 percentage points. Among those who remain, most still secure at least one offer, so capacity is used more fully relative to the opt-in baseline. The composition of applicants shifts toward families with lower application costs, producing a more advantaged pool and smaller treatment-on-the-treated (TOT) effects for ELA scores, even as district averages tick up slightly. In short, decentralization discourages entry overall and re-sorts seats toward cost-insensitive families, yielding small district gains.

Our next set of system-level results shows that a transition to a mandatory system with a deferred acceptance mechanism for assignment raises achievement levels by roughly 0.02σ in both Math and ELA. The TOT effects for choice school enrollees also increase substantially by about 0.10σ . These impacts are driven by a 69.1 percentage point increase in the share of all families applying to a non-neighborhood school as their most-preferred option (recall that “apply” is defined as ranking a non-neighborhood school first; families who comply with the mandate by listing only their neighborhood school are counted as non-applicants). This evidence suggests that application costs are the empirically dominant barrier to participation under opt-in systems. As a result of its impacts on participation, the mandatory deferred acceptance system eliminates essentially all of the excess capacity among choice schools in the district, with 98 percent occupied seats compared to only 55 percent at baseline. Taken at face value, the transition to mandatory participation reallocates students across schools, reaches students who stand to gain more in terms of causal effects, and thereby raises district-level achievement.

Having demonstrated that a mandatory system with deferred acceptance produces the largest system-level gains, our next set of simulations examines whether combinations of policies can further expand these gains. To explore this possibility, we augment the mandatory participation regime with the generous busing policy that eliminates travel costs. In this environment, 99 percent of families apply to a choice school, an important increase from the mandatory participation regime without the stronger busing regime. Although the bump in demand is sizable, this translates into only a modest change in the number of seats filled, given that capacity constraints are essentially binding in the case of mandatory participation. Consequently, district-wide and TOT achievement gains are marginally changed relative to the mandatory-only case. Simulations that pair the information campaign with mandatory participation (with or without busing) deliver similarly small incremental gains, reflecting the same capacity constraint.

Finally, the simulations above show that several implementable interventions raise achievement, which naturally prompts the question: can a district do even better? Put differently, what is the maximum achievement attainable if a planner had full information and the authority to match students to schools? To benchmark our results against this best case, we compute an “achievement-optimal” allocation that assigns students solely to maximize district-wide achievement. The procedure we follow has been coined the treatment effect maximizing allocation (TEMA) by Abdulkadiroğlu et al. (2025). As first proposed by Shapley and Shubik

(1971), this is the solution of a linear program that maximizes achievement subject to district imposed capacity constraints. This “achievement-optimal” allocation delivers average gains of 0.038σ in Math and 0.034σ in ELA at the district level, fills all remaining choice capacity (moving from 55% filled at baseline to nearly 100%), and raises treatment-on-the-treated effects to 0.37σ (Math) and 0.32σ (ELA). Therefore, the transition to mandatory participation generates about half of the gains of the infeasible achievement optimal policy. As has been documented in a host of settings, demand for school effectiveness will always be a limiting factor (Abdulkadiroğlu et al. 2020; Ainsworth et al. 2023; Beuermann et al. 2023; Rothstein 2006).

Why does mandatory participation produce the largest benefits relative to opt-in or other policies? Viewed from the perspective of questions on the “Who, What and How” of markets, our education setting with a baseline opt-in system is subject to several allocative inefficiencies in terms of achievement maximization. To begin, the evidence suggests that the “what” is a vertically differentiated option, so there are potential achievement gains to be distributed to students. The key question, then, is “who” accesses those gains. In Los Angeles and nationally, opt-in systems disproportionately attract academically and socioeconomically advantaged families, students who potentially stand to gain the least from access to more effective schools. Turning to the “how,” our analysis of application and enrollment patterns linked to causal effects confirms this: the families in the LAUSD who do opt in tend to experience smaller achievement gains than the students who do not. Once we layer on additional details of the “how”—high application costs, weak information environments, and barriers to participation—the result is a system where advantaged students secure the bulk of available seats. A second inefficiency follows from the same barriers: many high-quality schools are left with unfilled seats. Thus, opt-in systems simultaneously allow advantaged students to hoard scarce opportunities while leaving capacity at effective schools unused, both outcomes that mandatory participation corrects. The elimination of application barriers by de-facto mandatory participation broadens access and reaches a more representative set of students, all the while ensuring all seats in higher quality programs are filled.

9 Conclusion

Our analysis of LAUSD’s opt-in system demonstrates both the promise and the limits of expanding school choice in public school districts. We provide evidence that public choice schools are vertically differentiated and, on average, deliver meaningful achievement impacts. Yet, market design details—optional participation in particular—ultimately blunts the potential benefits. Because application costs deter many students, especially those who would gain the most from attendance, effective schools often operate below capacity. At the same time, the students who do navigate the application process are disproportionately advantaged. These twin features of the opt-in system contribute to both inequality in educational outcomes and allocative inefficiencies.

The central finding from our counterfactual analysis is that centralized, mandatory assignment can deliver substantial gains. Such reforms reallocate students to more effective schools, fill unused seats, and erode negative selection on gains. Centralization generates achievement improvements roughly half as large as an achievement-maximizing allocation, underscoring the

limiting influence of misaligned demand for school effectiveness. However, unlike lighter-touch interventions that provide information or reduce travel costs via generous busing schemes, centralization directly addresses the dominant frictions of opt-in systems.

Taken together, our results demonstrate that centralization is not merely a matter of administrative convenience, but a core design feature that strongly mediates school choice effectiveness. More broadly, our findings highlight how institutional details shape the distribution of educational opportunity: expanding families' options marks progress, but whether these options reach the students who stand to gain the most depends on how assignment mechanisms structure participation. Our findings also speak to the potential implications of recent state-level voucher programs, which have expanded opt-in models and may produce the same inequities and inefficiencies we document, with an unclear picture of their vertical quality. Overall, our findings demonstrate the scope of policies that keep supply of effective schools fixed, but can improve outcomes by system design that can lead to more effective allocations.

One question that we cannot answer is why opt-in systems remain so common in large U.S. school districts. A political economy view is that reforms to assignment rules face organized resistance from powerful incumbents and mobilized parent groups, making overhauls costly to attempt and easy to derail—a point developed in the literature on political constraints in K-12 governance (e.g., Chubb and Moe 1988).²⁴ A second, more market-based explanation is that opt-in designs help districts retain families most likely to exit to charters or private schools—typically higher-achieving students seeking high-achieving peers—by letting them self-select into choice without mandating system-wide participation. This incentive is consistent with ongoing debates about cream-skimming and the distributional consequences of decentralized or unified enrollment systems (Kho et al. 2022). These are, of course, hypotheses. Future work could aim to better understand the relative importance of political versus economic constraints. If opt-in is the binding political equilibrium, the central task becomes identifying choice designs that deliver system-wide gains despite that limitation.

²⁴Recent episodes illustrate how assignment-rule changes can trigger backlash: Wake County's diversity-driven reassignment provoked voter mobilization and a board takeover in 2009 (WRAL News 2009), and San Francisco's shift away from merit-based admissions at a flagship school prompted sustained controversy (The San Francisco Standard 2022).

References

- Abdulkadiroğlu, Atila, , Parag Pathak, and Christopher Walters**, “Who Gets What May Not Matter: Understanding School Match Effects,” *Working Paper*, 2025.
- **and Tayfun Sönmez**, “School choice: A mechanism design approach,” *American economic review*, 2003, *93* (3), 729–747.
- **, Joshua D Angrist, Susan M Dynarski, Thomas J Kane, and Parag A Pathak**, “Accountability and flexibility in public schools: Evidence from Boston’s charters and pilots,” *The Quarterly Journal of Economics*, 2011, *126* (2), 699–748.
- **, Nikhil Agarwal, and Parag A Pathak**, “The welfare effects of coordinated assignment: Evidence from the New York City high school match,” *American Economic Review*, 2017, *107* (12), 3635–3689.
- Abdulkadiroglu, Atila, Parag A Pathak, Alvin E Roth, and Tayfun Sönmez**, “Changing the Boston school choice mechanism,” 2006.
- Abdulkadiroğlu, Atila, Parag A Pathak, Jonathan Schellenberg, and Christopher R Walters**, “Do parents value school effectiveness?,” *American Economic Review*, 2020, *110* (5), 1502–1539.
- Agarwal, Nikhil and Paulo Somaini**, “Demand analysis using strategic reports: An application to a school choice mechanism,” *Econometrica*, 2018, *86* (2), 391–444.
- **, Charles Hodgson, and Paulo Somaini**, “Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys,” *Econometrica*, 2025, *93* (2), 395–438.
- Agte, Patrick, Claudia Allende, Adam Kapor, Christopher Neilson, and Fernando Ochoa**, “Search and biased beliefs in education markets,” Technical Report, National Bureau of Economic Research 2024.
- Ainsworth, Robert, Rajeev Dehejia, Cristian Pop-Eleches, and Miguel Urquiola**, “Why do households leave school value added on the table? The roles of information and preferences,” *American Economic Review*, 2023, *113* (4), 1049–1082.
- American Civil Rights Foundation v. Los Angeles Unified School District**, 169 Cal. App. 4th 436 (Cal. Ct. App. 2nd Dist., 5th Div.). Case No. B205943 2008.
- Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja**, “Report cards: The impact of providing school and child test scores on educational markets,” *American Economic Review*, 2017, *107* (6), 1535–1563.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro**, “Transparency in structural research,” *Journal of Business & Economic Statistics*, 2020, *38* (4), 711–722.
- Angrist, Joshua D, Parag A Pathak, and Christopher R Walters**, “Explaining charter school effectiveness,” *American Economic Journal: Applied Economics*, 2013, *5* (4), 1–27.
- **, Peter D Hull, Parag A Pathak, and Christopher R Walters**, “Leveraging lotteries for school value-added: Testing and estimation,” *The Quarterly Journal of Economics*, 2017, *132* (2), 871–919.
- Angrist, Joshua, Guthrie Gray-Lobe, Clemence M Idoux, and Parag A Pathak**, “Still worth the trip? school busing effects in boston and new york,” Technical Report, National Bureau of Economic Research 2022.
- Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman**, “Smart matching platforms and heterogeneous beliefs in centralized school choice,” *The Quarterly Journal of Economics*, 2022, *137* (3), 1791–1848.
- Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler**, “Disability benefits, consumption insurance, and household labor supply,” *American Economic Review*, 2019, *109* (7), 2613–2654.

- Avery, Christopher, Geoffrey Kocks, and Parag A Pathak**, “The Algorithm Advantage: Ranked Application Systems Outperform Decentralized and Common Applications in Boston and Beyond,” Technical Report, National Bureau of Economic Research 2025.
- Bau, Natalie**, “Estimating an equilibrium model of horizontal competition in education,” *Journal of Political Economy*, 2022, 130 (7), 1717–1764.
- Bergman, Peter**, “The risks and benefits of school integration for participating students: Evidence from a randomized desegregation program,” Technical Report, IZA Discussion Papers 2018.
- Betts, Julian R., Sami Kitmitto, Jesse Levin, Johannes Bos, and Marian Eaton**, “What Happens When Schools Become Magnet Schools? A Longitudinal Study of Diversity and Achievement,” Technical Report, American Institutes for Research May 2015. Contract ED-04-CO-0025/0013 to the Institute of Education Sciences, U.S. Department of Education.
- Beuermann, Diether W, C Kirabo Jackson, Laia Navarro-Sola, and Francisco Pardo**, “What is a good school, and can parents tell? Evidence on the multidimensionality of school output,” *The Review of Economic Studies*, 2023, 90 (1), 65–101.
- Blume, Howard**, “School busing and race tore L.A. apart in the 1970s. Now, Kamala Harris is reviving debate,” *Los Angeles Times*, June 28 2019. Accessed: 2025-10-05.
- Bruhn, Jesse M, Christopher Campos, Eric Chyn, and Anh Tran**, “Who Benefits from Remote Schooling? Self-Selection and Match Effects,” Technical Report, National Bureau of Economic Research 2023.
- Calsamiglia, Caterina, Chao Fu, and Maia Güell**, “Structural estimation of a model of school choices: The boston mechanism versus its alternatives,” *Journal of Political Economy*, 2020, 128 (2), 642–680.
- Campos, Christopher**, “Social Interactions, Information, and Preferences for Schools: Experimental Evidence from Los Angeles,” Technical Report, National Bureau of Economic Research 2024.
- **and Caitlin Kearns**, “The Impact of Public School Choice: Evidence from Los Angeles’s Zones of Choice,” *The Quarterly Journal of Economics*, 10 2024, 139 (2), 1051–1093.
 - **and Christopher Neilson**, “Parents’ Beliefs and Valuations of Academic and Socio-Emotional School Effectiveness,” Technical Report 2025.
- Carneiro, Pedro and James J. Heckman**, “The Evidence on Credit Constraints in Post-Secondary Schooling,” *The Economic Journal*, 2002, 112 (482), 989–1018. Published by the Royal Economic Society.
- Chade, Hector and Lones Smith**, “Simultaneous search,” *Econometrica*, 2006, 74 (5), 1293–1307.
- Chubb, John E and Terry M Moe**, “Politics, markets, and the organization of schools,” *American political science review*, 1988, 82 (4), 1065–1087.
- Chyn, Eric**, “Moved to opportunity: The long-run effects of public housing demolition on children,” *American Economic Review*, 2018, 108 (10), 3028–3056.
- Corcoran, Sean P, Jennifer L Jennings, Sarah R Cohodes, and Carolyn Sattin-Bajaj**, “Leveling the playing field for high school choice: Results from a field experiment of informational interventions,” Technical Report, National Bureau of Economic Research 2018.
- Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg**, “Who benefits from universal child care? Estimating marginal returns to early child care attendance,” *Journal of Political Economy*, 2018, 126 (6), 2356–2409.
- Corradini, Viola**, “Information and access in school choice systems: Evidence from new york city,” Technical Report, Technical report, working paper 2023.
- **and Clemence Idoux**, “Overcoming Racial Gaps in School Preferences: The Effect of Peer Diversity on School Choice,” 2025. Unpublished manuscript.

