

EQUILIBRIUM EFFECTS OF INFORMED SCHOOL CHOICE

By

Claudia Allende, Francisco Gallego and Christopher Neilson

COWLES FOUNDATION DISCUSSION PAPER NO. 2566



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

ABSTRACT

We study how information about school quality affects families' choices and how scaling up the intervention changes its impact. In a randomized trial in Chile, we offer families a short video and a personalized school report card before they enroll their children. Treated families choose schools with higher test scores and value added. Among families reached before enrollment, treated children score about 0.2 standard deviations higher in fourth grade, with positive but less precise estimates on later assessments. We embed the experimentally identified demand shift in a structural model of school choice and competition to study equilibrium effects at scale. Capacity constraints and centralized assignment attenuate the demand-side gains, but the same shift in demand strengthens schools' incentives to improve quality, and the supply response roughly offsets the congestion effect. More generally, equilibrium scaling can alter experimental effects in either direction because congestion and endogenous quality responses work against each other. We use the model to guide the design of a follow-up cluster-randomized experiment, focusing on the tradeoff between precision and contamination from equilibrium spillovers.

Claudia Allende
Stanford University
Graduate School of Business
and NBER
callende@stanford.edu

Christopher Neilson
Yale University
and Abdul Latif Jameel
Poverty Action Lab
and also NBER
christopher.neilson@yale.edu

Francisco Gallego
Pontificia Universidad Católica de Chile
and Abdul Latif Jameel
Poverty Action Lab
fgallego@puc.cl

The authors wish to thank Steve Berry, Ryan Cooper, Michael Dinerstein, Francisco Lagos, Chris Walters, and Román Andrés Zárate for useful comments and discussions. We thank the constructive comments from conference participants at the RESTUD Tour 40th Reunion, the NBER Summer Institute IO and Education group meetings and the Y-RISE inaugural conference, as well as seminar and workshop participants at MIT, UCL, UC San Diego, NYU, Princeton, University of Bergen, University of Oslo, PUC-Chile and Yale University. We thank Ryan Cooper for efficiently leading the project development at JPAL-LAC; Josefa Aguirre, Jorge Cariola, José I. Cuesta, Cristián Larroulet, and Cristián Ugarte for excellent research assistance during the fieldwork during 2009-2010; Magdalena Zahri, Ximena Poblete, and Francisca Zegers for help with the production of the information instruments; OPINA and EKHOS for field work; and FONDECYT (Grant No. 1100623) and Industrial Relations Section of Princeton University for financial support.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w35837>

1. INTRODUCTION

A growing body of evidence shows that providing families with information about school quality can shift their school choices (Hastings and Weinstein, 2008, Andrabi et al., 2017, Kapor et al., 2020, Campos, 2024, Corradini, 2024). But what happens when such a policy is implemented at scale? In a small-scale experiment, school characteristics are, to a first approximation, fixed, and each treated family's choice has negligible effects on market conditions. At scale, this approximation breaks down. Information can concentrate demand at particular schools, causing families to compete for a limited number of seats and changing schools' incentives to invest in quality. A small-scale experiment therefore identifies the direct choice response, but not the congestion and supply responses that emerge at scale. Whether these forces reinforce or undo the experimental gains determines what the small-scale experiment implies for implementing the information intervention more broadly.²

This paper studies these equilibrium effects in Chile, a well-studied setting with extensive school choice and a large private voucher sector. To do so, we first designed a school-choice information intervention with scale in mind, using a low-cost format that could be delivered through Chile's existing public preschool infrastructure. The intervention combines a short video with a personalized school report card: the video emphasizes the link between school quality and future educational and labor-market opportunities, while the report card compares nearby primary schools on test scores, prices, and location and includes a map. We examine the intervention's direct effects on school choice in a small-scale randomized controlled trial. We find that treated families choose schools with higher test scores and value added. These effects are concentrated among families reached before they had enrolled their child in a primary school, when the intervention could still influence the initial choice. These experimentally identified choice effects provide the demand-side inputs to the structural model. Separately, we follow children in administrative data to assess whether the changes in school choice were consequential for learning. Among families reached before enrollment, treated children score about 0.2 standard deviations higher on fourth-grade assessments. Point estimates remain positive in eighth grade and on the college entrance examination more than a decade later, but these

²The distinction between partial and equilibrium effects has long been recognized as important in education contexts (Heckman et al., 1998). Andrabi et al. (2017) show that information can induce equilibrium responses, including supply-side adjustment, in small village schooling markets in Pakistan, and Borger et al. (2024) replicate this market-level approach in rural Haiti. Experimental evidence on market-level responses in large urban school systems, however, remains limited.

later assessments are observed for smaller and potentially selected test-taking samples, so we interpret them as supporting evidence rather than as a separate long-run finding.

To move from the experiment to an equilibrium analysis at scale, we build on the demand model estimated by Neilson (2026).³ The baseline demand system features heterogeneous families choosing among schools based on quality, out-of-pocket price, and distance across 53 urban markets. Following the scaled-logit logic in Brown and Jeon (2024), we extend this system to allow the information intervention both to raise the weight on school characteristics targeted by the report cards and to make choices more responsive to school characteristics generally. We estimate the parameters governing these two responses by matching the experimental effects on chosen-school characteristics to their model-simulated counterparts using GMM. Once estimated, these parameters describe how information changes the way families evaluate the options they face. Because this response operates on each family’s own local alternatives, the model can carry it to families whose choice environments differ from those in the experiment.

The treatment-induced shift in demand also changes schools’ incentives to invest in quality. As families become more responsive to quality, schools’ quality markdowns shrink and the return to improving quality rises. We build the supply analysis on the quality first-order condition derived by Neilson (2026). That paper calculates schools’ markdowns and relates changes in them to the observed quality response following the SEP reform (*Subvención Escolar Preferencial*), but does not estimate the underlying supply-side cost parameters. We extend this analysis in two directions. First, we use a ten-year administrative panel to estimate the model’s composite quality-cost terms. Second, we derive a separate condition for schools near capacity. These schools cannot respond to stronger demand primarily by expanding enrollment, but can benefit when higher quality changes the composition of students and therefore per-student revenue. We estimate the unconstrained and capacity-composition conditions jointly. Identification combines reform-induced variation in school revenue with local labor-market shifters, including non-education earnings and a new inverse-distance-weighted index of nearby public-school teacher quality. In the policy counterfactual, centralized assignment imposes school capacities, and private schools adjust quality according to the anchored quality FOC, while public-school quality remains fixed.

³Related work on school choice and competition in education markets includes Neilson (2026), Dinerstein and Smith (2015), Walters (2018), Kapor et al. (2020). The model in this paper also connects to earlier work on school choice in Chile by Gallego and Hernando (2008) and Chumacero et al. (2011).

We build the counterfactual in three steps. The first is a demand-only benchmark that holds school characteristics fixed and ignores capacity constraints. We apply the experimentally identified demand shift at full take-up to all eligible families (types 1–4) in the 53-market population—not only to families in the experimental sample—while preserving each family’s socioeconomic characteristics, location, and local school options. The resulting population average need not equal the experimental-sample estimate because families in the two populations face different distributions of school quality, prices, and distance. The detailed spatial data also allow the model to recover the distribution of effects across families and locations, rather than only an average effect. In this benchmark, the demand shift raises the expected quality (value added) of the schools attended by low-income families by 0.191 SD. Adding capacity constraints attenuates the effect by about two-fifths because of congestion at popular schools (from 0.191 to 0.115 SD). The final layer allows incumbent private schools to adjust quality, while holding entry, marginal costs, capacity, and public-school quality fixed. This defines our medium-run equilibrium approximation. The supply response roughly offsets the attenuation from congestion, producing an equilibrium effect of 0.226 SD for SEP families. This estimate is statistically different from zero, but it is only 0.036 SD above, and not statistically distinguishable from, the demand-only benchmark. The mechanism behind this reversal is that more quality-responsive demand raises schools’ returns to improving quality. More generally, scaling a demand-shifting policy in a market with finite capacity and endogenous quality activates two opposing forces: congestion reallocates access to existing quality, while stronger competitive incentives can induce providers to produce more of it. Their relative magnitude is an empirical property of the market, so the direction of the equilibrium correction cannot be inferred from the small-scale experimental effect alone.

We make three contributions to the literature. First, we provide experimental evidence on families’ responses to school-choice information and study the policy’s equilibrium consequences in a large urban market. Andrabi et al. (2017) provide the closest precedent, showing that school report cards in Pakistan’s small village markets raised achievement and induced supply-side responses. In that setting, report cards reach entire local markets, so market-level variation arises naturally. Large urban markets, like those we study in Chile, pose a different challenge: schools draw from spatially dispersed populations and families face many alternatives, so demand shifts diffuse across a wide geography rather than concentrating in a few schools. An experimental design that saturates some urban markets while preserving others as credible control groups may not be feasible, which is why the structural model is needed to study

the policy at scale. Related work by Bobba et al. (2026) studies information about admission chances in Mexico City’s centralized assignment to public high schools, holding school characteristics fixed; we instead study information about school quality and endogenous provider responses in a mixed public–private market.