- Crema, Angela**, “School competition and classroom segregation,” *Unpublished Manuscript*, 2022.
- Deming, David J, Justine S Hastings, Thomas J Kane, and Douglas O Staiger**, “School choice, school quality, and postsecondary attainment,” *American Economic Review*, 2014, *104* (3), 991–1013.
- Dobbie, Will and Roland G Fryer Jr**, “Are high-quality schools enough to increase achievement among the poor? Evidence from the Harlem Children’s Zone,” *American Economic Journal: Applied Economics*, 2011, *3* (3), 158–187.
- Fahle, Erin M, Demetra Kalogrides, Benjamin R Shear, Sean F Reardon, and Andrew D Ho**, “Stanford Education Data Archive Technical Documentation Version 5.0 March 2024,” 2024.
- Figlio, David and Cassandra MD Hart**, “Competitive effects of means-tested school vouchers,” *American Economic Journal: Applied Economics*, 2014, *6* (1), 133–156.
- Figlio, David N and Cecilia Elena Rouse**, “Do accountability and voucher threats improve low-performing schools?,” *Journal of Public Economics*, 2006, *90* (1-2), 239–255.
- Fu, Chao**, “Equilibrium tuition, applications, admissions, and enrollment in the college market,” *Journal of Political Economy*, 2014, *122* (2), 225–281.
- Gilraine, Michael, Uros Petronijevic, and John D Singleton**, “Horizontal differentiation and the policy effect of charter schools,” *American Economic Journal: Economic Policy*, 2021, *13* (3), 239–276.
- Hastings, Justine S and Jeffrey M Weinstein**, “Information, school choice, and academic achievement: Evidence from two experiments,” *The Quarterly journal of economics*, 2008, *123* (4), 1373–1414.
- Hastings, Justine, Thomas J Kane, and Douglas O Staiger**, “Heterogeneous preferences and the efficacy of public school choice,” *NBER working paper*, 2009, *2145*, 1–46.
- Heckman, James J**, “Sample selection bias as a specification error,” *Econometrica: Journal of the econometric society*, 1979, pp. 153–161.
- **and Edward Vytlacil**, “Structural equations, treatment effects, and econometric policy evaluation 1,” *Econometrica*, 2005, *73* (3), 669–738.
- Hess, Frederick M., Robert Maranto, and Scott Milliman**, “Small Districts in Big Trouble: How Four Arizona School Systems Responded to Charter Competition,” *Teachers College Record*, 2001, *103* (6), 1102–1124.
- Hill, Paul and Christine Campbell**, “Portfolio School Districts Project,” 2011.
- Hill, Paul Thomas**, *Strife and progress: Portfolio strategies for managing urban schools*, Brookings Institution Press, 2013.
- Hoxby, Caroline M**, “Does competition among public schools benefit students and taxpayers?,” *American Economic Review*, 2000, *90* (5), 1209–1238.
- Hoxby, Caroline Minter**, “School choice and school productivity. Could school choice be a tide that lifts all boats?,” in “The economics of school choice,” University of Chicago Press, 2003, pp. 287–342.
- Hsieh, Chang-Tai and Miguel Urquiola**, “The effects of generalized school choice on achievement and stratification: Evidence from Chile’s voucher program,” *Journal of public Economics*, 2006, *90* (8-9), 1477–1503.
- Idoux, Clemence**, *Integrating new york city schools: The role of admission criteria and family preferences*, MIT Department of Economics, 2022.
- Kahlenberg, Richard D., Halley Potter, and Kimberly Quick**, “A Bold Agenda for School Integration,” Technical Report, The Century Foundation April 2019. Report.

- Kapor, Adam J, Christopher A Neilson, and Seth D Zimmerman**, “Heterogeneous beliefs and school choice mechanisms,” *American Economic Review*, 2020, 110 (5), 1274–1315.
- Kho, Adam, Ron Zimmer, and Andrew McEachin**, “A descriptive analysis of cream skimming and pushout in choice versus traditional public schools,” *Education Finance and Policy*, 2022, 17 (1), 160–187.
- Kline, Patrick and Christopher R Walters**, “Evaluating public programs with close substitutes: The case of Head Start,” *The Quarterly Journal of Economics*, 2016, 131 (4), 1795–1848.
- Kohli, Sonali, Anna M. Phillips, and Howard Blume**, “With magnet and single-gender schools, LAUSD hopes to boost enrollment this school year,” *Los Angeles Times*, 2016.
- Lallinger, Stefan Redding**, “How Charter, Magnet, and Innovative District Schools Can Help Overcome School Segregation,” Technical Report, The Century Foundation June 2024. Report.
- Laverde, Mariana**, “Distance to schools and equal access in school choice systems,” *Available at SSRN 4905690*, 2024.
- Li, Han**, “SF’s school lottery drives parents crazy and it’s about to change,” April 2025. The San Francisco Standard, Published Apr. 08, 2025, 6:00am.
- Lincove, Jane Arnold and Jon Valant**, “The Effects of Unified School Enrollment Systems on New Orleans Schools: Enrollment, Demographics, and Outcomes after the Transition to OneApp. Technical Report.,” *National Center for Research on Education Access and Choice*, 2023.
- Los Angeles Unified School District**, “Enrollment Trends (ASSESS9 Enrollment Trends, Version 1),” Technical Report, Los Angeles Unified School District 2022. Accessed: 2025-10-02.
- Machado, Cecilia and Christiane Szerman**, “Centralized college admissions and student composition,” *Economics of Education Review*, 2021, 85, 102184.
- Marsh, Julie A, Taylor N Allbright, Danica R Brown, Katrina E Bulkley, Katharine O Strunk, and Douglas N Harris**, “The process and politics of educational governance change in New Orleans, Los Angeles, and Denver,” *American Educational Research Journal*, 2021, 58 (1), 107–159.
- Mumma, Kirsten Slungaard**, “The effect of charter school openings on traditional public schools in Massachusetts and North Carolina,” *American Economic Journal: Economic Policy*, 2022, 14 (2), 445–474.
- Munteanu, Andrei**, “School choice, student sorting, and academic performance,” *Review of Economics and Statistics*, 2024, pp. 1–45.
- Neal, Derek and Joseph Root**, “The Provision of Information and Incentives in School Assignment Mechanisms,” Technical Report, National Bureau of Economic Research 2024.
- Neilson, Christopher A**, “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students,” Technical Report 2021.
- Otero, Sebastián, Nano Barahona, and Cauê Dobbin**, “Affirmative action in centralized college admissions systems,” *Unpublished, Working Paper*, 2023.
- Ovidi, Marco**, “Parents Know Better: Sorting on Match Effects in Primary School,” *Review of Economics and Statistics*, 2025, pp. 1–94.
- Porter, Jack and Ping Yu**, “Regression discontinuity designs with unknown discontinuity points: Testing and estimation,” *Journal of Econometrics*, 2015, 189 (1), 132–147.
- Prendergast, Canice**, “How food banks use markets to feed the poor,” *Journal of Economic Perspectives*, 2017, 31 (4), 145–162.

- Roth, Alvin E**, *Who gets what—and why: the new economics of matchmaking and market design*, Houghton Mifflin Harcourt, 2015.
- Rothstein, Jesse M**, “Good principals or good peers? Parental valuation of school characteristics, Tiebout equilibrium, and the incentive effects of competition among jurisdictions,” *American Economic Review*, 2006, *96* (4), 1333–1350.
- Sartain, Lauren and Lisa Barrow**, “The pathway to enrolling in a high-performance high school: Understanding barriers to access,” *Education Finance and Policy*, 2022, *17* (3), 379–407.
- Setren, Elizabeth**, “Busing to Opportunity? The Impacts of the METCO Voluntary School Desegregation Program on Urban Students of Color,” Technical Report, National Bureau of Economic Research 2024.
- Shapley, Lloyd S and Martin Shubik**, “The assignment game I: The core,” *International Journal of game theory*, 1971, *1* (1), 111–130.
- The San Francisco Standard**, “Selective Admissions Are Back at Lowell High After a Split Vote Reinstates the Controversial Policy,” 2022. Published June 22, 2022. Accessed November 19, 2025.
- Train, Kenneth E**, *Discrete choice methods with simulation*, Cambridge university press, 2009.
- Trajkovski, Samantha, Jeffrey Zabel, and Amy Ellen Schwartz**, “Do School Buses Make School Choice Work?,” *Regional Science and Urban Economics*, 2021, *86*, 103607.
- Urquiola, Miguel**, “Does school choice lead to sorting? Evidence from Tiebout variation,” *American Economic Review*, 2005, *95* (4), 1310–1326.
- U.S. Department of Education, National Center for Education Statistics**, “Common Core of Data (CCD),” <https://nces.ed.gov/ccd> 2024. Public elementary/secondary education data. Accessed 2025-08-14.
- Walters, Christopher R**, “The demand for effective charter schools,” *Journal of Political Economy*, 2018, *126* (6), 2179–2223.
- Willis, Robert J and Sherwin Rosen**, “Education and self-selection,” *Journal of political Economy*, 1979, *87* (5, Part 2), S7–S36.
- WRAL News**, “Wake County Voters Oust School Board Majority,” 2009. Published October 6, 2009. Accessed November 19, 2025.

Online Appendix

Who Chooses and Who Benefits? The Design of Public School Choice Systems


Jesse Bruhn Christopher Campos Eric Chyn Antonia Vazquez

Table of Contents

A	Appendix Figures and Tables	1
B	Data Appendix	20
B.1	Details of the Original Data Collection	20
B.2	List of Districts	20
B.3	Additional Choice Options in LAUSD	25
C	Oversubscribed School Lottery Details	26
C.1	Choice School Lottery Offers	26
D	Demand Model	27
D.1	Overview	27
E	Counterfactual Details	38
E.1	Information Interventions	38
E.2	Busing Program	38
E.3	Decentralized	38
E.4	Mandatory Application	39
E.5	Mandatory Application and Complementary Policies	39


A Appendix Figures and Tables

Figure A.1: District Choice and Charter School Enrollment Trends (National Data)




Notes: This figure reports enrollment shares for charter schools and district-operated choice schools (“District Option”) at the national level. Data are from the Common Core of Data (U.S. Department of Education, National Center for Education Statistics 2024). For a given district, a local Charter includes all Charter schools that have the district’s Local Education Agency ID or operated in a zip-code served by a traditional public school with that district’s LEAID. To identify district operated choice schools, we use the Common Core’s magnet flag, since other forms of choice are not reliably recorded. In a few states (e.g., Ohio), magnet status is not reported. For that reason, we also count a school as a within district choice option if the word “magnet” appears in its name. The district operated choice numbers in this figure are potentially conservative, since many districts have choice options (e.g., gifted and talented schools) which are not officially named as a magnet and otherwise not easily identified.

Figure A.2: Expansions in Access to LAUSD Choice Schools




Notes: This figure reports three maps covering the region of Los Angeles County that is served by the Los Angeles Unified School District (LAUSD). In each map, we plot census blocks with their shading determined by different measures of proximity or changes in proximity, whose legend is reported in each panel. White-bordered polygons overlaid on top of each map are middle school attendance zone boundaries for LAUSD schools. Panel (a) and Panel (b) display block-level distances to nearest choice schools in 2004 and 2023, respectively. Panel (c) reports Census block level changes in relative distance to choice schools between 2004 and 2023, where relative distance is defined as the difference between a block's distance to the nearest choice program and the nearest non-choice neighborhood school.

Figure A.3: Lottery Effects by Similarity Index




(a) Math




(b) ELA

Notes: This figure reports reduced form effects of receiving an offer on Math and ELA achievement in Panel (a) and Panel (b), respectively. We report estimates for different quintiles of a similarity index constructed using the observable characteristics in Figure 7. Specifically, we standardize each covariate to have a mean of zero and a unit variance, and then calculate the Euclidean distance of each student to the average applicant to that choice program in that same year. The estimates come from a regression of achievement—Math or ELA—on lottery strata fixed effects, main effects for preference index quintiles, and interactions of offers and preference quintiles. We report estimates of the interactions in each sub-figure. We report 95-percent confidence intervals that use robust standard errors clustered at the lottery strata level. We also report p -values from hypothesis tests that test the null hypothesis that the preference-index effects are equal and another test that tests the null hypothesis that the preference-index effects are jointly equal to zero.

Figure A.4: Subgroup Heterogeneity in Application Rates and Lottery Effects




(a) Math



(b) ELA


Notes: This figure reports subgroup heterogeneity in lottery-based treatment effects and application rates. As indicated in the legend, each marker type represents results for a given student subgroup and its complement (e.g., English versus Non-English Learner students). With the exception of baseline achievement, all subgroups are based on binary pre-determined characteristics. In the case of achievement, we define high and low-ability sub-groups as those above and below the median of baseline achievement. We omit creating a subgroup based on suspensions, as it is a relatively rare outcome. The y -axis reports the given subgroups treatment effect (e.g., the effect of receiving an offer to attend a choice school for English Learners). All estimates are from separate subgroup specific models where we regress an achievement outcome on the lottery offer Z_i and lottery fixed effect. The x -axis reports the corresponding choice-school application rate for a given subgroup, which we measure as the share of all 5th graders who submitted at least one choice school application over our sample period. Panels (a) and (b) report results for Math and ELA, respectively. Detailed estimates are included in Appendix Table A.4.

Figure A.5: Application Cost Identification Exercise




Notes: This figure reports results from an identification exercise for the application costs in the model specified in Section 6.2. The vertical axis corresponds to application rates in the simulated data while the horizontal axis varies estimated application costs ($c(a, X_i, \eta_i)$), where we scale costs by corresponding number on the horizontal axis. The black line reports the corresponding application rate as we scale costs while holding other parameters of the model constant. For example, at scale factor equal to one, the value of the black line corresponds to the simulated application rate when cost parameters are evaluated at the values that maximize the likelihood. Similarly, at scale factor equal to 0.5, we scale costs by 0.5 and hold other parameters of the model constant at the values that maximize the likelihood. The dashed black line corresponds to application rates observed in the data. The intuition is that a moment is important for a parameter's identification if, as we move across the scaling factor axis, the values of simulated application rates change and cross the horizontal dashed line (i.e., the value of that moment in the real-world data). Overall, this exercise provides graphical evidence of how application costs are identified.

Figure A.6: Model Fit: Application and Enrollment




(a) Applications




(b) Enrollment

Notes: This figure reports an assessment of the overall fit of the model based on out-of-sample validation. We estimate our preferred model on a 25% random training sample and assess its fit on the remaining 75%. We run 200 simulations in which we use parameter estimates from the training sample to simulate applications, offers, and enrollment in the hold-out sample. The figure plots school-level aggregates based on simulated averages and the actual data. Panel (a) reports school-level observed number of applications in the holdout sample (y -axis) against the model-predicted applications in the holdout sample (x -axis). Maroon markers correspond to schools, or in other words, assess the fit at of school-level statistics, while the grey markers further interact school with student observable attributes. Panel (b) reports similar statistics for enrollment.

Figure A.7: Distribution of Lottery Applicant Choice School Preferences




(a) θ Percentile




(b) Preference Index Percentile

Notes: This figure focuses on the sample of students who are in the lottery estimation sample (i.e., students who applied to oversubscribed choice schools) and reports the distribution of their estimated posterior means of θ_i and a preference index measure in Panels (a) and (b), respectively. The preference index is defined as $\mathcal{I}_i \equiv \beta_c X_i + \theta_i$, which can be interpreted as student i 's utility from enrolling in a choice school as a function of both observables and unobservables. Both figures demonstrate that students that are actively engaging in the choice process are drawn from regions of the distribution that are associated with the smallest learning gains.

Figure A.8: Model Validation: Visual Instrumental Variables Plot




(a) Student Achievement Model 1




(b) Student Achievement Model 2

Notes: This figure reports results from a validation of our student achievement outcome model following Angrist et al. (2013). Each dot in the figure represents results from an oversubscribed choice school lottery (specifically, the $N = 420$ lotteries that overlap with the 2004-2013 time window used to estimate the structural model). The “first-stage” (x -axis) is the estimated effect of a randomized offer Z_i in the given lottery on the model-based prediction of student achievement based on the model from Section 6.3. The “reduced form” (y -axis) is the estimated effect of the same randomized offer Z_i in the given lottery on observed achievement. Panels (a) and (b) report results for the model that models selection according to Equations M1 and M2. The slope (i.e., the IV “forecast coefficient”) equals 0.91 (0.06), so we cannot reject 1, indicating the outcome model predicts causal impacts up to sampling error. All specifications include lottery-by-grade strata fixed effects.

Figure A.9: Choice School Peer Effect Analysis




(a) Math




(b) ELA

Notes: This figure reports lottery-based reduced form effects on own student achievement and school peer quality for Math and ELA in Panel (a) and Panel (b), respectively. Each marker in a given plot corresponds to a choice school, where the vertical axis reports the effect of winning an offer from the school on own achievement, and the horizontal axis reports the effect of winning an offer from the school on enrolled school peer quality. All estimated effects are conditional on lottery strata.

Figure A.10: Estimated Effects of Counterfactual Results on District-Level Achievement by Race



(a) Change in Group-Specific Achievement (Math)



(b) Change in Group-Specific Achievement (ELA)

Notes: This figure reports counterfactual achievement results from several scenarios for Black, White, Hispanic, and Asian students. Panel (a) and Panel (b) provide results for Math and ELA achievement, respectively. We consider participation-targeted policies: General Information, which raises the information parameter θ_i for a randomly selected 50% of students; Targeted Information, which raises θ_i for a randomly selected 50% of students who are enrolled in schools whose average achievement is below median; and Busing, which sets distance costs to zero. We also consider two system-level designs: Decentralized choice, which allows families to submit multiple applications and hold simultaneous offers while schools run independent lotteries, and Mandatory application, which sets application costs to zero, requires a rank-ordered list (which may include the neighborhood school), and assigns seats via deferred acceptance. We also report combinations of these components (Mandatory + Busing; Mandatory + Information; Mandatory + Information + Busing) and an Achievement-Optimal benchmark that assigns students in order of modeled match quality. Each scenario is simulated 100 times, and we report average effects; for cases with school-run lotteries, we recompute equilibrium best-response admission probabilities under policy-induced demand (see Appendix Section E).

Table A.1: Summary Statistics for Main Analysis Samples

	Sample		
	All		
	Grade 5	Lottery	Structural
	(1)	(2)	(3)
Student Demographics			
Hispanic	0.74	0.57	0.76
Black	0.10	0.14	0.09
White	0.09	0.13	0.08
Asian	0.04	0.12	0.04
Female	0.49	0.51	0.50
Poverty	0.72	0.63	0.70
English Learner	0.27	0.08	0.32
Standardized Test Scores			
Baseline ELA	-0.01	0.67	-0.05
Baseline Math	-0.01	0.65	-0.04
Distance to District Schools			
Nearest	0.75	0.78	0.76
Nearest Choice (Rel. Dist.)	0.57	0.55	0.59
Students	593,550	27,203	337,661

Notes: This table reports summary statistics for the various analysis samples we use throughout the paper. Column 1 reports statistics for the baseline sample consisting of all fifth-grade students in LAUSD during our sample period. Column 2 restricts the baseline sample to students who show up in our lottery sample, indicating they applied to a choice program and the choice program was oversubscribed. Column 3 restricts the baseline sample to students who meet the restrictions defined in Section 3.3.

Table A.2: Oversubscribed Choice Program Lottery Balance Analysis

Distance Quintile	Q1		Q2		Q3		Q4		Q5	
Variable	Effect	Control Mean	Effect	Control Mean	Effect	Control Mean	Effect	Control Mean	Effect	Control Mean
Female	-0.005 (0.014)	0.507	0.021 (0.015)	0.511	0.004 (0.015)	0.506	0.004 (0.015)	0.497	-0.006 (0.013)	0.512
Black	0.005 (0.004)	0.057	0.002 (0.006)	0.065	-0.010 (0.008)	0.108	-0.010 (0.009)	0.168	0.010 (0.010)	0.277
Hispanic	-0.006 (0.008)	0.765	-0.001 (0.009)	0.648	0.009 (0.010)	0.545	0.011 (0.010)	0.475	0.000 (0.012)	0.501
English Learner	0.005 (0.007)	0.074	-0.001 (0.007)	0.071	0.011* (0.006)	0.054	-0.003 (0.007)	0.054	-0.004 (0.006)	0.061
Special Education	0.000 (0.005)	0.027	-0.003 (0.004)	0.021	-0.001 (0.004)	0.020	-0.002 (0.003)	0.018	0.000 (0.003)	0.016
Poverty	0.018* (0.010)	0.758	-0.019* (0.011)	0.686	0.009 (0.012)	0.614	-0.018 (0.012)	0.573	-0.003 (0.013)	0.631
Parent reports going to college	-0.001 (0.012)	0.238	-0.010 (0.012)	0.254	-0.004 (0.013)	0.269	0.002 (0.013)	0.274	-0.006 (0.012)	0.262
Speaks english at home	0.016 (0.012)	0.322	-0.001 (0.012)	0.356	-0.007 (0.012)	0.435	-0.001 (0.013)	0.508	0.017 (0.012)	0.519
Speaks Spanish at home	-0.021* (0.012)	0.609	0.010 (0.011)	0.524	0.009 (0.010)	0.415	-0.001 (0.011)	0.353	-0.011 (0.011)	0.377
Baseline Suspensions	-0.002 (0.003)	0.007	0.004 (0.002)	0.007	0.001 (0.004)	0.011	0.003 (0.004)	0.015	-0.001 (0.005)	0.024
Baseline Math Score	-0.013 (0.021)	0.365	-0.033 (0.022)	0.480	-0.008 (0.022)	0.606	-0.019 (0.025)	0.569	-0.061*** (0.022)	0.414
Baseline ELA Score	-0.031 (0.021)	0.353	-0.015 (0.022)	0.482	-0.001 (0.023)	0.619	-0.007 (0.023)	0.605	-0.030 (0.020)	0.470
Missing Baseline Score	-0.003 (0.007)	0.089	-0.005 (0.006)	0.099	-0.002 (0.005)	0.083	-0.010* (0.005)	0.094	0.000 (0.005)	0.088
Joint test	0.582		0.275		0.850		0.390		0.520	
Observations	7,794		8,255		8,701		9,121		9,729	

Notes: This table tests for lottery balance across baseline and other exogenous covariates. Columns in this table are grouped according to samples defined by quintiles of distance to choice school at baseline (Q1-Q5). Each row in this table denotes a predetermined covariate. Each cell under the columns labeled "Effect" gives the results of a separate regression of the relevant sub-sample. Standard errors, clustered at the level of the lottery, are included in parentheses. Columns labeled "Control Mean" give the mean of the variable in that row for that specific sample. The second to last row of this table gives p -values from a joint test that all of the coefficients in that column are identically zero. The final row gives the number of observations.

Table A.3: Oversubscribed School Lottery First Stage Estimates by Quintile

	Quintile of Distance to Choice School				
	Q1	Q2	Q3	Q4	Q5
Lottery Offer	0.381 (0.022)	0.409 (0.018)	0.459 (0.015)	0.434 (0.018)	0.455 (0.019)

Notes: This table reports the first stage effects of receiving an offer. We report estimates for different quintiles of distance to choice school at baseline. The estimates come from a regression of enrolling in a choice school on lottery strata fixed effects, main effects for distance quintiles, and interactions of offers and distance quintiles. We report estimates of the interactions. Standard errors are clustered at the lottery level.

Table A.4: Lottery Estimates by Subgroup

Covariate	Math		ELA	
	Focal Group	Complement Group	Focal Group	Complement Group
	Covariate = 1	Covariate = 0	Covariate = 1	Covariate = 0
High Achiever	-0.041*** (0.013)	0.025* (0.014)	-0.049*** (0.013)	0.016 (0.016)
Female	-0.014 (0.015)	-0.038** (0.015)	-0.031** (0.015)	-0.039*** (0.015)
Poverty	-0.016 (0.013)	-0.064*** (0.019)	-0.030** (0.013)	-0.059*** (0.019)
Special Education	0.091 (0.079)	-0.030*** (0.011)	0.093 (0.074)	-0.039*** (0.011)
Gifted	-0.026 (0.025)	-0.030** (0.012)	-0.033 (0.027)	-0.038*** (0.012)
English Learner	0.085*** (0.031)	-0.033*** (0.011)	0.035 (0.032)	-0.040*** (0.011)
English at home	-0.065*** (0.017)	-0.000 (0.014)	-0.074*** (0.018)	-0.009 (0.013)
Spanish at home	0.013 (0.014)	-0.061*** (0.015)	0.005 (0.014)	-0.070*** (0.016)
Born in USA	-0.027** (0.011)	-0.040 (0.040)	-0.038*** (0.011)	-0.028 (0.038)

Notes: This table presents heterogeneous effects by student subgroups. Each row in this table denotes one of the pre-determined characteristics on display in Figure 7. With the exception of baseline achievement, all subgroups are based on binary pre-determined characteristics. In the case of achievement, we define high and low-ability sub-groups as those above and below the median of baseline achievement. We omit creating a subgroup based on suspensions, as it is a relatively rare outcome. All estimates are from separate subgroup specific models where we regress an achievement outcome on the lottery offer Z_i and lottery fixed effect. The first two columns correspond to treatment effects on Math; the second two columns correspond to treatment effects on ELA. Columns labeled “Focal Group” denote estimates from the sub-group of students where the covariate equals 1. “Complement Group” denotes estimates from the sub-group of students where the covariate equals zero. Standard errors, clustered at the level of the lottery, are included in parentheses. Stars denote conventional levels of significance. A visualization of these estimates, and their relationship to application rates by sub-group, is displayed in Appendix Figure A.4.