Second, we show how a demand shift identified in a small-scale experiment can inform a prospective equilibrium policy evaluation. Our approach lies between two familiar uses of structural models. In one, the policy has already been implemented at scale and the model decomposes observed effects (e.g., Neilson, 2026, Dinerstein and Smith, 2015). In the other, the entire counterfactual rests on the model, with no experimental variation to pin down the demand response. Here, the experiment identifies a credible demand shift that is observed in the data, but the full policy environment—congestion, assignment constraints, supply-side responses—must be studied through the model, because these forces emerge only when the treated population is large enough to change market conditions. A comparison with Neilson (2026) makes the distinction concrete: that paper studies Chile’s SEP voucher reform after it was implemented at scale and uses the model to decompose an observed equilibrium response. For the information policy studied here, no such implementation exists. We combine the experimentally identified demand shift with administrative capacities and a transparent markdown-based provider response to construct a prospective policy calculation, and use the historical supply estimates and cost sensitivities to interpret it.⁴

Third, we show how the estimated model can inform the design of a larger-scale experiment that would measure these equilibrium effects directly. One middle ground between randomizing entire markets and randomizing individual families is to define smaller geographic clusters within markets and assign them independently to treatment or control. The challenge is that equilibrium responses propagate through schools shared across cluster boundaries, so treated and control clusters need not generate sharply separated exposures. In Section 7, we use solved no-treatment and full-treatment equilibria to construct a transparent endpoint-exposure approximation and study how cluster size, provider-core size, and buffer zones shape attenuation and precision in a cluster-randomized trial in Santiago. The exercise identifies a stable preference

⁴Todd and Wolpin (2023) survey the literature combining experiments with structural models for policy evaluation, including cases where experiments estimate or validate structural parameters (Todd and Wolpin, 2006, Walters, 2018). A related approach to ours appears in Lise et al. (2015), who combine experimental variation with equilibrium considerations in the context of employment.

for larger geographic clusters while showing that the exact core choice depends on how provider responses are interpolated between the equilibrium endpoints.

Taken together, the paper offers a way to assess demand-shifting policies whose broad implementation could change both access and supply. The approach applies more broadly when an experiment identifies individual responses, but broad implementation would change market access or provider incentives.

The rest of the paper proceeds as follows. Section 2 sets out the conceptual framework. Sections 3–5 describe the setting, experimental evidence, and empirical model. Section 6 reports the equilibrium counterfactuals, Section 7 applies the model to the design of a larger experiment, and Section 8 concludes.

2. CONCEPTUAL FRAMEWORK

The experimental evidence. The experiment and the policy exercise answer different questions. Random assignment identifies how the information intervention changes the characteristics of the primary school a family chooses. We estimate this response among families reached before enrollment, when the intervention could still affect the initial choice. It is a partial-equilibrium object in a precise sense: the experiment was small enough that school quality, prices, and seat availability were, to a first approximation, fixed. The experiment therefore does not identify what happens when many families respond at once. The model extrapolates choices rather than achievement. The fourth-grade results show that the induced changes in school choice were followed by learning gains, but no achievement coefficient enters the model.

The experimental effect is also an average over local environments. Each report card described the schools near a particular family. Two families with the same socioeconomic characteristics and voucher eligibility but living in different neighborhoods therefore received different information. One may have learned about substantial quality differences among schools available at no cost; another may have faced nearby schools of similar quality, with better alternatives only at considerably greater distance. The experiment identifies the average response across the environments represented in its sample.

From the experimental sample to population choices. The object carried to the target population is not the experimental average change in chosen-school quality. It is the estimated change in how families evaluate school attributes. Section 5 parameterizes the intervention as changing the weights on the characteristics the report card targeted—quality, out-of-pocket price, and

distance—and, potentially, the overall responsiveness of choices to school attributes. We estimate those parameters by requiring model-simulated effects on chosen-school characteristics to match the corresponding experimental effects.

Once estimated, the treatment parameters operate on each family’s own choice problem. For the target population of first-grade entrants in the 53 model markets, administrative data provide the distribution of family types, residential locations, and available schools, while the baseline demand estimates describe how those families choose among quality, price, and distance. The main counterfactual applies the intervention at full take-up to family types 1–4—families in which the mother has at most completed high school—and leaves the two higher-education types untreated. The maintained transport assumption is that the demand shift estimated within SEP status applies to families of the same status in the target markets. Their choice environments, by contrast, are taken from the population data rather than imposed from the experiment.

This distinction explains why experimental subgroup analysis cannot replace the model. The relevant environment is the configuration of quality, prices, and distances for all schools a family might attend—a multidimensional spatial object that no single subgroup indicator summarizes. It also clarifies the appropriate model-fit requirement. The model must reproduce the experimental effects in the experimental markets, which we verify in Section 6. It need not reproduce the same average effect in the target population, where the estimated behavioral change operates on a different distribution of alternatives.

Equilibrium forces. Two additional forces arise when many families respond, and neither is identified by the experiment. First, seats are finite. If informed families concentrate their demand on the same schools, assignment rather than preferences alone determines who attends. Quantifying this congestion channel requires school capacities and an assignment mechanism. We use capacities imputed from administrative records and a deferred-acceptance assignment rule representing the centralized policy environment.

Second, schools may adjust quality. More quality-responsive demand changes the return to providing quality, but the experiment does not identify that supply relationship. Building on the quality first-order condition in Neilson (2026), we use a ten-year administrative panel to estimate historical quality conditions and composite cost terms, including a separate condition for schools near capacity. For the prospective benchmark, we use that evidence to motivate a common anchored-FOC system: each private school’s competitive-quality term is fixed at its observed-equilibrium value, and schools’ qualities are solved jointly as treatment changes their

quality markdowns. The experiment disciplines the demand shift, while the historical panel motivates the markdown mechanism; a separate sensitivity rescales marginal quality costs proportionally. Because congestion restricts access to existing quality while stronger incentives can induce schools to produce more quality, the two forces work in opposite directions. Section 6 reports them in sequence rather than presenting only their net effect.

Scope of the counterfactual. The resulting equilibrium predictions remain conditional on the maintained structure. The medium-run counterfactual holds school entry, tuition, capacity, school-specific competitive-quality terms, and public-school quality fixed. Assignment allows student composition to change, and the historical supply model allows SEP composition to affect school incentives through targeted voucher revenue. The benchmark quality counterfactual, however, uses the common anchored FOC rather than switching schools into a separate capacity-composition equation. It also does not allow peer composition itself to change measured school quality or enter family utility or school objectives directly. Section 6 discusses these restrictions and reports selected sensitivities. Finally, the solved no-treatment and full-treatment quality equilibria provide endpoints for the prospective design exercise in Section 7, which studies how cluster geometry affects attenuation and simulated precision under transparent alternative exposure mappings.

3. SETTING, INTERVENTION, AND DATA

3.1. *Institutional Setting*

Primary schools in Chile are either free public schools, private voucher schools, or private non-voucher schools. The system features a high degree of school choice and a large private sector: in 2016, the market share for first-grade students was 36% for public schools, 55% for voucher schools, and 8% for private schools. Voucher schools may charge copayments, but these are capped and interact with targeted subsidies for disadvantaged students. In particular, in 2008 the government introduced the SEP reform (*Subvención Escolar Preferencial*), which increased the voucher for the poorest 40% of students. Participating schools were required to charge zero tuition to eligible families and were prohibited from screening applicants, including on academic grounds. The 2009 General Education Law extended the prohibition on selection more broadly across subsidized elementary schools. The 2015 Inclusion Law subsequently froze existing copayments and began replacing them with public funding. Chile later replaced school-by-school applications with a centralized application system that lets families rank multiple

schools in a common process. When a school is oversubscribed, assignments follow centrally specified priorities and lottery numbers. Centralization made applications and the allocation of scarce seats more transparent; the prohibition on selection predates it.⁵


In spite of the variety of schools and choices available, students from poorer families tend to attend schools with lower academic outcomes and lower inputs. Figure 1 shows model-predicted baseline expected school quality for the structural experiment sample by SEP eligibility, which proxies for socioeconomic status. One reason for this gap is that poorer families may place less weight on academic quality when choosing schools, either because quality is harder to observe or because the returns to choosing a better school are less salient. Chilean policymakers have long published school performance data, but this information is rarely tailored to the schools in a family's immediate vicinity or timed to reach parents when they are actively choosing where to enroll. Our intervention targets both gaps.