Table A.5: Comparison of Alternative Models of Idiosyncratic Choice School Preferences

	One Type ($K = 1$)	Two Types ($K = 2$)	Three Types ($K = 3$)
Number of parameters	94	97	100
Log likelihood	-64,260.9	-62,872.4	-62,625.2
BIC	129,587.9	126,844.9	126,384.6
Likelihood ratio tests:			
χ^2 statistic (df)		2,777.1 (3)	494.3 (3)
p -value		0	0

Notes: This table reports estimation statistics for different models we estimate. For each, we report the number of parameters associated with the model, the negative log-likelihood, a Bayesian Information Criterion (BIC) measure, and the statistic and p -value of a likelihood ratio test comparing a given model against the previous.

Table A.6: Demand Estimates: Mixture Model ($K = 2$)

Panel (a): Estimates for Observable Parameters

	Utility (1)	Distance Cost (2)	Log Cost (3)
Main Effects	-0.485 [1.132]	0.321 (0.005)	0.298 (0.013)
Female	0.020 (0.019)	0.001 (0.003)	0.003 (0.003)
Black	-0.072 (0.047)	-0.021 (0.007)	-0.032 (0.008)
Hispanic	0.006 (0.032)	0.008 (0.005)	0.068 (0.006)
White	-1.106 (0.036)	0.104 (0.007)	-0.080 (0.007)
Poverty	0.056 (0.022)	0.014 (0.003)	0.013 (0.003)
LEP	-0.308 (0.035)	-0.023 (0.005)	0.006 (0.004)
Speaks English at Home	-0.151 (0.025)	-0.022 (0.004)	-0.010 (0.004)
Baseline ELA	0.130 (0.014)	0.001 (0.002)	-0.033 (0.002)
Baseline Math	0.158 (0.014)	0.007 (0.002)	-0.019 (0.002)
Neighborhood Median Income	-0.050 (0.004)	0.001 (0.001)	0.006 (0.001)
Baseline Choice Enrollment	1.055 (0.031)	-0.082 (0.004)	-0.093 (0.006)
Baseline Peer Quality	0.125 (0.022)	-0.025 (0.004)	-0.023 (0.004)

Panel (b): Estimates for Unobservable Parameters

	μ (1)	σ (2)	$Pr(K_i = k)$ (3)
Type 1	-0.220 (0.008)	0.059 (0.025)	0.906
Type 2	2.120 (0.045)	0.546 (0.033)	0.094
Cost Heterogeneity		0.400 (0.006)	

Notes: This table reports estimates of an alternative two-type ($K = 2$) version of the demand model in Section 6.2, estimated via simulated maximum likelihood with 300 draws for taste heterogeneity (θ_i) and cost heterogeneity (η_i). Panel (a) reports observable heterogeneity: the first row gives main effects for distance and log cost; school mean utilities are shown as averages with noise-adjusted SDs in brackets; remaining rows are heterogeneity terms. Panel (b) reports unobservables: a two-type mixture-of-normals for tastes (means in col. 1, SDs in col. 2, type probabilities in col. 3) and mean-zero normal costs with the SD in col. 2. Standard errors are in parentheses.

Table A.7: Outcome Model Estimates ($K = 2$)

	Math		ELA	
	Neighborhood School (1)	Choice School (2)	Neighborhood School (3)	Choice School (4)
Main Effects		0.257 [0.164]		0.235 [0.115]
Female	-0.019 (0.002)	-0.014 (0.007)	0.117 (0.002)	-0.004 (0.006)
Black	-0.215 (0.011)	0.046 (0.022)	-0.182 (0.010)	0.043 (0.021)
White	-0.007 (0.012)	-0.031 (0.022)	-0.001 (0.009)	-0.031 (0.022)
Hispanic	-0.154 (0.010)	-0.003 (0.020)	-0.123 (0.008)	0.015 (0.019)
Asian	0.212 (0.014)	-0.066 (0.026)	0.076 (0.010)	-0.050 (0.023)
Poverty	0.008 (0.003)	-0.008 (0.010)	0.000 (0.002)	-0.002 (0.007)
LEP	-0.115 (0.004)	0.056 (0.013)	-0.207 (0.004)	0.030 (0.014)
English Home	-0.085 (0.004)	0.024 (0.012)	-0.076 (0.004)	0.061 (0.011)
Baseline Math	0.524 (0.004)	-0.034 (0.009)	0.183 (0.002)	-0.023 (0.006)
Baseline ELA	0.201 (0.003)	0.026 (0.007)	0.576 (0.003)	0.001 (0.008)
Baseline Peer Quality	-0.143 (0.012)	0.035 (0.019)	-0.091 (0.009)	0.024 (0.014)
Choice School Preference θ_i	0.046 (0.004)	-0.107 (0.011)	0.050 (0.003)	-0.098 (0.008)
Neighborhood Effects	✓		✓	
Year Effects	✓		✓	
Sub-district Effects	✓		✓	
H_0 : No selection on unobservables (p -values)	0.000		0.000	
H_0 : No treatment effect heterogeneity (p -values)	0.000		0.000	
Observations	334,166		334,166	

Notes: This table reports estimates of the achievement model in Section 6.3 for Math and ELA assuming that a two-type mixture distribution ($K = 2$). Cols. 1 and 2 correspond to Math estimates and Cols. 3 and 4 correspond to ELA estimates. The first row labeled main effects reports the enrollment-weighted-average of the choice school effects with noise-adjusted standard deviation reported in brackets below. Cols. 1 and 3 report main effects for corresponding covariates labeled in each row, while Cols. 2 and 4 report choice school treatment effect heterogeneity with respect to the labeled row variable. Both models include neighborhood fixed effects—defined as Census blocks–year fixed effects, and sub-district fixed effects. Bottom rows report p -values for (i) no selection on unobservable tastes (the main effect of θ and the effect heterogeneity are jointly zero) and (ii) no treatment-effect heterogeneity (e.g., choice-enrollee interactions in Col. 2 are jointly zero in each model). Robust standard errors clustered at the enrolled school level are reported in parentheses.

Table A.8: Outcome Model Bootstrapped Estimates

	Model 1: Type Probabilities				Model 2: Linear Control Function (θ_i)			
	Math		ELA		Math		ELA	
	Neighborhood School	Choice School	Neighborhood School	Choice School	Neighborhood School	Choice School	Neighborhood School	Choice School
	(γ_0) (1)	Diff (γ_c) (2)	(γ_0) (3)	Diff (γ_c) (4)	(γ_0) (5)	Diff (γ_c) (6)	(γ_0) (7)	Diff (γ_c) (8)
Main Effects		0.252 (0.003) [0.173]		0.212 (0.001) [0.117]		0.223 (0.002) [0.171]		0.198 (0.001) [0.115]
Female	-0.024 (0.002)	-0.012 (0.007)	0.109 (0.002)	0.001 (0.006)	-0.024 (0.002)	-0.012 (0.007)	0.109 (0.002)	0.001 (0.006)
Black	-0.276 (0.016)	0.080 (0.026)	-0.222 (0.013)	0.058 (0.020)	-0.276 (0.016)	0.080 (0.026)	-0.222 (0.013)	0.058 (0.020)
White	0.092 (0.018)	-0.072 (0.027)	0.064 (0.013)	-0.052 (0.022)	0.092 (0.018)	-0.076 (0.027)	0.064 (0.013)	-0.054 (0.022)
Hispanic	-0.174 (0.013)	0.030 (0.022)	-0.142 (0.010)	0.033 (0.016)	-0.175 (0.013)	0.026 (0.022)	-0.142 (0.010)	0.031 (0.016)
Asian	0.270 (0.017)	-0.065 (0.025)	0.107 (0.011)	-0.052 (0.020)	0.270 (0.017)	-0.063 (0.026)	0.106 (0.011)	-0.050 (0.020)
Poverty	-0.016 (0.004)	0.000 (0.009)	-0.021 (0.003)	0.001 (0.007)	-0.016 (0.004)	-0.001 (0.009)	-0.021 (0.003)	0.000 (0.007)
LEP	-0.084 (0.004)	0.033 (0.013)	-0.170 (0.004)	0.001 (0.014)	-0.083 (0.004)	0.037 (0.013)	-0.170 (0.004)	0.003 (0.014)
Median Income		0.015 (0.002)		0.013 (0.002)		0.015 (0.002)		0.012 (0.002)
English Home	-0.037 (0.005)	-0.020 (0.012)	-0.028 (0.005)	0.020 (0.011)	-0.037 (0.005)	-0.020 (0.012)	-0.027 (0.005)	0.020 (0.011)
Baseline Math	0.518 (0.004)	-0.032 (0.009)	0.173 (0.003)	-0.016 (0.006)	0.518 (0.004)	-0.031 (0.010)	0.173 (0.003)	-0.016 (0.006)
Baseline ELA	0.227 (0.003)	0.004 (0.006)	0.607 (0.003)	-0.024 (0.008)	0.227 (0.003)	0.006 (0.006)	0.607 (0.003)	-0.024 (0.008)
Baseline Peer Quality	-0.063 (0.013)	-0.035 (0.021)	-0.029 (0.010)	-0.025 (0.017)	-0.062 (0.013)	-0.032 (0.022)	-0.028 (0.010)	-0.024 (0.017)
Pr(Type 2)	-0.026 (0.042)	-0.152 (0.055)	0.001 (0.030)	-0.108 (0.042)				
Pr(Type 3)	0.088 (0.022)	-0.341 (0.045)	0.105 (0.015)	-0.269 (0.032)				
Choice School Preference θ_i					0.025 (0.005)	-0.068 (0.010)	0.029 (0.004)	-0.060 (0.007)
Neighborhood Effects	✓		✓		✓		✓	
Year Effects	✓		✓		✓		✓	
Sub-district Effects	✓		✓		✓		✓	
Bootstrap Iterations	100		100		✓		✓	

Notes: This table reports bootstrapped estimates of the achievement model in Section 6.3 for Model 1 (Cols. 1–4) and Model 2 (Cols. 5–8). This analysis addresses concern over estimation error in either type probability estimates (in Model 1) or the posterior θ_i (in Model 2) estimates. We assume a three-type ($K = 3$) mixture model for the distribution of θ_i . For each bootstrap iteration, we draw from the asymptotic distribution of the underlying demand estimates, recompute either the posterior type probabilities (Model 1) or posterior means of θ_i (Model 2), and then re-estimate the outcome model. The entries in the table are bootstrap averages of the corresponding point estimates; bootstrapped standard errors clustered by enrolled school are reported in parentheses. Cols. 1, 3, 5, and 7 report baseline coefficients (γ_0) for the covariates listed in each row; Cols. 2, 4, 6, and 8 report the corresponding choice-enrollee interactions (γ_c). The “Main effects” row gives the enrollment-weighted average effect of attending a choice school; noise-adjusted SDs across schools are shown in brackets. All specifications include Census-block, year, and sub-district fixed effects.

Table A.9: Outcome Model Estimates with Distance Heterogeneity

	Model 1: Type Probabilities				Model 2: Linear Control Function (θ_i)			
	Math		ELA		Math		ELA	
	Neighborhood School	Choice School	Neighborhood School	Choice School	Neighborhood School	Choice School	Neighborhood School	Choice School
	(γ_0) (1)	Diff (γ_c) (2)	(γ_0) (3)	Diff (γ_c) (4)	(γ_0) (5)	Diff (γ_c) (6)	(γ_0) (7)	Diff (γ_c) (8)
Main Effects		0.288 [0.160]		0.243 [0.113]		0.251 [0.158]		0.227 [0.111]
Female	-0.022 (0.002)	-0.009 (0.007)	0.111 (0.002)	0.002 (0.006)	-0.022 (0.002)	-0.010 (0.007)	0.111 (0.002)	0.001 (0.006)
Black	-0.196 (0.011)	0.029 (0.022)	-0.161 (0.010)	0.027 (0.020)	-0.197 (0.011)	0.029 (0.022)	-0.161 (0.010)	0.027 (0.020)
White	-0.004 (0.012)	-0.027 (0.021)	0.001 (0.009)	-0.027 (0.021)	-0.004 (0.012)	-0.033 (0.021)	0.001 (0.009)	-0.030 (0.021)
Hispanic	-0.149 (0.010)	-0.003 (0.020)	-0.119 (0.008)	0.012 (0.019)	-0.150 (0.010)	-0.008 (0.020)	-0.119 (0.008)	0.010 (0.019)
Asian	0.211 (0.014)	-0.062 (0.025)	0.075 (0.010)	-0.042 (0.022)	0.210 (0.014)	-0.059 (0.025)	0.075 (0.009)	-0.041 (0.022)
Poverty	0.010 (0.003)	-0.011 (0.010)	0.003 (0.002)	-0.006 (0.007)	0.010 (0.003)	-0.012 (0.010)	0.003 (0.002)	-0.007 (0.007)
LEP	-0.087 (0.004)	0.037 (0.013)	-0.171 (0.003)	0.010 (0.014)	-0.087 (0.004)	0.042 (0.013)	-0.171 (0.003)	0.012 (0.014)
Median Income		-0.005 (0.002)		-0.006 (0.002)		-0.006 (0.002)		-0.006 (0.002)
English Home	-0.070 (0.004)	0.009 (0.012)	-0.059 (0.004)	0.044 (0.010)	-0.070 (0.004)	0.009 (0.012)	-0.059 (0.004)	0.044 (0.010)
Baseline Math	0.526 (0.004)	-0.038 (0.009)	0.182 (0.002)	-0.023 (0.006)	0.526 (0.004)	-0.037 (0.009)	0.182 (0.002)	-0.023 (0.006)
Baseline ELA	0.210 (0.003)	0.012 (0.007)	0.591 (0.003)	-0.016 (0.008)	0.210 (0.003)	0.014 (0.007)	0.591 (0.003)	-0.016 (0.008)
Baseline Peer Quality	-0.149 (0.012)	0.032 (0.018)	-0.095 (0.009)	0.024 (0.014)	-0.147 (0.012)	0.036 (0.019)	-0.094 (0.009)	0.025 (0.014)
Relative Distance		0.004 (0.006)		-0.002 (0.005)		0.004 (0.006)		-0.002 (0.005)
Pr(Type 2)	-0.053 (0.032)	-0.168 (0.047)	-0.006 (0.023)	-0.118 (0.040)				
Pr(Type 3)	0.076 (0.015)	-0.387 (0.039)	0.101 (0.012)	-0.308 (0.031)				
Choice School Preference θ_i					0.023 (0.003)	-.076 (.009)	0.029 (0.003)	(-.070) (.007)
Neighborhood Effects	✓		✓		✓		✓	
Year Effects	✓		✓		✓		✓	
Sub-district Effects	✓		✓		✓		✓	

Notes: This table reports estimates of the achievement model in Section 6.3 for Math (Columns 1–4) and ELA (Cols. 5–8), augmented to allow treatment-effect heterogeneity in relative distance to the nearest choice school. Neighborhood fixed effects absorb all cross-sectional variation in relative distance, so identification comes from within-block changes in distance. Cols. 1, 3, 5, and 7 report baseline coefficients (γ_0) for the covariates listed in each row; Cols. 2, 4, 6, and 8 report the corresponding choice-enrollee interactions (γ_c). The “Main effects” row gives the enrollment-weighted average effect of attending a choice school; noise-adjusted SDs are in brackets. All specifications include Census-block, year, and sub-district fixed effects. Robust SEs clustered by enrolled school are in parentheses.