We reach families through the public preschool system. Pre-K attendance is widespread before primary school entry, with net enrollment rates of 55% at age 3 and 87% at age 4 (OECD, 2016). We partnered with *Fundación Integra*, a national network of approximately 1,000 public preschool centers serving predominantly low- and lower-middle-income families throughout the country. We delivered the intervention at preschool parent meetings, reaching families during the period when they are actively deciding which primary school to enroll their child in for the following year. At the time of our experiment (2010), each school managed its own admissions and timing varied across schools. Under the current centralized enrollment window, the timing of information provision would be more predictable in a scaled-up policy. The Online Appendix provides additional institutional detail on preschool provision.

3.2. *Information Intervention and Experimental Design*

The intervention was designed to make school choice more salient and to provide parents with simple, comparable information about nearby schools. Its goal was to increase both awareness of local school characteristics and the perceived returns to school quality. The treatment was delivered at regular *Integra* parent-teacher meetings and combined a short video with a personalized report card. The video emphasized that choosing a school carefully can shape a child's future opportunities: it linked school quality to preparation for higher education and better jobs, and featured role models from low-income neighborhoods who had accessed better opportuni-

⁵See Neilson (2026) for a detailed description of the market structure and policy environment.



Notes: Distribution of model-predicted baseline expected school quality (normalized VA) for the 815 families in the structural experiment sample, evaluated in their 2011 model choice environments. Each family receives equal weight within its SEP group. SEP families are eligible for Chile’s targeted voucher (*Subvención Escolar Preferencial*), which proxies for low socioeconomic status. Vertical lines mark group means.

ties through education (Figure O-2 and Figure O-3). Each family also received a report card listing test scores, prices, and locations of nearby primary schools, together with a map to help parents locate and compare options (Figure O-1).

We implemented the study in the three largest regions of Chile: Valparaíso, Biobío, and Santiago. Preschools were eligible if they were in urban areas, had at least 10 primary schools within 2 km, and were located in municipalities with at least two primary schools per preschool. These criteria were dictated by the intervention: each family received a personalized comparison of nearby schools, so the experiment required a setting with enough local alternatives for the report card to present a meaningful choice. The criteria yielded 143 eligible preschools, which we randomly assigned to treatment and control before arranging visits, stratifying by region, the number of grades offered, and a measure of local competition given by the number of primary schools within 2 km. Ten could not be visited because of principal refusal, scheduling conflicts, or low expected attendance, leaving 133 participating preschools. The original design distinguished between two treatment intensities, but imperfect tracking during implementation prevents us from separating the two arms. We therefore pool both versions and interpret the estimates as intent-to-treat effects of assignment to the bundled information package. Additional details on sample selection, including a comparison of participating and nonparticipating eligible centers, treatment materials, and implementation appear in Online Appendix Sections O-2 and O-3. Figure O-5 maps the spatial distribution of treatment and control preschools and

student homes in Santiago, and Figure O-6 shows how treated students are distributed across destination schools.

The experiment was implemented between August and December 2010 by trained staff who attended preschool parent meetings. In the 133 preschools that agreed to participate, 1,832 parents signed informed consent forms and completed a baseline survey before any information was provided. The baseline survey collected contact information and information on the application process, including whether the parent expected the child to enter primary school in 2011, whether a school had already been chosen, and whether the family had older children already attending primary school. Parents in treatment preschools then received the intervention. Because our presence at the meeting was not announced in advance, parents could not attend in anticipation of receiving the intervention. Between May and July 2011 we conducted a follow-up survey on enrollment decisions and successfully re-interviewed 1,612 of the 1,832 baseline respondents (88%). We also matched 1,795 of the 1,832 parents in the original sample to administrative records using individual student identifiers (98%).

3.3. *Data, Outcomes, and Analysis Samples*

We use several complementary data sources. First, we use administrative data on preschools from *Integra*, including location, enrollment, attendance, and socioeconomic composition. Second, we use self-collected data from the baseline and follow-up surveys in the experimental preschools. These data include contact information, individual identifiers, family location, and questions regarding the application process. We match this information to administrative data from Chile's Ministry of Education.

The Ministry of Education data include student-by-year matriculation records, information on grade progression and demographics, eligibility for the SEP voucher, and standardized test scores from SIMCE. We complement these data with birth records from the Ministry of Health, which contain information on child health at birth and on parental background, and with administrative records on all schools in the country, including school type, enrollment, address, tuition, and other characteristics. For the supply analysis, we also use a school-year panel with transfers, expenditures, and measures of teacher and principal quality.⁶

Different parts of the paper use different slices of these data. The reduced-form analysis uses the experimental cohort treated in 2010 and their observed primary-school choices in 2011;

⁶For further details, see the Online Appendix and Neilson (2026).

some children entered their chosen school before first grade. We take the baseline demand preference parameters and the school–market–year mean utilities for 2006, 2007, 2011, and 2012 from Neilson (2026), who estimates them using administrative records on the full population of first-grade choices and school market shares in 53 urban markets. To extend the school panel for the supply estimation, we hold the demand parameters fixed and recover mean utilities for 2008–2010 and 2013–2015 by inverting observed market shares at those fixed parameters. The Integra eligibility rules therefore do not enter either market-share inversion. Our second-step treatment estimation uses the subset of experimental families who had not yet enrolled and can be matched to the six urban markets shared with the baseline demand model. Holding the baseline school attributes and mean utilities fixed, this step estimates how the intervention changes the weights families place on observed school characteristics and the scale of choice behavior so as to match the experimental changes in chosen-school characteristics. The supply estimation uses the resulting 53-market urban school-year panel for 2006–2015, and the main counterfactuals are simulated in the 2011 market environment, when the experimental cohort entered primary school. Online Appendix O-15.1 summarizes how these samples map into each empirical exercise, and Table O-8 reports summary statistics for the main estimation variables.

Our main short-run outcomes are the characteristics of the schools families choose, including distance, fees, test scores, and value added. Our longer-run outcomes are student academic outcomes measured in administrative test-score and attainment records. Because the intervention was delivered during the period in which parents were enrolling children in primary school, timing matters for interpretation. The follow-up survey records whether enrollment had already been completed by the preschool meeting; families reporting that it had been completed had less scope to respond to the treatment. We therefore report results for both the full sample and the subgroup reporting that enrollment was not yet complete at the meeting, which is also the sample used for the structural treatment-effect estimation. Treatment and control are balanced on most measured characteristics within these subsamples, but not all: the exceptions, including child sex and home ownership, are reported in the balance tables and included in the standard expanded family-control block throughout the reduced-form analysis. Detailed balance tables, additional evidence on enrollment timing, and further discussion of treatment implementation are reported in Table O-5, Table O-6, Table O-7, and Online Appendix Section O-4.

We use the school-level test-score value-added measure developed in Neilson (2026) as our main proxy for academic quality, extending its coverage as needed for the empirical exercises in this paper. This measure captures the school’s contribution to later student achievement after

conditioning on a rich set of student and school characteristics. We show in the Online Appendix that it is strongly correlated with teacher quality—in particular, with teachers’ own scores on Chile’s national university entrance exam (PSU, now PAES)—and that rankings are similar whether or not prior scores are available as controls. The report cards given to treated families displayed raw test scores (SIMCE averages), not value added, since raw scores are publicly available and easy for parents to interpret. The two measures are strongly correlated across schools (Figure O-7), so the quality information provided to parents is closely aligned with the measure used in the structural model. The estimation equation and validation exercises are presented in Online Appendix Section O-7.

4. EXPERIMENTAL RESULTS

Table I reports effects on the characteristics of the schools families chose, and Table II reports effects on student outcomes. All specifications include the randomization controls and an expanded family-control block: child sex, SEP priority eligibility, maternal education and marital status, household size, durable-goods ownership, and home ownership, with missing-value indicators. The same block is used for every outcome and does not change the estimation samples. The Online Appendix reports results under alternative control sets and documents control coverage.

Because the intervention was delivered while families were choosing a primary school, we report results for the full sample and by reported enrollment status at the preschool meeting. Panel B comprises families whose children were already enrolled; Panel C comprises those whose children were not yet enrolled. Enrollment status is observed for 1,571 of the 1,780 children in the analysis sample, so the pooled specification uses a slightly larger sample than Panels B and C combined.

The experimental evidence on school choice is concentrated among families reached before they had committed to a school (Panel C of Table I). In this group, treated families travel about 0.2 km farther, are 12.3 percentage points more likely to choose a fee-charging school, and choose schools with stronger academic profiles.⁷ The clearest shifts are in chosen-school value added and school-average fourth-grade mathematics scores, with a weaker effect on school-average second-grade language scores and an imprecise one on school-average fourth-grade

⁷For the distance outcome we restrict to home-to-school distances below 4 km, which excludes families who appear to have moved between the intervention and the following school year. Results are robust to alternative distance thresholds.