B Data Appendix

B.1 Details of the Original Data Collection

We collected data for the 150 highest enrollment school districts in the United States that cover roughly 27% of public school enrollment nationwide. The data we collected was obtained through both online research and calls to district administrators made by research assistants. For each school district, we were interested in obtaining measures of the availability of intra-district choice, the assignment process, the timing, participation details, the type of choice offerings a district provided, and an assessment of the difficulty in obtaining this information via the web and phone. A full list of the variables that we coded is provided in Appendix Table B.1. The process began with a research assistant examining the district’s webpage and collecting all relevant information that was online. Next, research assistants made calls to school districts to obtain information via conversations with administrators. Most school districts were researched by three research assistants. In the case of discrepancies, the principal investigators verified the accuracy of the information.

B.2 List of Districts

Appendix Table B.2 reports the districts that are flagged as districts that use a centralized algorithm for assignments and have de facto mandatory participation. Two popular districts in this sample are New York City and Boston, both the subject of extensive literature (Abdulkadiroğlu et al. 2017). Appendix Table B.3 reports the districts that use a centralized algorithm for participation but have an opt-in design. Last, Table B.4 reports the districts that do not use a centralized algorithm for assignments and run independent lotteries for each one of their choice schools. For each table, we report the state the district belongs to along with its enrollment based on 2021-2022 Common Core data maintained by the NCES, and its rank in terms of enrollment. These are the districts we use to define groups in Figure 2 in the main body of the paper.

Table B.1: 150 Largest School Districts, Choice Program Survey Variables

Variable Name	Description	Type
District	District name (state)	String
Zoned School	Indicates if families are assigned a zoned school	Boolean
School Choice	Indicates if district has open or specialized program enrollment	Boolean
Open Enrollment	Indicates if the district has any open enrollment process	Boolean
Open Offerings	If the district has open enrollment, what schools can families apply	String
Timing	If the district has open enrollment, how often families can apply	Numeric
Mandatory	Indicates if open enrollment is mandatory in the district	Boolean
Mandatory Notes	If the district has mandatory open enrollment, how is it mandatory	String
Mandatory Grades	If mandatory, what grades are mandatory	String
Centralized/Algorithm	Are school choice offers determined via an algorithm	Boolean
Max Apply	If there is open enrollment, max number of programs you can apply	Numeric
Ranked Choice	If there is open enrollment, do you rank the programs you apply	Boolean
Num Offers	If there is open enrollment, how many offers can you receive	Numeric
Timeline	If there is open enrollment, what months are applications open	String
Easy to Learn	Indicates if info. is easy to gather from the district website or a call	Boolean
Website Easy	Indicates if it is easy to gather choice info. from the district website	Boolean
Call Easy	Indicates if it is easy to gather choice info. by calling the district	Boolean
Magnet	Indicates if the school district has any magnet programs	Boolean
Dual Language	Indicates if the district offers any dual language programs	Boolean
CTE	Indicates if the district has specialized career-technical education	Boolean
General Theme Based	Indicates if the district has any general or other themed programs	Boolean
Virtual Schools	Indicates if the district offers full-time only virtual school options	Boolean
Max Apply Special	If there are any specialized programs, max number you can apply	Numeric
Ranked Choice Special	If there are specialized programs, do you rank choices	Boolean
Offer Algorithm Special	If there are specialized programs, what is the algorithm for offers	String
Num Offers Special	If there are specialized programs, max number of specialized offers	Numeric

Table B.2: List of Mandatory Districts

State	District	Rank	Enrollment
New York	New York City Public Schools	1	845,509
Texas	Houston ISD	10	184,109
Florida	Lee	27	100,064
Kentucky	Jefferson County	29	94,793
Colorado	Denver	34	88,258
Maryland	Baltimore City Public Schools	45	75,811
Colorado	Jefferson County School District No. R-1	47	74,251
Texas	Austin ISD	49	72,830
Wisconsin	Milwaukee School District	55	66,864
California	Long Beach Unified	60	63,966
Texas	Garland ISD	80	51,659
California	San Francisco Unified	94	48,736
Massachusetts	Boston	109	45,742

Table B.3: List of Opt-In Districts

State	District	Rank	Enrollment
California	Los Angeles Unified	2	419,929
Florida	Miami-Dade	3	335,500
Illinois	Chicago Public Schools Dist 299	4	322,809
Nevada	Clark County	5	304,565
Florida	Broward	6	251,408
Florida	Hillsborough	7	224,152
Florida	Orange	8	206,815
Florida	Palm Beach	9	189,777
North Carolina	Wake County Schools	14	161,481
North Carolina	Charlotte-Mecklenburg Schools	16	144,116
Florida	Duval	19	127,971
Texas	Cypress-Fairbanks ISD	20	118,470
Pennsylvania	Philadelphia City SD	21	117,907
Florida	Polk	22	111,041
Tennessee	Memphis-Shelby County Schools	24	110,057
Georgia	Cobb County	25	105,368
Texas	Northside ISD	26	101,095
California	San Diego Unified	28	95,492
Georgia	Dekalb County	31	91,938
Florida	Pinellas	32	91,021
Georgia	Fulton County	35	88,043
Florida	Pasco	37	85,855
Tennessee	Davidson County	40	80,468
South Carolina	Greenville 01	42	78,371
Florida	Osceola	46	74,289
Florida	Brevard	48	73,810
Utah	Davis District	50	72,499
California	Fresno Unified	53	68,568
Florida	Seminole	57	65,443
Nevada	Washoe County	61	63,777
Florida	Volusia	62	62,742
California	Elk Grove Unified	63	62,603
Colorado	Douglas County School District No. RE 1	65	61,409
Tennessee	Knox County	66	60,604
Utah	Granite District	67	60,270
Utah	Jordan District	68	59,145
Texas	North East ISD	71	57,343
Texas	Arlington ISD	74	54,750
Texas	Klein ISD	75	53,093
Florida	Manatee	76	52,895
North Carolina	Winston-Salem / Forsyth County Schools	77	52,157
Colorado	Cherry Creek School District No. 5 (Arapahoe)	78	51,808
Nebraska	Omaha Public Schools	79	51,693
Tennessee	Rutherford County	81	51,595
District of Columbia	District of Columbia Public Schools	84	50,831
Georgia	Clayton County	85	50,832
Washington	Seattle School District No. 1	86	50,770
Louisiana	Jefferson Parish	87	50,467
California	Corona-Norco Unified	90	50,256
Georgia	Atlanta Public Schools	91	49,660
Florida	Lake	97	48,285
Florida	Collier	98	48,262
Michigan	Detroit Public Schools Community District	99	48,271
South Carolina	Horry 01	100	48,205
Texas	Pasadena ISD	103	47,486
Florida	St. Lucie	105	46,987
Kansas	Wichita	106	46,516
Texas	Round Rock ISD	107	46,197
Tennessee	Hamilton County	108	45,790
Florida	Marion	110	45,547
Florida	Sarasota	111	45,314
Ohio	Columbus City Schools District	112	45,181
California	San Bernardino City Unified	113	44,712
Texas	San Antonio ISD	114	44,670
Utah	Nebo District	116	44,446
Oregon	Portland SD 1J	117	44,039
Alaska	Anchorage School District	119	43,563
California	Clovis Unified	120	43,291
Georgia	Henry County	121	43,258
California	Kern High	122	43,116
Arizona	Chandler Unified District #80 (4242)	123	42,832
New Jersey	Newark Public School District	124	42,791
Georgia	Cherokee County	127	42,016
Arizona	Tucson Unified District (4403)	133	40,929
California	Capistrano Unified	134	40,836
Texas	Alief ISD	137	39,474
Florida	Clay	138	39,362
California	San Juan Unified	146	38,488
California	Sacramento City Unified	147	38,268
Texas	Mesquite ISD	148	38,265

Table B.4: List of Decentralized Districts

State	District	Rank	Enrollment
Georgia	Gwinnett County	11	182,214
Virginia	Fairfax County Public Schools	12	180,714
Maryland	Montgomery County Public Schools	15	160,223
Texas	Dallas ISD	17	143,196
Maryland	Prince George's County Public Schools	18	131,310
Maryland	Baltimore County Public Schools	23	110,215
Texas	Katy ISD	30	94,785
Virginia	Prince William County Public Schools	33	90,550
Utah	Alpine District	36	87,051
Maryland	Anne Arundel County Public Schools	38	84,436
Virginia	Loudoun County Public Schools	39	81,636
Texas	Fort Bend ISD	40	80,694
New Mexico	Albuquerque	44	76,756
Texas	Conroe ISD	51	72,352
Texas	Fort Worth ISD	52	71,060
North Carolina	Guilford County Schools	54	67,832
Virginia	Virginia Beach City Public Schools	58	64,896
Virginia	Chesterfield County Public Schools	59	64,132
Texas	Aldine ISD	69	57,844
Maryland	Howard County Public Schools	70	57,643
Arizona	Mesa Unified District (4235)	72	57,263
Georgia	Forsyth County	73	54,984
Florida	St. Johns	82	54,134
Alabama	Mobile County	83	50,929
Virginia	Henrico County Public Schools	88	50,463
South Carolina	Charleston 01	89	49,864
North Carolina	Cumberland County Schools	92	49,314
Texas	El Paso ISD	93	49,139
Texas	Humble ISD	95	48,552
Texas	Lewisville ISD	96	48,440
Texas	Plano ISD	101	47,899
Maryland	Frederick County Public Schools	102	47,681
Texas	Socorro ISD	104	47,304
Texas	Killeen ISD	118	43,864
Texas	Leander ISD	125	42,593
Louisiana	East Baton Rouge Parish	126	42,509
Tennessee	Williamson County	128	41,909
North Carolina	Union County Public Schools	130	41,487
Nebraska	Lincoln Public Schools	131	41,654
Texas	United ISD	132	41,203
Virginia	Chesapeake City Public Schools	135	40,576
Texas	Clear Creek ISD	136	40,200
Tennessee	Montgomery County	139	39,345
South Carolina	Berkeley 01	140	39,265
Idaho	Joint School District No. 2	141	38,991
California	Riverside Unified	142	38,855
Louisiana	St. Tammany Parish	144	38,734
Minnesota	Anoka-Hennepin School District	145	38,631
Colorado	Aurora Joint District No. 28 of the Counties of Adams and A	149	38,178
Maryland	Harford County Public Schools	150	38,158

B.3 Additional Choice Options in LAUSD

In the main text, we study the effects of choice schools between 2004 and 2017. The choice landscape has continued to evolve since the end of our sample period, and we discuss some additional options that do not play a major role in our analysis. Dual Language Programs are an increasingly popular elementary school option. Schools for Advanced Studies are a type of neighborhood school tracking program that academically advantaged students are automatically enrolled in, but out-of-zone students may also petition for enrollment. Due to our focus on middle schools, our analysis does not include Dual Language programs in any of the empirical analyses, including participation patterns, lottery-based estimates, or structural model estimates reported in Sections 4–8. Although Schools for Advanced Studies are part of the choice landscape, they differ from magnet programs and affiliated charter schools in that most eligible students are automatically enrolled rather than apply through an opt-in process.

Dual Language Programs Dual Language Education (DLE) schools offer grade-level instruction, biliteracy development, and sociocultural learning to students in two languages: English and a target language. LAUSD launched its first Spanish-English dual immersion program in 1991, followed by a Korean-English program in 1992. The programs gained steady momentum through the 1990s and 2000s, weathering significant challenges including California’s English-only movement under Proposition 227 in 1998, which severely restricted bilingual education statewide. Despite these constraints, LAUSD has maintained a commitment to dual language instruction, growing from 13 Two-Way Dual Language programs in 2005 to 146 by 2022. This persistence positioned the district to rapidly expand when California voters passed Proposition 58—a repeal of Proposition 227—in 2016, restoring flexibility for multilingual programming. Currently, LAUSD offers more than 230 DLE programs.