TABLE I
EFFECTS ON CHARACTERISTICS OF SCHOOLS CHOSEN ONE YEAR AFTER TREATMENT

	Distance (kms) (1)	Positive Price (2)	Lang 2nd (3)	Lang 4th (4)	Math 4th (5)	V. Added ^a (6)
Panel A: Full Sample						
Treatment	0.063 (0.062)	0.050 (0.034)	0.009 (0.020)	0.013 (0.024)	0.021 (0.026)	0.035 (0.030)
Control mean	1.081	0.494	-0.119	-0.095	-0.082	-0.024
N obs.	1,427	1,701	1,693	1,692	1,692	1,701
Panel B: Already enrolled at the time of the PreK visit						
Treatment	-0.161 (0.133)	0.000 (0.055)	-0.026 (0.041)	-0.021 (0.047)	-0.042 (0.056)	-0.025 (0.058)
Control mean	1.192	0.589	-0.031	-0.007	0.046	0.067
N obs.	512	578	576	577	577	578
Panel C: Not enrolled at the time of the PreK visit						
Treatment	0.198*** (0.073)	0.123*** (0.040)	0.052* (0.028)	0.043 (0.029)	0.082** (0.033)	0.099*** (0.037)
Control mean	1.010	0.434	-0.179	-0.137	-0.146	-0.067
N obs.	800	934	928	926	926	934

Notes: ITT estimates from regressions of chosen-school characteristics on treatment assignment. Panel A: full sample; Panel B: families already enrolled at the time of treatment; Panel C: families not yet enrolled. All specifications include randomization controls and the expanded family controls listed in Appendix A1. Standard errors in parentheses are clustered at the preschool-center level. ^a Column 6 uses the 2011 school value-added measure Q_2 from Neilson (2026). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

language scores. The value-added effect—about 0.10 in the normalized quality measure used in the model—is the object most directly relevant for the structural analysis. Among families who had already enrolled (Panel B), the estimates are imprecise and none is statistically significant; taken together, they do not show a systematic shift in chosen-school characteristics after the decision had already been made. In the full sample (Panel A), the estimates are attenuated accordingly.

Reported enrollment status refers to a decision made before treatment, but the resulting subsamples need not be balanced on observables. Most measured characteristics are balanced within each subgroup, with some exceptions in Panel B, including child sex and home own-

ership. The expanded family controls account for these observed differences and improve precision. Treatment does not predict enrollment-status classification or its observability.⁸ The not-yet-enrolled group is also the feasible policy margin under Chile’s current centralized admissions system: its common application calendar provides a predictable point at which families can receive information before submitting their rankings. The structural analysis therefore focuses on this group.

Table II examines whether the changes in school choice translated into better academic outcomes. Among not-yet-enrolled families, fourth-grade SIMCE scores are 0.24 SD higher in language and 0.21 SD higher in mathematics. The eighth-grade estimates are 0.18 and 0.26 SD, respectively. The treatment does not affect reaching twelfth grade or participation in the college entrance exam. Conditional on taking the PAES, the language and mathematics estimates are 0.19 and 0.15 SD.

The later outcomes are observed for smaller and potentially selected test-taking samples. Students remain observable in administrative enrollment records, but not all students sit each assessment, and the 2020 eighth-grade SIMCE was cancelled for one cohort. Treatment does not predict score observability, either overall or within the enrollment-status groups. Nevertheless, because test participation selects the observed samples and substantially reduces their size, we interpret the eighth-grade and PAES estimates cautiously.⁹

The structural analysis uses the treatment-induced school-choice response among not-yet-enrolled families; achievement coefficients do not enter its estimation or counterfactuals.

Finally, we examine choices in relation to the report card’s design features. As reported in the Online Appendix, treated families are less likely to choose a listed school. Conditional on choosing a listed school, the effect on choosing a green-coded school is small and statistically insignificant, while the high-ranked-school indicator increases among families who had not yet enrolled. This conditional association does not separately identify a causal effect of highlight-

⁸Adding older-sibling status leaves the estimated school-choice effects essentially unchanged; Online Appendix O-2.1 reports the full comparison.

⁹The long-horizon estimates for already-enrolled families in Panel B are imprecise. Neither PAES coefficient is statistically significant. Eighth-grade mathematics is positive and significant, but it is estimated on 243 children drawn almost entirely from one cohort and is the only one of the 14 Panel-B school-choice and student-outcome coefficients that rejects zero at the 5-percent level. Under 14 independent tests, the probability of at least one rejection is 51 percent; the Holm-adjusted p -value for the surviving estimate is 0.319. The Online Appendix reports alternative control sets and detailed coverage.

TABLE II
EFFECTS ON STUDENT OUTCOMES

	SIMCE Test Scores				Reached 12th Grade (5)	College Entrance Exam (PAES)		
	4th Grade		8th Grade			Took PAES (6)	(Cond. on taking PAES)	
	Lang (1)	Math (2)	Lang (3)	Math (4)			Lang (7)	Math (8)
<i>Panel A: Full Sample</i>								
Treatment	0.096* (0.052)	0.109** (0.054)	0.034 (0.082)	0.235** (0.094)	-0.004 (0.029)	0.027 (0.030)	0.166** (0.072)	0.133** (0.053)
Control mean	-0.222	-0.145	-0.059	-0.103	0.728	0.592	-0.367	-0.225
N obs.	1,583	1,584	610	615	1,775	1,775	1,083	1,060
<i>Panel B: Already enrolled at the time of the PreK visit</i>								
Treatment	-0.082 (0.108)	-0.119 (0.098)	-0.195 (0.229)	0.450** (0.193)	-0.036 (0.049)	0.013 (0.054)	0.165 (0.109)	0.159 (0.098)
Control mean	-0.039	0.058	0.169	0.044	0.770	0.613	-0.287	-0.145
N obs.	543	533	241	243	596	596	373	364
<i>Panel C: Not enrolled at the time of the PreK visit</i>								
Treatment	0.241*** (0.071)	0.209*** (0.068)	0.184 (0.112)	0.256** (0.115)	0.002 (0.036)	0.029 (0.031)	0.187** (0.082)	0.147** (0.071)
Control mean	-0.339	-0.252	-0.188	-0.162	0.706	0.579	-0.447	-0.284
N obs.	857	865	301	301	975	975	581	568

Notes: ITT estimates from regressions of student outcomes on treatment assignment. Panel A: full sample; Panel B: families already enrolled at treatment; Panel C: families not yet enrolled. Columns 1–2: fourth-grade SIMCE; Columns 3–4: eighth-grade SIMCE; Columns 5–6: reaching twelfth grade and college entrance exam participation; Columns 7–8: PAES scores, conditional on taking the exam. The eighth-grade samples consist almost entirely of the 2012 entry cohort because the 2020 assessment was cancelled. All specifications include the controls listed in Appendix A1; standard errors are clustered at the preschool-center level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ing: school rank and visual presentation were not independently randomized, and choosing a listed school is itself affected by treatment.

In sum, personalized information shifts families toward higher-quality schools and raises student achievement among families reached before enrollment. Estimates on later assessments remain positive, although they are observed for smaller samples selected by test participation. The treatment-induced shifts in chosen school quality, price, and distance provide the experimental targets for the structural analysis in the next section.

5. EMPIRICAL MODEL FOR POLICY ANALYSIS

This section develops the structural model. We adopt the baseline demand system—utility specification, market shares, and family types—from Neilson (2026) and summarize it briefly in subsection 5.1. Our main contribution on the demand side is a treatment-effects parameterization (subsection 5.2) that allows the information intervention to shift school choice through characteristic-specific channels and a common rescaling of choice behavior. A Bayesian signal model (subsection 5.3) provides one interpretation of the scale parameter, following Brown and Jeon (2024). On the supply side (subsection 5.4), we recover cost parameters from schools’ quality first-order conditions and introduce, for historical supply estimation, a capacity extension that changes the moment condition contributed by schools that cannot easily expand enrollment. subsection 5.5 describes the three-step estimation.