The district operates three distinct program models: Two-Way Immersion (TWI) programs that integrate native English speakers with native speakers of the target language; One-Way Immersion (OWI) programs designed specifically for English Learners whose home language matches the target language; and World Language Immersion (WLI) programs that provide English-speaking students access to second language instruction. DLE programming encompasses seven languages, Arabic, Armenian, French, Japanese, Korean, Mandarin, and Spanish, with Spanish programs comprising the largest share due to both demographic demand and the district’s substantial Latino enrollment. An overwhelming majority of these programs are elementary school programs, even though in recent years there has been some minor expansion into the secondary grade levels.

Schools for Advanced Studies LAUSD’s Schools for Advanced Studies (SAS) initiative began as a district-level strategy to recognize and expand exemplary gifted and talented education (GATE) at resident (i.e. non-magnet, non-charter) schools. The rationale was that many high-ability or gifted students were being underserved by the standard GATE framework, and that by designating certain schools as SAS “demonstration sites,” LAUSD could concentrate resources around models that might then influence practices across the district. Importantly, magnet schools, independent charters, and dual-language programs are explicitly ineligible for SAS designation, so SAS remains rooted in the resident school infrastructure.

Within LAUSD’s broader choice ecosystem, SAS programs satisfy a distinctive niche. Unlike magnet programs, SAS was not born of desegregation court orders or integration mandates; its purpose is not explicitly to draw students from across zones (though cross-boundary permitting is possible in some cases) but rather to strengthen the quality of gifted education within neighborhood schools. In practice this means SAS cannot be lumped with magnets or charters without obscuring key differences: SAS sites remain “resident schools first,” meaning that students living in their boundaries who qualify for GATE/SAS are typically placed automatically (unless opting out).

C Oversubscribed School Lottery Details

C.1 Choice School Lottery Offers

When applications to non-neighborhood schools (e.g., magnets and other specialized programs) exceed available seats, LAUSD allocates seats by lottery. Lotteries are run within strata defined by the application year, student grade, school/program, race, and priority-points level, so applicants only compete against others in the same stratum. The data available for our analysis report each applicant’s randomly assigned number, whether the applicant ultimately received an offer from the program, and all fields needed to define the lottery strata. Programs set an undisclosed cutoff on the ranking induced by the random numbers to fill seats; applicants on the offer side of this cutoff receive initial offers. As seats open, programs may also extend offers to waitlisted students and, in a small number of cases, make post-lottery offers to late applicants if capacity remains.

For the analysis, we construct an indicator for having a random number on the offer side of the cutoff. Because cutoffs are not recorded, we infer them from the joint distribution of random numbers and final offers, following Porter and Yu (2015), who extend RD methods to cases with unknown discontinuity points. Concretely, we estimate the threshold that maximizes the discontinuity in the offer probability.

Identification of Potential Lotteries: Lotteries occur at the intersection of year, school, grade, race, and priority points levels. We first identify which of these cells experienced lottery-based assignment using the following procedure. We begin by calculating offers and applications. For each combination of application year, student grade, school, race, and priority points level, we compute the total number of offers made and the total number of applications received. Next, we identify oversubscribed cells by flagging cells as oversubscribed when applications exceed offers. We then define potential lottery cells. A cell represents a potential lottery if it satisfies three conditions: it has multiple applicants, is oversubscribed, and has at least one offer made. Each unique combination of year, school, grade, race, and priority points level that meets the potential lottery criteria is considered a lottery.

Cutoff Estimation: For each potential lottery identified above, we estimate the lottery number cutoff using a regression-based approach that searches for the point of maximum discontinuity in assignment probability. Within each lottery cell, we rank applicants by their lottery number. We exclude the minimum and maximum values to avoid boundary issues. For each potential cutoff rank c , we create an indicator variable:

$$\text{above}_c = 1[\text{rank_lottery_number} \geq c]. \quad (16)$$

This indicator equals one if an applicant’s lottery number rank is at or above the candidate cutoff. Then, we estimate the discontinuity at each candidate cutoff. Formally, the specification is:

$$\text{offer_magnet}_i = \alpha + \beta \cdot \text{above}_c + \epsilon_i. \quad (17)$$

The coefficient β captures the change in probability of receiving an offer at the candidate cutoff c . Standard errors are heteroskedasticity robust. We select the cutoff value c that maximizes the R^2 among all candidates satisfying:

- $\beta > 0$ (offers increase above the cutoff)
- $p\text{-value} \leq 0.01$ (statistically significant at the 1% level) or $p\text{-value}$ is missing (which occurs when $\beta = 1$ with perfect prediction)
- Model is identified (rank condition satisfied)

D Demand Model

This appendix section provides technical details related to the estimation of the model, primarily discussing the simulated likelihood function and its associated gradient.

D.1 Overview

The model outlined in Section 6 has a likelihood function given by Equation 10 that depends on a vector of parameters $\Omega = (\{\delta_j\}, \beta_x, \psi_D, \psi_{D2}, \mu_c, \mu_{cx}, \{\mu_k, \sigma_k, p_k\}, \sigma_\eta)$. Note that $\omega \in \Omega$ denotes an arbitrary element of the parameter vector. This section provides details associated with the likelihood maximization procedure. Likelihoods are calculated via simulation, where we use $r = 1, \dots, R$ as the index for simulation draws.

In practice, we maximize the following function:

$$\mathcal{L} = \sum_{i=1}^N \ln \bar{P}_i(\Omega), \quad (18)$$

where the term $\bar{P}_i(\Omega) = (1/R) \sum_{r=1}^R P_{ir}(\Omega)$ is the simulated likelihood contribution for individual i .

The partial derivative of the likelihood function with respect to parameter ω is given by:

$$\frac{\partial \mathcal{L}}{\partial \omega} = \sum_{i=1}^N \frac{1}{\bar{P}_i(\Omega)} \underbrace{\left(\frac{1}{R} \sum_{r=1}^R \frac{\partial \ln P_{ir}(\Omega)}{\partial \omega} \cdot P_{ir}(\Omega) \right)}_{\bar{g}_i(\Omega)}, \quad (19)$$

where $\bar{g}_i(\omega)$ is the average partial derivative for individual i , and the average is calculated across simulation draws. Therefore, for a given parameter ω , the partial derivative of the likelihood can be written as:

$$\frac{\partial \mathcal{L}}{\partial \omega} = \sum_{i=1}^N \frac{\bar{g}_i(\Omega)}{\bar{P}_i(\Omega)}.$$

D.1.1 Log-Likelihood and Gradient

For individual i and simulation draw r , the per-draw log-likelihood is defined as:²⁵

$$\ell_{ir}(\Omega) \equiv \ln P_{ir}(\Omega) = \ln P_{A,ir}(A_i | \Omega) + \ln P_{E,ir}(E_i | \Omega), \quad (20)$$

where A_i is an application vector and E_i is the school that a child enrolls in. For parsimony, we suppress the conditioning on Ω in the expressions that follow.

Recall that Equation 8 in the main body of the paper defines the following term:

$$P_{A,ir}(A_i) = \frac{\exp(V_{ir}(A_i)/\lambda)}{\sum_{a'} \exp(V_{ir}(a')/\lambda)},$$

where the term $V_{ir}(A_i) \equiv e_{ir}(A_i) - c_{ir}(A_i)$ is based on Equations 5 and 6. In addition, note

²⁵Conditioning on data such as X_i, D_i, θ_i is suppressed from the following notation.

that Equation 9 in the main body defines the term:

$$P_{E,ir}(E_i = j) = \frac{\exp(v_{ijr})}{\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r})},$$

where E_i is the school in which a student enrolls, $\mathcal{O}_i$ is the school offer set, and the systematic utility term is defined as $v_{ijr} = \delta_j + \theta_{ir} + \beta_c X_i - \psi_D D_{ij} - \psi_{D2} D_{ij}^2 - \psi_{Dx} (D_{ij} \times X_i)$.

Here, $P_{A,ir}$ is the portion of the likelihood contribution that comes from the application stage, and $P_{E,ir}$ is the portion of the likelihood contribution that comes from the enrollment stage. At the application stage, families are weighing their expected utilities from an application, $e_{ir}(A_i)$, against the costs from applying, $c_{ir}(A_i)$, which we denote by V_{ir} in the notation above. The kernel smoothing factor λ allows us to convert the application choice into a multinomial logit problem following (Walters 2018).

Returning to Equation 20, we know that the per-draw log-likelihood for individual i is:

$$\ell_{ir}(\Omega) = \ln P_{ir}(\Omega).$$

The partial derivative with respect to parameter ω , conditional on a draw r , is:

$$\frac{\partial \ell_{ir}}{\partial \omega} = \frac{1}{\lambda} \left[\frac{\partial V_{ir}(A_i)}{\partial \omega} - \sum_a P_{A,ir}(a) \frac{\partial V_{ir}(a)}{\partial \omega} \right] + \frac{\partial \ln P_{E,ir}(E_i)}{\partial \omega}. \quad (21)$$

Note that A_i denotes the application vector submitted by i in data, and a denotes one of the possible application vectors that could be submitted by i .

The per-draw gradient vector $g_{ir}(\Omega)$ stacks these partial derivatives and has M -many entries based on the dimension of Ω :

$$g_{ir}(\Omega) = \begin{pmatrix} \frac{\partial \ell_{ir}}{\partial \omega_1} \\ \vdots \\ \frac{\partial \ell_{ir}}{\partial \omega_m} \\ \vdots \\ \frac{\partial \ell_{ir}}{\partial \omega_M} \end{pmatrix},$$

where the subscript m is the parameter index.

For each element of Ω , Equation (21) implies a decomposition into application and enrollment components. We therefore define $g_{ir}^A(\Omega)$ and $g_{ir}^E(\Omega)$ so that we have:

$$g_{ir}(\Omega) = g_{ir}^A(\Omega) + g_{ir}^E(\Omega),$$

where the m -th entries of $g_{ir}^A(\Omega)$ and $g_{ir}^E(\Omega)$ correspond to the first and second terms on the right-hand side of Equation (21), respectively. In this case, the simulated gradient contribution of individual i is:

$$\bar{g}_i(\Omega) = \frac{1}{R} \sum_{r=1}^R \left(g_{ir}^A(\Omega) + g_{ir}^E(\Omega) \right).$$

D.1.2 Application-Stage Gradient

For any feasible application vector a , we have that $V_{ir}(a) = e_{ir}(a) - c_{ir}(a)$. The partial derivative with respect to ω_m is:

$$\frac{\partial V_{ir}(a)}{\partial \omega_m} = \frac{\partial e_{ir}(a)}{\partial \omega_m} - \frac{\partial c_{ir}(a)}{\partial \omega_m}.$$

Based on Equation (21), we denote the m -th component of the application-stage gradient vector $g_{ir}^A(\Omega)$ as $g_{ir,m}^A$. This can be written as:

$$g_{ir,m}^A = \frac{1}{\lambda} \left[\frac{\partial e_{ir}(A_i)}{\partial \omega_m} - \sum_a P_{A,ir}(a) \frac{\partial e_{ir}(a)}{\partial \omega_m} \right] - \frac{1}{\lambda} \left[\frac{\partial c_{ir}(A_i)}{\partial \omega_m} - \sum_a P_{A,ir}(a) \frac{\partial c_{ir}(a)}{\partial \omega_m} \right], \quad (22)$$

where A_i denotes the application vector submitted by i in the data.

Note that if an individual chooses *not* to apply, then $V_{ir}(A_i) = 0$, so the terms involving $\partial e_{ir}(A_i)/\partial \omega_m$ and $\partial c_{ir}(A_i)/\partial \omega_m$ drop out. In that case,

$$g_{ir,m}^A = -\frac{1}{\lambda} \left[\sum_a P_{A,ir}(a) \frac{\partial e_{ir}(a)}{\partial \omega_m} \right] + \frac{1}{\lambda} \left[\sum_a P_{A,ir}(a) \frac{\partial c_{ir}(a)}{\partial \omega_m} \right]. \quad (23)$$

D.1.3 Enrollment-Stage Gradient

Based on Equation (9) in the main text, we can write the probability of enrolling at school j as:

$$P_{E,ir}(E_i = j) = \frac{\exp(v_{ijr})}{\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r})},$$

where j' is an offered school, and $\mathcal{O}_i$ denotes the offer set for student i . Our specification for utility assumes

$$v_{ijr} = \delta_j + \theta_{ir} + \beta_c X_i - \psi_D D_{ij} - \psi_{D2} D_{ij}^2 - \psi_{Dx} (D_{ij} \times X_i).$$

Note that the outside option is included in $\mathcal{O}_i$ and normalized appropriately.²⁶

For individual i and draw r , the enrollment-stage per-draw log-likelihood can be written as

$$\begin{aligned} \ell_{ir}^E(\Omega) &= \ln P_{E,ir}(E_i | \Omega) \\ &= \ln \left(\frac{\exp(v_{iEr})}{\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r})} \right) \\ &= v_{iEr} - \ln \left(\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r}) \right). \end{aligned}$$

The first term depends only on the utility index for the school that student i actually enrolls in (the observed $E_i = j$), while the second term depends on the utilities of all schools in the offer set $\mathcal{O}_i$.

²⁶For students without any offer, we treat enrollment as degenerate on the outside option; in that case the enrollment-stage log-likelihood contribution is constant and the gradient is zero.