5.1. Baseline Demand Model

Following Neilson (2026), the utility for family i from school j is

$$U_{i,j} = \bar{\beta}x_j + \xi_j + \beta_i q_j + \alpha_i \text{op}_{k(i),j} + \lambda_i d_{\text{loc}(i),j} + \epsilon_{i,j}, \quad (1)$$

where q_j is school academic quality (value added), $\text{op}_{k(i),j}$ is the out-of-pocket payment for a family of type k at school j given the prevailing voucher policy, and $d_{\text{loc}(i),j}$ is distance from the family’s residential location to school j . The vector x_j collects observable school characteristics over which families share common preferences $\bar{\beta}$: administration type, for-profit status, religious affiliation, and whether the school serves grades K through 12. The term ξ_j is the standard BLP-style unobserved school attribute, and $\epsilon_{i,j}$ is i.i.d. extreme value type 1. The distinction between the school’s posted tuition p_j and the family’s out-of-pocket payment $\text{op}_{k(i),j}$ matters for the supply side: schools set a single tuition, while out-of-pocket payments depend on how that tuition interacts with the voucher schedule; in particular, SEP-eligible families at SEP-participating schools pay zero out of pocket.¹⁰

Preferences are heterogeneous across six discrete family types $k \in \{1, \dots, 6\}$, defined by mother’s education and family income (Neilson, 2026). Quality preferences include a type-specific component β_k and a standard-normal draw $z_i \sim N(0, 1)$, so that $\beta_i = \beta_{k(i)} + \sigma z_i$, while price and distance coefficients are type-specific: $\alpha_i = \alpha_{k(i)}$ and $\lambda_i = \lambda_{k(i)}$. We use signed price

¹⁰We use “location” for readability; these are the same discrete residential nodes used in Neilson (2026).

coefficients: a negative α_k means that a higher out-of-pocket payment lowers utility. We group families by SEP eligibility: types $\{1, 3\}$ are poor and eligible for Chile’s targeted voucher (*Subvención Escolar Preferencial*), while types $\{2, 4, 5, 6\}$ are not.

Families choose the school that maximizes $U_{i,j}$ from the set of schools in their market m . The market share of school j aggregates over types, residential locations, and the random coefficient:

$$s_j(\mathbf{q}, \mathbf{op}, \theta) = \sum_{k=1}^K \Pi_k^m \sum_{\text{loc}}^{L^m} w_k^{\text{loc}} \int s_j^{\text{loc},k}(\theta, z) d\Phi(z), \quad (2)$$

where Π_k^m is the proportion of type k families in market m , w_k^{loc} is the within-type distribution across locations, and $s_j^{\text{loc},k}(\theta, z)$ is the standard logit choice probability conditional on type, location, and standard-normal draw, with Φ denoting the standard-normal distribution. We take the estimated preference parameters and the mean utilities for the demand-estimation years from Neilson (2026). For the additional school-years needed in the supply panel, we hold those preference parameters fixed and recover mean utilities—and hence ξ_j —by inverting observed market shares at those fixed parameters.

5.2. Treatment Effects Parameterization

The information treatment may affect school choice through multiple channels. Providing families with school-specific report cards can shift the weight they place on school quality, price, and distance. It can also sharpen choice behavior more broadly, increasing families’ responsiveness to all school attributes. We identify both channels by matching the randomized treatment effects on chosen-school characteristics to their model-implied counterparts, allowing the experimental variation to determine their relative importance.

Define the baseline (control) deterministic utility as

$$V_{ij}^C = \tilde{\delta}_j + \bar{\phi}_k^{q,C} q_j + \bar{\phi}_k^{op,C} \text{op}_{ij} + \bar{\phi}_k^{d,C} d_{ij} + \bar{\sigma}^C z_i q_j, \quad (3)$$

where $\delta_j \equiv \bar{\beta}^q q_j + \bar{\beta}^C x_j + \bar{\xi}_j^C$ denotes the full recovered school mean utility from Neilson (2026), and $\tilde{\delta}_j \equiv \delta_j - \bar{\beta}^q q_j$ removes its common quality component. Here $\bar{\phi}_k^{q,C}$ is the full type-specific quality coefficient, including $\bar{\beta}^q$, while $\bar{\phi}_k^{op,C}$ and $\bar{\phi}_k^{d,C}$ are the signed price and distance

coefficients. The draw z_i is standard normal and $\bar{\sigma}^C$ is its baseline quality-taste dispersion.¹¹ The treated utility adds two components:

$$V_{ij}^T = \kappa_k \cdot V_{ij}^C + \psi_k^q q_j + \psi_k^{op} \text{op}_{ij} + \psi_k^d d_{ij}. \quad (4)$$

The first term rescales all baseline utility components by a common factor κ_k —including residual mean utility $\tilde{\delta}_j$, which contains the common preferences $\bar{\beta}^C x_j$ and the unobserved school attribute $\bar{\xi}_j^C$. The ψ_k^c terms add characteristic-specific treatment effects on top.

The model restricts treatment effects on the coefficients of other school attributes, including sector and residual school attributes, to the common rescaling:

$$\bar{\beta}_k^T = \kappa_k \cdot \bar{\beta}_k^C. \quad (5)$$

For attributes targeted by the information intervention—quality, out-of-pocket price, distance—the effective coefficient shifts through both the common scale and a characteristic-specific treatment effect:

$$\bar{\phi}_k^{c,T} = \kappa_k \cdot \bar{\phi}_k^{c,C} + \psi_k^c. \quad (6)$$

Because Equation 5 rules out separate shifts in these coefficients, moments on voucher-school attendance and residualized mean utility help identify κ_k (subsection 5.5.1). When $\kappa_k = 1$, the model reduces to characteristic-specific treatment effects with no common scaling, nesting the standard (unscaled) specification.

Interpretation. The treatment-effect parameters ψ_k^c are summary parameters that capture characteristic-specific shifts in how much characteristic c influences school choice, beyond common rescaling. Each ψ_k^c combines two mechanisms that are observationally equivalent in our setting. The first is an *information channel*: more precise signals raise the weight families place on the corresponding characteristic in expectation, even if underlying preferences are unchanged (Brown and Jeon, 2024). The second is a *preference channel*: by making quality differences salient and easy to compare, the report cards may cause families to place more weight on academic quality independent of any change in signal precision. The experiment identifies the

¹¹Under the standard logit normalization, these estimates represent structural preference parameters divided by the error scale. subsection 5.3 develops this point formally; the normalization is innocuous for the baseline demand but becomes consequential when treatment may change the scale.

combined effect of these channels on school choice; separately identifying them would require data on subjective beliefs.

For the policy counterfactuals in this paper, separately identifying these underlying channels is not necessary. The counterfactual exercises apply the same treatment observed in the experiment to eligible families (types 1–4) at full take-up, and simulate the resulting equilibrium. The demand shifts we estimate are identified directly from experimental variation, so the estimated total shift—however it decomposes across the information and preference channels—is the relevant sufficient statistic for these simulations.

To allow for heterogeneous treatment effects, we estimate separate ψ_k^c parameters and scale parameters κ_k for SEP-eligible families ($k \in \{1, 3\}$) and non-SEP families ($k \in \{2, 4\}$). The full treatment parameter vector is

$$\theta_2^T = \left(\psi_{\text{SEP}}^q, \psi_{\text{SEP}}^{op}, \psi_{\text{SEP}}^d, \psi_{\text{nSEP}}^q, \psi_{\text{nSEP}}^{op}, \psi_{\text{nSEP}}^d, \kappa_{\text{SEP}}, \kappa_{\text{nSEP}} \right). \quad (7)$$

We incorporate the treatment into the demand model by creating *treated types* in addition to the baseline six types from Neilson (2026). The experimental sample includes families of types 1–4 (types 5–6 have too few observations in the experiment). Treated families have utility V_{ij}^T from Equation 4, while control families retain the baseline utility V_{ij}^C .

5.3. Imperfect Information and the Scaled Logit

The parameterization in subsection 5.2 can be interpreted through a Bayesian signal model of school characteristics. When families observe school quality, price, and distance with noise, the posterior mean for each characteristic is attenuated toward the prior by a signal-to-noise ratio $\rho_k^c \in (0, 1)$ that depends on signal precision. The effective preference weight on characteristic c becomes a product of the structural preference and the signal precision: $\phi_k^c = \beta_k^c \cdot \rho_k^c$. Crucially, as emphasized by Brown and Jeon (2024), signal noise also enters the composite error term of the choice model. Improving signal precision on specific characteristics can therefore change not only the weight on those characteristics but also the error variance of the entire choice model. We adopt a scaled logit approximation that captures this possibility while preserving the logit functional form.¹²

¹²This variance-matching approximation is standard in the heteroskedastic logit literature; see (Train, 2009), Chapters 2 and 6. It is most accurate when the EV1 taste shock dominates the signal noise, a condition consistent with our empirical finding that the estimated scale correction is modest (κ not significantly different from 1).

Dividing the expected utility by the type-specific error scale μ_k (derived in the appendix from the composite error variance) recovers standard logit choice probabilities with *effective* parameters that embed both preferences and signal precision. The baseline demand estimation in Neilson (2026) identifies only these effective parameters under the standard normalization $\mu_k^C = 1$. This normalization is innocuous for the baseline demand but becomes consequential when treatment may change the scale. Within this framework, the common scale parameter from subsection 5.2 is

$$\kappa_k \equiv \frac{\mu_k^C}{\mu_k^T}. \quad (8)$$

When improved signal precision reduces the composite error variance, $\mu_k^T < \mu_k^C$ and $\kappa_k > 1$: the composite error shrinks, sharpening the weight on all school attributes, including those about which the treatment provides no new information. The model nests the standard logit at $\kappa_k = 1$, i.e., when treatment does not change the error variance. We do not interpret $\kappa_k > 1$ as parents literally learning about every residual school characteristic; rather, it is a parsimonious way to allow for any common sharpening of choice behavior beyond the characteristic-specific treatment shifts. Appendix A3 provides a simplified type-level illustration, including the posterior, expected utility, composite error, and variance-matching formulas.

5.4. Supply Side

We model privately owned schools as choosing quality q_j to maximize a weighted objective that reflects both profit motives and other institutional goals. On the supply side, the relevant price variable is the school's posted tuition p_j (the sticker price). Family out-of-pocket payments op_{ij} are derived from that posted tuition after applying the voucher schedule and therefore vary by student type and school participation status. Revenue per student depends on the sticker price and the voucher schedule $v_j^k(p_j)$, where j indexes schools and k indexes student type. In the baseline demand system, posted tuition is treated as an endogenous observed school characteristic: schools that invest in unobserved quality attributes may also set higher prices, inducing a correlation between tuition and the unobserved school attribute ξ_j . The demand estimation addresses this through instrumental variables, so the recovered ξ_j already accounts for this correlation. For the supply-side exercise, we take tuition and the recovered ξ_j as given and do not model a pricing first-order condition. This does not imply that historical tuition was exogenous: observed prices reflect schools' prior choices and therefore substantial path

dependence. The institutional case for fixing future prices is instead that subsidized private schools operated under regulated shared financing, and the 2015 Inclusion Law subsequently froze existing copayments and began phasing them out. The observed tuition vector is therefore a natural medium-run benchmark for a prospective information policy. The Online Appendix displays the statutory schedule and the distinction between posted tuition and families' out-of-pocket payments; see also Sanchez (2017). We hold observed tuition fixed and focus on the quality first-order condition as the relevant adjustment margin. The per-student cost of providing quality q_j is

$$C_{jt}(q_{jt}) = \gamma_{\text{SEP}} \text{SEP}_{jt} + \gamma_{mt}^{\text{type}(j)} + c_{jt} q_{jt}, \quad c_{jt} = c_j + \delta_{\text{type}(j)}(t - 2006). \quad (9)$$

Here SEP_{jt} indicates program participation, $\gamma_{mt}^{\text{type}(j)}$ is a market-year effect specific to voucher or private non-voucher schools, c_j is a school-specific unit cost of quality, and $\delta_{\text{type}(j)}$ allows that unit cost to trend differently by school type.¹³ The first-order condition for optimal quality at an unconstrained school yields

$$q_{jt}^* = \underbrace{\left[\frac{R_{jt} - \gamma_{\text{SEP}} \text{SEP}_{jt} - \gamma_{mt}^{\text{type}(j)}}{c_{jt}} \right]}_{\text{Competitive quality}} - \underbrace{s_{jt} \left[\frac{\partial s_{jt}}{\partial q_{jt}} \right]^{-1}}_{\text{Quality markdown}}, \quad (10)$$

where R_{jt} is observed revenue per student (the regular subsidy, SEP subsidy, and posted tuition, as applicable). This expression decomposes optimal quality into a competitive component and a markdown reflecting market power. The markdown $s_{jt}/(\partial s_{jt}/\partial q_{jt})$ is inversely related to the semi-elasticity of demand with respect to quality: schools with more elastic demand face stronger competitive pressure and provide quality closer to the competitive level. Neilson (2026) derives this quality FOC and uses it to interpret the observed response to the SEP reform. We estimate this condition on our extended panel, allowing a capacity and revenue-composition condition for schools near capacity, as described below. The counterfactual response rule in Section 6 is built on the same markdown term.

¹³The quality index q_j recovered from the demand estimation is identified only up to differences. In the supply estimation, we normalize q_j so that its minimum is zero and impose a small positive floor on c_{jt} .

5.4.1. *Interpreting Supply-Side Parameters*

Schools in the data belong to three administrative sectors: municipally managed public schools, privately managed voucher schools, and private non-voucher schools. Within the voucher sector, an administrative indicator records whether the operator's legal form was for-profit or nonprofit; the latter include foundations and religious organizations, including Catholic providers. This was a substantive legal distinction during our estimation period; the 2015 Inclusion Law subsequently required subsidized private providers to become nonprofit. Legal form need not, however, reveal the provider's complete objective. Following Hoxby (2003), we distinguish the distribution of operating surplus from the incentive to generate resources that can be used by the school.

To make this distinction precise, let operating surplus be

$$\Pi_j(q_j) = n_j(q_j)[R_j - C_j(q_j)] - F_j.$$

A for-profit provider may maximize Π_j . If a nonprofit provider instead values both operating surplus and educational quality delivered to its students,

$$V_j(q_j) = a_j \Pi_j(q_j) + (1 - a_j) n_j(q_j) B_j(q_j), \quad a_j > 0,$$

then its problem can be written in the same form as surplus maximization after replacing input cost with the effective cost

$$\tilde{C}_j(q_j) = C_j(q_j) - \frac{1 - a_j}{a_j} B_j(q_j).$$

When the mission payoff is per student, linear in quality, and unchanged by the information policy, direct concern for quality changes the net marginal cost of providing it but leaves the form of the quality FOC in Equation 10 unchanged. The school-specific term c_j can therefore reflect persistent differences in both input costs and schools' objective functions. Type-specific trends allow this term to evolve differently for voucher and private non-voucher schools, and the supply instruments are interacted with the administrative for-profit and nonprofit-voucher indicators.

For the historical supply estimation, we write the quality FOCs around the observed equilibrium. Equivalently, we use observed quality to recover the school-specific unobservable that

rationalizes each school's choice. Capacity determines which moment condition applies in the estimation panel: unconstrained schools can attract additional enrollment, whereas schools near capacity can change the composition of enrollment. Under either regime, the moments use variation in the marginal return to quality while absorbing persistent differences in input costs and objective functions.

This argument does not cover every possible nonprofit objective. A mission payoff unrelated to the number of students served, depending directly on student composition, determined by a binding break-even constraint, or changing with the information policy generally adds terms to the FOC that need not remain fixed when demand changes. We therefore do not interpret c_j as a deep technological parameter. Separately identifying schools' production technology and objective functions is beyond the scope of this paper.

Public schools are different institutionally. Municipalities manage portfolios of schools, can pool resources and supplement per-pupil transfers, and face rigidities in adjusting teaching staff. The school-level budget constraint is therefore softer than for private providers (Gallego, 2013). We exclude public schools from supply estimation and hold their quality fixed in the main counterfactual, treating this as a scope restriction rather than assuming that a private-provider FOC describes municipal decisions.

5.4.2. *Capacity Constraints and the Supply-Side FOC*

The baseline FOC in Equation 10 corresponds to the standard case in which a school can expand enrollment when higher quality raises demand. This is the same approximation used in the baseline demand system estimated by Neilson (2026), who argues that capacity constraints were not a first-order feature of household choice in the historical estimation period.¹⁴ We maintain that approximation for the baseline demand inputs, while allowing a subset of school-year observations to contribute a capacity-composition condition to the historical supply estimation. In the policy exercise, capacities instead enter the assignment stage directly.

Capacity-constrained schools may still have incentives to improve quality through a *composition channel*. When a school near capacity raises its quality, it need not expand total enrollment to benefit: by changing which families apply and accept available seats, it can instead attract

¹⁴Neilson (2026) argues in the Online Appendix that this approximation is reasonable in the baseline data because primary-level selection by voucher schools was illegal during the estimation period, only 2% of urban primary schools reached the legal class-size cap, and only 2% of surveyed parents reported preferring another school but being turned away.

a more SEP-intensive student body. Since SEP students generate higher per-student revenue ($v_b + v_{\text{SEP}}$ versus $v_b + p_j$), this compositional shift creates a financial incentive for quality improvement even when the standard enrollment channel is muted.

Formally, for a school operating at capacity $\bar{c}_j$, the first-order condition for quality becomes

$$c_{jt} = \Delta R_{jt} \cdot \frac{\partial s_{jt}^{\text{SEP}} / \partial q_{jt}}{s_{jt}}, \quad (11)$$

where $\Delta R_{jt} = R_{jt}^{\text{SEP}} - R_{jt}^{\text{nonSEP}}$ is the per-student revenue gap between SEP and non-SEP students, and $\partial s_{jt}^{\text{SEP}} / \partial q_{jt}$ is the derivative of the SEP-weighted share with respect to quality. The right-hand side captures the marginal revenue gain from improving composition: higher quality attracts more SEP students (who generate ΔR_{jt} more revenue each), scaled by the fraction of the student body affected. Under the linear cost function, the marginal cost of quality is c_{jt} , so the FOC equates this marginal cost to the marginal composition benefit.