For a scalar parameter $\omega_m \in \Omega$, the derivative of ℓ_{ir}^E is therefore:

$$\begin{aligned}\frac{\partial \ell_{ir}^E}{\partial \omega_m} &= \frac{\partial v_{iE_i r}}{\partial \omega_m} - \frac{\partial}{\partial \omega_m} \ln \left(\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r}) \right) \\ &= \frac{\partial v_{iE_i r}}{\partial \omega_m} - \frac{1}{\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r})} \sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r}) \frac{\partial v_{ij'r}}{\partial \omega_m}.\end{aligned}$$

We can simplify this expression by noting that the first term can be written as a sum over the offer set:

$$\frac{\partial v_{iE_i r}}{\partial \omega_m} = \sum_{j' \in \mathcal{O}_i} \mathbf{1}\{E_i = j'\} \frac{\partial v_{ij'r}}{\partial \omega_m},$$

where the indicator term is only non-zero when the index of the summation takes the value of the school that a child enrolls in (i.e., $j' = j$). It will also be helpful to define the following expression:

$$P_{E,ir}(j) \equiv \frac{\exp(v_{ijr})}{\sum_{j' \in \mathcal{O}_i} \exp(v_{ij'r})}.$$

Returning to the partial derivative, we can use the immediately preceding results to simplify:

$$\begin{aligned}\frac{\partial \ell_{ir}^E}{\partial \omega_m} &= \sum_{j' \in \mathcal{O}_i} \mathbf{1}\{E_i = j'\} \frac{\partial v_{ij'r}}{\partial \omega_m} - \sum_{j' \in \mathcal{O}_i} P_{E,ir}(j') \frac{\partial v_{ij'r}}{\partial \omega_m} \\ &= \sum_{j' \in \mathcal{O}_i} \left(\mathbf{1}\{E_i = j'\} - P_{E,ir}(j') \right) \frac{\partial v_{ij'r}}{\partial \omega_m}.\end{aligned}\tag{24}$$

The gradient vector $g_{ir}^E(\Omega)$ stacks these partial derivatives, with the m -th component being:

$$g_{ir,m}^E(\Omega) \equiv \frac{\partial \ell_{ir}^E}{\partial \omega_m}.$$

To more fully consider these expressions, consider school s and calculate $\partial \ell_{ir}^E / \partial \delta_s$. Since the parameter δ_s enters only the utility index for school s , we know that:

$$\frac{\partial v_{ij'r}}{\partial \delta_s} = \mathbf{1}\{j' = s\} \quad \text{for each } j' \in \mathcal{O}_i.$$

Substituting into Equation (24) yields:

$$\begin{aligned}g_{ir,\delta_s}^E &= \frac{\partial \ell_{ir}^E}{\partial \delta_s} = \sum_{j' \in \mathcal{O}_i} \left(\mathbf{1}\{E_i = j'\} - P_{E,ir}(j') \right) \mathbf{1}\{j' = s\} \\ &= \mathbf{1}\{E_i = s\} - P_{E,ir}(s).\end{aligned}$$

where only the term with $j' = s$ survives in the sum.

More generally, when a parameter ω_m enters v_{ijr} linearly and only for school s , we have

$$g_{ir,\omega_m}^E = \frac{\partial \ell_{ir}^E}{\partial \omega_m} = \left(\mathbf{1}\{E_i = s\} - P_{E,ir}(s) \right) \frac{\partial v_{isr}}{\partial \omega_m},\tag{25}$$

where the derivative $\partial v_{isr} / \partial \omega_m$ is given by the corresponding covariate or transformation. For

example, consider the cases:

$$\frac{\partial v_{isr}}{\partial \psi_D} = -D_{is}, \quad \frac{\partial v_{isr}}{\partial \psi_{D2}} = -D_{is}^2,$$

and the corresponding gradient components follow from Equation (25).

Finally, if a student receives no offer and enrolls in the outside option by construction, then $P_{E,ir}(E_i = 0)$ is identically one and does not depend on Ω . In this case, the enrollment-stage gradient contribution is:

$$\frac{\partial \ell_{ir}^E}{\partial \omega_m} = 0 \quad \text{for all } \omega_m.$$

D.1.4 Estimating a Mixture of Normals for Idiosyncratic Choice School Preferences

The likelihood contribution of individual i is:

$$\mathcal{L}(A_i, Z_i, E_i \mid X_i, D_i) = \int P_A(A_i \mid X_i, D_i, \theta, \eta) \times P_E(E_i \mid Z_i, X_i, D_i, \theta) dF(\theta, \eta \mid X_i, D_i),$$

where it is implicit that we must integrate out the unobserved taste and cost shocks θ_i and η_i (we suppress the i subscript on (θ, η) in the notation below).

In a model with a single normal component, we would assume $\theta \sim N(0, \sigma_\theta^2)$ and $\eta \sim N(0, \sigma_\eta^2)$, independent of (X_i, D_i) , so

$$dF(\theta, \eta \mid X_i, D_i) = \phi(\theta; 0, \sigma_\theta^2) \phi(\eta; 0, \sigma_\eta^2) d\theta d\eta,$$

where $\phi(x; \mu, \sigma^2)$ denotes a normal density for x with mean μ and variance σ^2 .

To allow for more flexible heterogeneity in tastes for choice schools, we generalize this to a finite mixture of normals. In the mixture model with K types, we have:

$$\theta \sim \sum_{k=1}^K p_k N(\mu_k, \sigma_k^2),$$

where p_k is the share of type k in the population and each type k has its own mean μ_k and standard deviation σ_k . We maintain a mean-zero normalization for θ ,

$$E[\theta] = \sum_{k=1}^K p_k \mu_k = 0,$$

to pin down the overall location of the taste distribution and avoid confounding the mean of θ_i with the school fixed effects $\{\delta_j\}$ in the utility index. This restriction imposes one linear constraint across the K component means. For example, we can express:

$$\mu_K = -\frac{\sum_{j=1}^{K-1} p_j \mu_j}{p_K},$$

so that there are $K - 1$ unrestricted mean parameters. We also estimate K standard deviations σ_k for each type. Keeping the assumptions on the distribution of η constant, the joint density

becomes:

$$dF(\theta, \eta \mid X_i, D_i) = \left(\sum_{k=1}^K p_k \phi(\theta; \mu_k, \sigma_k^2) \right) \phi(\eta; 0, \sigma_\eta^2) d\theta d\eta.$$

Under this mixture specification, the likelihood contribution of individual i is:

$$\begin{aligned} \mathcal{L}(A_i, Z_i, E_i \mid X_i, D_i) = \sum_{k=1}^K p_k & \left(\int P_A(A_i \mid X_i, D_i, \theta, \eta) \right. \\ & \times P_E(E_i \mid Z_i, X_i, D_i, \theta) \phi(\theta; \mu_k, \sigma_k^2) \phi(\eta; 0, \sigma_\eta^2) d\theta d\eta \Big). \end{aligned}$$

In the $K > 1$ case, we must also parameterize the type probabilities p_k . These probabilities must satisfy $p_k \geq 0$ for all k and $\sum_{k=1}^K p_k = 1$. Rather than estimate the p_k directly under these constraints, we follow a standard approach and work with an unconstrained parameter vector $(\alpha_1, \dots, \alpha_K) \in \mathbb{R}^K$ and map it into probabilities using the soft-max (multinomial logit) transformation:

$$p_k = \frac{\exp(\alpha_k)}{\sum_{k'=1}^K \exp(\alpha_{k'})}, \quad k = 1, \dots, K$$

This guarantees that each p_k is strictly positive and that the probabilities sum to one for any values of $(\alpha_1, \dots, \alpha_K)$.

The soft-max transformation is invariant to adding a constant to all components; that is, $p_k(\alpha_1 + c, \dots, \alpha_K + c) = p_k(\alpha_1, \dots, \alpha_K)$ for any scalar c . To obtain an identified parameterization, we therefore normalize one component, which we take to be $\alpha_K = 0$. Under this normalization, the probabilities can be written as:

$$p_k = \frac{\exp(\alpha_k)}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})},$$

and

$$p_K = \frac{1}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})}.$$

Thus, there are $K - 1$ free parameters governing the type probabilities.

Combining the $K - 1$ unrestricted means, the K standard deviations, and the $K - 1$ type-probability parameters, the mixture distribution for θ is characterized by $3K - 2$ free parameters.

D.1.5 Changes to the Gradient in the Mixture Case

With the mixture-of-normals specification for θ , the simulated choice probability for individual i is:

$$P_i(\Omega) = \sum_{k=1}^K p_k \bar{P}_{i,k}(\Omega).$$

In this expression, the term:

$$\bar{P}_{i,k}(\Omega) = \frac{1}{R} \sum_{r=1}^R P_{ir,k}(\Omega)$$

is the simulated probability for individual i under mixture component k , averaged over R simulation draws. The sample log-likelihood is:

$$\mathcal{L}(\Omega) = \sum_{i=1}^N \ln P_i(\Omega).$$

Consider a scalar parameter $\omega \in \Omega$ that does not parameterize the type probabilities p_k (i.e., ω is not one of the $\{\alpha_k\}$, so $\partial p_k / \partial \omega = 0$). The derivative of the log-likelihood with respect to ω is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \omega} &= \sum_{i=1}^N \frac{1}{P_i(\Omega)} \frac{\partial P_i(\Omega)}{\partial \omega} \\ &= \sum_{i=1}^N \frac{1}{P_i(\Omega)} \frac{\partial}{\partial \omega} \left(\sum_{k=1}^K p_k \bar{P}_{ik}(\Omega) \right) \\ &= \sum_{i=1}^N \frac{1}{P_i(\Omega)} \left(\sum_{k=1}^K p_k \frac{\partial \bar{P}_{ik}(\Omega)}{\partial \omega} \right) \\ &= \sum_{i=1}^N \frac{1}{P_i(\Omega)} \left(\sum_{k=1}^K p_k \frac{1}{R} \sum_{r=1}^R \frac{\partial P_{irk}(\Omega)}{\partial \omega} \right). \end{aligned}$$

Using the identity

$$\frac{\partial P_{irk}(\Omega)}{\partial \omega} = P_{irk}(\Omega) \frac{\partial \ln P_{irk}(\Omega)}{\partial \omega},$$

we can rewrite the inner sum as:

$$\frac{1}{R} \sum_{r=1}^R \frac{\partial P_{irk}(\Omega)}{\partial \omega} = \frac{1}{R} \sum_{r=1}^R P_{irk}(\Omega) \frac{\partial \ln P_{irk}(\Omega)}{\partial \omega} \equiv \bar{g}_{ik}(\Omega),$$

where $\bar{g}_{ik}(\Omega)$ is the simulated gradient contribution for individual i under mixture component k . Substituting back, we obtain:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \omega} &= \sum_{i=1}^N \frac{1}{P_i(\Omega)} \left(\sum_{k=1}^K p_k \bar{g}_{ik}(\Omega) \right) \\ &= \sum_{i=1}^N \frac{\bar{g}_i(\Omega)}{P_i(\Omega)}, \end{aligned}$$

where the term

$$\bar{g}_i(\Omega) \equiv \sum_{k=1}^K p_k \bar{g}_{ik}(\Omega)$$

is the mixture-weighted average of the K simulated gradient contributions for individual i .

Relative to the case of a single-normal model, the structure of the individual gradient contribution is qualitatively similar, but we must now keep track of the K component-specific simulated probabilities $\bar{P}_{ik}(\Omega)$ and their associated gradients $\bar{g}_{ik}(\Omega)$. The overall gradient is a probability-weighted average of these K simulated gradients for each individual.

D.1.6 Gradients with Respect to Type Probability Parameters in the Mixture Case

We now consider the gradient of the log-likelihood with respect to the type-probability parameters $\{\alpha_k\}_{k=1}^{K-1}$. Recall that the simulated probability for individual i under the mixture specification is:

$$P_i(\Omega) = \sum_{k=1}^K p_k \bar{P}_{ik}(\Omega),$$

where $p_k = p_k(\alpha_1, \dots, \alpha_{K-1})$ and $\bar{P}_{ik}(\Omega) = \frac{1}{R} \sum_{r=1}^R P_{irk}(\Omega)$ is the simulated probability for individual i under component k . The individual log-likelihood contribution is:

$$\ell_i(\Omega) \equiv \ln P_i(\Omega).$$

For a given α_k (with $k = 1, \dots, K-1$), the derivative of ℓ_i is:

$$\begin{aligned} \frac{\partial \ell_i}{\partial \alpha_k} &= \frac{1}{P_i(\Omega)} \frac{\partial P_i(\Omega)}{\partial \alpha_k} \\ &= \frac{1}{P_i(\Omega)} \frac{\partial}{\partial \alpha_k} \left(\sum_{k'=1}^K p_{k'} \bar{P}_{ik'}(\Omega) \right) \\ &= \frac{1}{P_i(\Omega)} \left(\sum_{k'=1}^K \frac{\partial p_{k'}}{\partial \alpha_k} \bar{P}_{ik'}(\Omega) + \sum_{k'=1}^K p_{k'} \frac{\partial \bar{P}_{ik'}(\Omega)}{\partial \alpha_k} \right). \end{aligned}$$

Only the mixture weights p_k and the component K contribution $\bar{P}_{iK}(\Omega)$ depend on the type-probability parameters. For $k < K$, we have $\partial \bar{P}_{ik'}(\Omega)/\partial \alpha_k = 0$, so we can simplify this expression to:

$$\frac{\partial \ell_i}{\partial \alpha_k} = \frac{1}{P_i(\Omega)} \left(\sum_{k'=1}^K \frac{\partial p_{k'}}{\partial \alpha_k} \bar{P}_{ik'}(\Omega) + p_K \frac{\partial \bar{P}_{iK}(\Omega)}{\partial \alpha_k} \right). \quad (26)$$

Derivatives of the mixture weights: To obtain $\partial p_{k'}/\partial \alpha_k$, recall that the multinomial-logit (soft-max) parameterization is:

$$p_k = \frac{\exp(\alpha_k)}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})} \quad \text{for } k = 1, \dots, K-1,$$

The normalization $\alpha_K = 0$ implies:

$$p_K = \frac{1}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})}.$$

Differentiating with respect to α_k (for a fixed $k \in \{1, \dots, K-1\}$) yields:²⁷

$$\begin{aligned} \frac{\partial p_k}{\partial \alpha_k} &= p_k(1 - p_k), \\ \frac{\partial p_{k'}}{\partial \alpha_k} &= -p_{k'}p_k \quad \text{for all } k' \neq k. \end{aligned}$$

²⁷These are the familiar soft-max derivatives that follow from applying the quotient rule to $p_k = \exp(\alpha_k)/(1 + \sum_{n=1}^{K-1} \exp(\alpha_n))$.

Derivatives of the mixture weights. To obtain $\partial p_{k'}/\partial \alpha_k$, recall that under the multinomial-logit (soft-max) parameterization with normalization $\alpha_K = 0$ we have

$$p_k = \frac{\exp(\alpha_k)}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})} \quad \text{for } k = 1, \dots, K-1,$$

and the normalization $\alpha_K = 0$ implies

$$p_K = \frac{1}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})}.$$

Differentiating with respect to α_k (for a fixed $k \in \{1, \dots, K-1\}$) yields:²⁸

$$\frac{\partial p_k}{\partial \alpha_k} = p_k(1 - p_k),$$

$$\frac{\partial p_{k'}}{\partial \alpha_k} = -p_{k'}p_k \quad \text{for all } k' \neq k$$

(where the second expression applies in particular to $k' = K$).

Derivatives of the K -th component contribution. The next term of interest in Equation (26) captures the dependence of the K -th component contribution $\bar{P}_{iK}(\Omega)$ on α_k through the mean-zero restriction on θ . Recall that

$$\bar{P}_{iK}(\Omega) = \frac{1}{R} \sum_{r=1}^R P_{irK}(\Omega),$$

and hence:

$$\frac{\partial \bar{P}_{iK}(\Omega)}{\partial \alpha_k} = \frac{1}{R} \sum_{r=1}^R \frac{\partial P_{irK}(\Omega)}{\partial \alpha_k} = \frac{1}{R} \sum_{r=1}^R P_{irK}(\Omega) \frac{\partial \ln P_{irK}(\Omega)}{\partial \alpha_k}.$$

The derivative of the per-draw log-probability can be decomposed into application and enrollment components:

$$\frac{\partial \ln P_{irK}(\Omega)}{\partial \alpha_k} = g_{ir,K}^A(\alpha_k) + g_{ir,K}^E(\alpha_k),$$

where $g_{ir,K}^A(\alpha_k)$ and $g_{ir,K}^E(\alpha_k)$ collect the terms arising from the application and enrollment stages, respectively, when the type is K .²⁹

The only channel through which α_k affects these terms is the normalized mean of the K -th type. Under the mean-zero restriction $E[\theta] = 0$ and the soft-max parameterization of the p_k ,

²⁸These are the familiar soft-max derivatives that follow from applying the quotient rule to $p_k = \exp(\alpha_k)/(1 + \sum_{n=1}^{K-1} \exp(\alpha_n))$.

²⁹As in Equation (22), $g_{ir,K}^A(\alpha_k)$ is built from derivatives of the application index $V_{ir}(a) = e_{ir}(a) - c_{ir}(a)$ with respect to α_k , and $e_{ir}(a)$ is the expected payoff from applying to portfolio a .

we can write:³⁰

$$\mu_K = - \sum_{k'=1}^{K-1} \exp(\alpha_{k'}) \mu_n,$$

so that:

$$\frac{\partial \mu_K}{\partial \alpha_k} = - \exp(\alpha_k) \mu_k. \quad (27)$$

In the application-stage expected utility for type K , this implies:

$$\frac{\partial e_{ir,K}(a)}{\partial \alpha_k} = \pi_{ia} P_{E,irK} \frac{\partial \mu_K}{\partial \alpha_k} = -\pi_{ia} P_{E,irK} \exp(\alpha_k) \mu_k.$$

Since the application-cost function $c_{ir}(a)$ does not depend on the type-probability parameters $\{\alpha_k\}$, the derivative of the application-stage index $V_{ir}(a) = e_{ir}(a) - c_{ir}(a)$ with respect to α_k is entirely driven by $\partial e_{ir,K}(a)/\partial \alpha_k$.

Substituting these derivatives into the general application-stage gradient expression in Equation (22) yields $g_{ir,K}^A(\alpha_k)$. A similar substitution into the enrollment-stage gradient expression in Equation (24), using $\partial v_{ijr}/\partial \alpha_k = \partial \mu_K/\partial \alpha_k$ for type K , yields $g_{ir,K}^E(\alpha_k)$.

Putting these pieces together, we can summarize the gradient of the individual log-likelihood with respect to α_k as

$$\frac{\partial \ell_i}{\partial \alpha_k} = \frac{1}{P_i(\Omega)} \left(\sum_{k'=1}^K \frac{\partial p_{k'}}{\partial \alpha_k} \bar{P}_{ik'}(\Omega) + p_K \bar{g}_{iK}(\alpha_k) \right),$$

where the last term in the parentheses is given by:

$$\bar{g}_{iK}(\alpha_k) \equiv \frac{1}{R} \sum_{r=1}^R P_{irK}(\Omega) [g_{ir,K}^A(\alpha_k) + g_{ir,K}^E(\alpha_k)],$$

which is the simulated gradient contribution for individual i arising from the impact of α_k on the K -th mixture component.

D.1.7 Gradients with Respect to Mixture Means

For the component means $\{\mu_m\}_{m=1}^{K-1}$, the application-stage per-draw gradient for type k with respect to μ_m follows the general form in Equation (22). For draw r , we can write:

$$g_{ir,k}^A(\mu_m) = \frac{1}{\lambda} \left[\frac{\partial e_{ir,k}(A_i)}{\partial \mu_m} - \sum_a P_{A,irk}(a) \frac{\partial e_{ir,k}(a)}{\partial \mu_m} \right],$$

³⁰Under the soft-max parameterization with normalization $\alpha_K = 0$, we have that $\mu_K = -(\sum_{k=1}^{K-1} p_k \mu_k)/p_K$. In addition, we have:

$$p_k = \frac{\exp(\alpha_k)}{D(\alpha)} \quad (k = 1, \dots, K-1), \quad p_K = \frac{1}{D(\alpha)},$$

where $D(\alpha) \equiv 1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'})$. Substituting these expressions into the μ_K above yields:

$$\mu_K = - \frac{\sum_{k=1}^{K-1} [\exp(\alpha_k)/D(\alpha)] \mu_k}{1/D(\alpha)} = - \sum_{k=1}^{K-1} \exp(\alpha_k) \mu_k.$$

where $e_{ir,k}(a)$ denotes the application-stage expected utility for type k and application vector a , and $P_{A,irk}(a)$ is the corresponding application probability for type k .

For types $k < K$, the expected utility $e_{ir,k}(a)$ contains a term that is linear in μ_k with coefficient $\pi_{ia}P_{E,irk}$, so the derivative with respect to μ_m is:

$$\frac{\partial e_{ir,k}(a)}{\partial \mu_m} = \begin{cases} \pi_{ia} P_{E,irk}, & \text{if } m = k, \\ 0, & \text{if } m \neq k. \end{cases}$$

For type K , the mean μ_K is not a free parameter but is pinned down by the mean-zero restriction $E[\theta] = 0$. As noted above, this implies:

$$\mu_K = -\frac{\sum_{k'=1}^{K-1} p_{k'} \mu_{k'}}{p_K},$$

so that for $m < K$,

$$\frac{\partial \mu_K}{\partial \mu_m} = -\frac{p_m}{p_K}.$$

In turn, the μ_K -dependent part of the expected utility for type K is proportional to $\pi_{ia}P_{E,irK} \mu_K$, so by the chain rule we obtain:

$$\begin{aligned} \frac{\partial e_{ir,K}(a)}{\partial \mu_m} &= \pi_{ia}P_{E,irK} \frac{\partial \mu_K}{\partial \mu_m} \\ &= -\pi_{ia}P_{E,irK} \frac{p_m}{p_K}, \quad m = 1, \dots, K-1. \end{aligned}$$

There is no free parameter μ_K , so we do not compute a gradient with respect to μ_K itself.

Substituting these derivatives into the general application-stage gradient expression above yields $g_{ir,k}^A(\mu_m)$ for each type k and mean parameter μ_m . The corresponding enrollment-stage gradients with respect to μ_m are analogous, obtained by replacing $e_{ir,k}(a)$ with the enrollment utility index v_{ijr} and applying the same derivative rules. For brevity, we do not display them here explicitly.

E Counterfactual Details

This appendix section provides details for our counterfactual analysis. Each counterfactual policy exercise involves simulating environments using our estimated demand and potential outcome models. The outcome model estimates are based on the linear control function approach (Equation 14). All simulations assume the supply of choice schools is held fixed to the observed years included in our sample (2004-2013). The specific alterations to the environment allow us to isolate the policy effect of interest relative to the existing opt-in (voluntary) system. For each counterfactual, we simulate $R = 100$ economies and report averages of statistics across simulations. We discuss the details for each counterfactual next.

E.1 Information Interventions

As mentioned in the main body of the paper, information interventions involve boosting a student’s θ_i by one standard deviation of their estimated type-specific distribution. Importantly, we only allow the shift in θ_i to affect application and enrollment decisions but use the original θ_i when calculating counterfactual potential outcomes.

Because changes to θ_i lead to changes in application rates relative to the baseline scenario, we must calculate equilibrium admission probabilities that students consider when making application decisions to choice schools. We follow the logic of Walters (2018) and Avery et al. (2025) to solve for equilibrium admission probabilities under the new environment. For each counterfactual simulation, $r = 1, \dots, R$, we solve for the equilibrium admission probabilities given the draw of idiosyncratic preference (θ_i) as well as their post-lottery (ξ_{ij} and cost shocks (η_i) that govern application and enrollment decisions. Once we solve for the equilibrium π_{ij}^{BR} for each school j , we then calculate the resulting statistics for the counterfactual simulation. We repeat this $R = 100$ times and average statistics across simulation rounds.

E.2 Busing Program

The busing policies that we consider are approximated in our model as an elimination of travel costs. As in the information interventions, this leads to a change in demand, so we must estimate equilibrium admission probabilities in this new environment. We follow the same procedure outlined above.

E.3 Decentralized

The decentralized market allows families to submit applications to multiple schools and potentially receive multiple offers, in contrast to our baseline model, which restricts households to a single application to mirror LAUSD policy during the sample period. To implement this counterfactual, we modify how application vectors are constructed. A key change is to augment the application cost structure by introducing marginal costs: whereas the baseline model includes only a fixed application cost, the decentralized setting requires a per-application marginal cost to prevent students from submitting excessively large portfolios. We follow the prior literature and calibrate this marginal cost to match the fixed-to-marginal cost ratio in Avery et al. (2025). In this environment, students form application portfolios by weighing the fixed cost of initiating the application process together with a constant marginal cost for each additional school they choose to include.

Because fixed application costs introduce non-convexities, we cannot directly apply the off-the-shelf marginal improvement algorithm from Chade and Smith (2006). To address this, we use a modified procedure: we first “seed” the algorithm by starting with the single school that delivers the highest stand-alone expected utility, and then apply the standard Chade-Smith marginal-improvement rule to consider adding additional schools to their portfolio. At the end

of the process, we compare the expected utility from the resulting application portfolio to the utility from not applying to any school.

To verify that this modified procedure produces the optimal application portfolio, we conduct Monte Carlo simulations comparing the algorithm’s output to brute-force enumeration of all feasible portfolios for a representative sample of students. In all cases, the modified greedy algorithm recovered the same optimal portfolio as the exhaustive search, confirming that the procedure correctly identifies the utility-maximizing application vector. As in our other counterfactuals, we then compute equilibrium admission probabilities implied by the demand induced by the transition to a decentralized market.

E.4 Mandatory Application

The mandatory application policy involves an application mandate and the introduction of a deferred acceptance mechanism to govern assignments. Importantly, we allow families to rank multiple schools. In practice, this means we eliminate application costs and require families to rank schools in order of preferences governed by the estimated demand parameters and their observables. Our use of an uncapped choice school list avoids having to estimate equilibrium admission probabilities. Therefore, we simulate $R = 100$ environments and average statistics across simulation rounds. We assume students enroll in their assigned school.

E.5 Mandatory Application and Complementary Policies

For counterfactuals that combine the mandatory application policy with other complementary reforms, we follow the procedures discussed above. Because each involves the use of a deferred acceptance mechanism with unrestricted list lengths, we do not estimate equilibrium admission probabilities. We assume students enroll in their assigned school.