In practice, we classify each school-year observation as operating under one of the two regimes—unconstrained or capacity-constrained—based on observed enrollment relative to capacity, and estimate the supply parameters using both FOCs simultaneously with the school-specific cost parameter c_j shared across regimes. Section 5.5.5 describes the classification rule and estimation details. The policy counterfactual does not dynamically switch between these historical estimation equations. It anchors the quality FOC at the observed equilibrium and solves jointly for private-school quality as described in Section 6, while the assignment stage separately imposes school capacities.

5.5. Estimation

Our estimation proceeds in three steps: (i) recovering baseline demand parameters from Neilson (2026), (ii) estimating treatment effect parameters using the experimental variation, and (iii) recovering supply-side cost parameters.¹⁵ We begin with an informal discussion of identification.

5.5.1. Identification

Baseline demand. The baseline demand parameters $\theta = \{\bar{\beta}, \alpha, \beta, \lambda, \sigma, \xi\}$ are identified by three complementary sources of variation in Neilson (2026): the BLP inversion (Berry et al.,

¹⁵See Appendix A4 for additional estimation details and standard error formulas.

1995) recovers full school-level mean utilities δ_j by matching predicted to observed market shares; micro-moments matching the average quality, out-of-pocket payment, and distance chosen by each family type identify heterogeneity across the six family types (Petrin, 2002); and IV moments address the endogeneity of price and quality using policy-driven revenue variation from the voucher and SEP reforms and local labor-cost shifters that proxy for the cost of hiring skilled workers. Appendix A4 lists these moments and instruments explicitly. As discussed in subsection 5.3, the baseline estimates identify the *effective* parameters that embed the normalization $\mu_k^C = 1$.

Treatment effects on informed characteristics. The ψ_k^c parameters — capturing how much treatment shifts the weight on quality, price, and distance beyond the common scale effect — are identified by the random assignment of the information treatment. Preschool centers were randomly assigned to treatment or control, stratified by region, grades offered, and the number of primary schools within 2 km. Treatment effects compare families across those assigned centers; preschool-municipality fixed effects and market controls are regression adjustments. The moments match these experimental treatment effects, separately for each characteristic and type group.

Scale parameter. The scale parameters κ_g and characteristic-specific shifts are jointly disciplined by the full system of 12 moment entries. Voucher-school attendance and residualized mean utility $\tilde{\delta}_j$ provide information about choices along dimensions whose coefficients have no separate treatment shift. Changes in quality, price, or distance weights also change the probabilities of choosing schools with these attributes. Together with the treatment effects on informed characteristics (q , op , d), these moments discipline the 8-parameter vector; no single moment identifies κ_g in isolation.

Supply. Supply parameters are identified from the first-order conditions for quality (Equation 10 and Equation 11). SEP participation is the observable cost shifter in Equation 9; it is distinct from the instruments used to identify its coefficient. Those instruments carry over the excluded variation used in the baseline demand estimation—simulated transfer changes, reform timing interacted with pre-policy neighborhood SEP exposure, and local non-education earnings deviations—and add an inverse-distance-weighted index of nearby public-school teacher quality. The implementation also interacts the instrument matrix with for-profit and nonprofit-voucher indicators. School-specific quality costs and market-year-type effects are concentrated

out, leaving γ_{SEP} and the voucher- and private-school cost trends as the three freely estimated parameters.

5.5.2. Step 1: Baseline Demand Parameters

The baseline demand parameters are taken from Neilson (2026), who estimates the first-step demand system using administrative first-grade choice data from 2006, 2007, 2011, and 2012. We refer the reader to that paper for a detailed description of the estimation strategy and the resulting parameter estimates.

5.5.3. Estimation Sample

The treatment effect parameters are estimated using a subsample of the experimental families that can be matched to the market-level demand model. We begin with the full experimental sample and apply a series of restrictions, each motivated by the requirements of the structural estimation. First, we restrict attention to families that had *not yet enrolled* their children at the time of the Pre-K visit, since the treatment effects on school choice are concentrated entirely among these families (section 4). Second, we restrict to the six urban markets that overlap between our experiment and the market structure in Neilson (2026), which is necessary because the baseline demand parameters—mean utilities δ_j , quality measures, and spatial choice sets—are only available within Neilson’s market definitions. Third, we apply a geographic consistency filter: we keep families if they remain in the same municipality as at baseline or if the chosen school is within 4 km of the family’s home, excluding only observations that are both in a different municipality and implausibly far from the chosen school.

We then drop families with missing information on the school they ultimately chose (value-added, out-of-pocket price, or distance). Finally, family types are assigned using six categories defined by mother’s education and SEP eligibility; because the experiment contains few families of types 5–6 (technical and college-educated), we collapse these into types 3–4 using actual JUNAEB SEP eligibility records to ensure correct classification.

Table III reports the treatment effects on chosen school characteristics in the resulting estimation sample, separately for SEP-eligible and non-SEP families. These are the experimental moments that the structural model must match. The treatment significantly increases the value-added of schools chosen by SEP families (0.220, $p < 0.05$); the point estimate on driving distance is also positive (0.134 km) though not statistically significant. These treatment effects form the empirical targets for the GMM estimation described below.

TABLE III
TREATMENT EFFECTS IN THE MODEL SAMPLE (EXPERIMENT MOMENTS)

	Quality	Sticker Price	OOP Price	Distance	Voucher	N
Pooled	0.195*** (0.073)	0.039** (0.017)	0.027 (0.017)	0.118 (0.078)	0.114*** (0.042)	815
SEP	0.220** (0.087)	0.025 (0.017)	0.006 (0.016)	0.134 (0.092)	0.106** (0.043)	505
Non-SEP	0.114 (0.104)	0.048 (0.036)	0.048 (0.036)	0.184 (0.119)	0.096 (0.067)	310
Preschool-municipality FE	Yes	Yes	Yes	Yes	Yes	

Notes: Experimental moments used in the demand estimation. Regressions use the model-overlap sample described in Section 5.5.3, preschool-municipality fixed effects, and standard errors clustered at the preschool-center level; exact controls are summarized in Appendix A2. Quality is normalized value added Q_2 . Distance is Google Maps driving distance, restricted to observations within 4 km. Voucher indicates a private voucher rather than a public school. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.5.4. Step 2: Treatment Effect Parameters

Conditional on the Step 1 estimates, we estimate the treatment parameter vector θ_2^T defined in Equation 7 using GMM. The identification strategy exploits the random assignment of the information treatment within the experimental sample.

We define *RCT moments* that match the treatment effects on school characteristics observed in the experimental data. For each characteristic $c \in \{q, p^{\text{sticker}}, \text{op}, d\}$ and each type group $g \in \{\text{SEP}, \text{nonSEP}\}$, the moment condition requires the model-simulated treatment effect to match the experimental estimate:

$$G_{c,g}(\theta_2^T) = \hat{\beta}_{c,g}^{\text{RCT}} - \hat{\beta}_{c,g}^{\text{sim}}(\theta_2^T), \quad (12)$$

where $\hat{\beta}_{c,g}^{\text{RCT}}$ is the coefficient on the treatment indicator from a regression of characteristic c on treatment status (with preschool-municipality fixed effects and the controls listed in Appendix A2) in the experimental data, and $\hat{\beta}_{c,g}^{\text{sim}}(\theta_2^T)$ is the analogous coefficient computed from model-simulated school choices. The simulated choices are generated using the treated and control probabilities implied by the model, with the random coefficient integrated via Gauss-Hermite quadrature. To help identify the scale parameters κ_g , we add moments on voucher-school attendance and residualized mean utilities $\tilde{\delta}_j \equiv \delta_j - \bar{\beta}^q q_j$, both of which capture treatment effects on already-known school attributes. Together, the full system of moments is used

to estimate all eight parameters (seven freely estimated in the benchmark specification, with $\psi_{\text{SEP}}^{\text{op}}$ restricted to zero as discussed in the results section).

We estimate using two-step efficient GMM: the first step uses the identity weighting matrix, and the second step uses $W = \hat{V}_{\text{CR}}^{-1}$, where $\hat{V}_{\text{CR}}$ is the cluster-robust covariance matrix of the data moments (clustering at the preschool center level). Standard errors are computed using the standard GMM sandwich formula.

5.5.5. Step 3: Supply Parameters

Given the demand parameters from Steps 1 and 2, we estimate the supply-side cost parameters using GMM on the quality first-order conditions. We classify a school-year observation as capacity-constrained when its enrollment-to-capacity ratio exceeds 0.95 and the school receives differential SEP revenue. This classification uses observed enrollment and imputed capacity and should therefore be understood as a transparent empirical rule rather than an exact measure of when capacity begins to bind. In our preferred specification, 2,414 of 22,882 school-year observations (10.5%) are classified as constrained, so the two-regime extension modifies the supply problem for a meaningful but limited subset of the sample.

For each school-year observation, we compute the quality markdown $\text{MD}_{jt} = s_{jt}/(\partial s_{jt}/\partial q_j)$ from the estimated demand system and define the “competitive quality” residual:

$$\text{CompQ}_{jt} \equiv q_{jt} + \text{MD}_{jt} = \frac{R_{jt} - \gamma_{\text{SEP}} \text{SEP}_{jt} - \gamma_{mt}^{\text{type}(j)}}{c_{jt}}, \quad (13)$$

where $c_{jt} = c_j + \delta_{\text{type}(j)}(t - 2006)$ and $\gamma_{mt}^{\text{type}(j)}$ absorbs common shocks within each market, year, and school type. This is the left-hand-side object implied by rearranging the unconstrained quality FOC; it is not a direct measure of “quality investment” or mean utility. For capacity-constrained observations, the estimation instead uses the composition-based condition in Equation 11.

We concentrate out c_j and $\gamma_{mt}^{\text{type}(j)}$ jointly. The three freely estimated parameters are γ_{SEP} , which captures the additive cost differential associated with SEP participation, and the voucher- and private-school trends in the unit cost of quality. The estimation sample consists of voucher and private non-voucher schools observed over 10 years (2006–2015) across 53 markets (22,882 school-year observations, 2,389 schools). Public schools are excluded for the reasons discussed in Section 5.4.1.

The moment conditions interact the regime-specific structural residual Ω_{jt} with the instrument matrix. This matrix carries over the baseline demand instruments and predetermined school characteristics, adds interactions with for-profit and nonprofit-voucher indicators, and includes a new Nearby Teacher Quality Index—the inverse-distance-weighted average of public-school teacher PSU scores within each market-year. Table O-15 documents the relationship between these instruments and observed school quality; the table is an instrument-relevance diagnostic rather than the supply GMM equation itself. We estimate using two-step optimal GMM. We view this two-regime supply specification as an extension of the baseline model rather than a fully unified model of demand, assignment, and capacity: it uses the same latent-demand primitives as the rest of the paper, but lets the marginal incentive to raise quality depend on whether a school appears to have slack enrollment capacity.

5.6. *Parameter Estimates*

Table IV presents the demand-side parameter estimates. Panel A reports the baseline demand parameters from Neilson (2026); Panel B reports the treatment effect parameters estimated in this paper. The baseline model is defined over six family types based on mother’s education and SEP eligibility. Our counterfactual applies the intervention only to types 1–4—families in which the mother has at most completed high school, separately by SEP eligibility—and leaves the two higher-education types untreated. All four treated types are substantially represented in the 815-family structural experiment sample: types 1–4 contain 265, 92, 240, and 218 families, respectively, with both treatment and control observations in every type (in this sample, the few families with technical- or college-educated mothers are grouped with types 3–4, as described in subsubsection 5.5.3). Because treatment heterogeneity is summarized more parsimoniously by SEP status, the experimental demand shift is identified from 505 SEP families (types 1 and 3) and 310 non-SEP families (types 2 and 4).

The experimental type shares need not equal their population shares. The counterfactual combines the experimentally estimated demand shift with the target population’s own distribution of family types, residential locations, and available schools. Administrative data for an adjacent Integra cohort provide a useful support check: 90.1 percent of Integra children entering first grade in the 53 model markets belong to types 1–4, and each of the four types is well represented. The exercise therefore changes the weights and choice problems assigned to family types observed in the experimental population; it does not extend the treatment to the two higher-education types absent from the policy counterfactual. The maintained transport assump-

TABLE IV
DEMAND MODEL ESTIMATES

	<i>Panel A: Baseline Demand (Neilson 2026)</i>				<i>Panel B: Treatment Effects (this paper)</i>		
	Quality (β^q)	Price (α)	Distance (λ)		Quality (ψ^q)	Price (ψ^{op})	Distance (ψ^d)
<HS	1.40 (0.03)	−2.70 (1.10)	−1.29 (0.05)	SEP	0.500 (0.230)	—	0.427 (0.204)
HS	0.56 (0.10)	−0.56 (0.38)	−1.09 (0.04)	nonSEP	0.399 (0.289)	−0.148 (0.067)	0.061 (0.269)
Tech/2-Year	0.86 (0.16)	−0.25 (0.38)	−1.04 (0.04)				
College+	1.13 (0.22)	−0.00 (0.38)	−0.96 (0.04)				
SEP	−0.33 (0.08)	−1.47 (0.31)	−0.06 (0.02)	κ_{SEP}	1.410 (0.291)		
				κ_{nonSEP}	1.258 (0.492)		
Rand. coeff. (σ_q)	0.85 (0.31)						


Notes: Panel A reports baseline demand parameters from Neilson (2026); the first quality entry is the common coefficient $\bar{\beta}^q$, subsequent education entries are deviations from the less-than-high-school category, and the SEP row is an additional shift. Price coefficients are signed. Panel B reports treatment-effect parameters estimated here using the sample in subsection 5.5.3. Panel C reports scale parameters and their standard errors. Standard errors in parentheses: Panel A reports the standard errors from Neilson (2026); Panels B and C are clustered at the preschool-center level. $\psi_{\text{SEP}}^{\text{op}}$ is set to zero in the benchmark because the prevalence of zero out-of-pocket prices among SEP choices leaves this shift poorly identified (see Online Appendix for the unrestricted specification).

tion is that the demand shift estimated within SEP status applies to families of the same status in the target markets. We refer the reader to Neilson (2026) for a thorough discussion of the baseline demand estimates and focus here on the new treatment parameters.

The main demand-side result is that the information treatment increases the weight families place on school quality. This effect is clearest for SEP families, where the estimated shift is economically meaningful relative to the baseline quality coefficient, but the pattern is similar for non-SEP families as well. In other words, the treatment does not simply move families mechanically toward specific schools; it makes school quality matter more in their choices.

The other treatment parameters point in the same direction. For SEP families, treatment raises the quality weight relative to the magnitude of the distance coefficient, making higher quality more attractive relative to proximity. For non-SEP families, the treatment makes out-of-pocket price matter somewhat more. Taken together, these estimates line up with the descriptive experimental evidence from Section 4: treated families place more weight on academic quality and are more willing to trade off convenience, and in some cases price, when choosing schools. Zero out-of-pocket prices at SEP schools limit the information available to estimate the SEP price shift.

The scale estimates provide only modest evidence that the intervention changed the overall sharpness of choice behavior. Both κ estimates are above one, but neither is statistically distinguishable from one at conventional two-sided levels under the displayed standard errors. The broader message is therefore not that the treatment worked mainly by reducing choice noise



Notes: Panel A: distribution of the estimated school-specific term $\hat{c}_j$ for voucher schools; dashed line marks the mean. Panels B and C: bin-averaged $\hat{c}_j$ against per-student expenditure and average teacher PSU scores for voucher schools. The term $\hat{c}_j$ may reflect persistent differences in both input costs and schools' objective functions. 22,882 school-year observations (2,389 schools), 53 markets, 2006–2015. Full parameter estimates in Table O-14.

across the board. Rather, the data suggest that the main action comes through the characteristic-specific shifts, especially the increased weight on quality. Allowing for the scaled-logit possibility emphasized by Brown and Jeon (2024) is still useful conceptually, but in this application it does not materially overturn the main demand-side interpretation.

5.7. Supply Model Estimates

Figure 2 presents the supply-side parameter estimates from the two-regime GMM procedure described in subsubsection 5.5.5, estimated on voucher and private non-voucher schools over 10 years (2006–2015) across 53 markets; Online Appendix O-15.1 gives the corresponding data map. The estimation concentrates out school-specific terms c_j and market-year-type effects $\gamma_{mt}^{\text{type}}$, with type-specific time trends $\delta_{\text{type}} \cdot t$ allowing these terms to evolve differently across school types. Identification uses the policy and local-labor-market variation from the baseline demand instrument set, interactions with for-profit and nonprofit-voucher status, and a new Nearby Teacher Quality Index (Table O-15).

The supply estimates indicate substantial heterogeneity across schools, and the two-regime specification applies the composition-channel first-order condition to the constrained observations defined in subsubsection 5.5.5. The time trends by school type, reported with the full parameter estimates in Table O-14, are comparatively modest, so the main variation of interest comes from cross-school differences in competitive incentives and the school-specific terms in the quality FOC.

The relationship between competitive incentives and observed quality also has direct historical support. Using the SEP reform, Neilson (2026) relates changes in school value added to the

model-implied pre- and post-reform quality markdowns and to changes in marginal revenue. The estimated coefficients on the two markdown terms are 0.96 and -0.98 , close to the values implied by the quality FOC, and the coefficient on the change in marginal revenue is 0.11 (s.e. 0.02). Our supply estimation uses the same economic relationship while allowing school-specific composite cost terms, sector-specific trends, capacity constraints, and the additional excluded variation described above.

Panel A of Figure 2 shows the distribution of the estimated school-specific terms for voucher schools. The distribution is concentrated at low values, with a mean of 0.35 and a tight dispersion. As discussed in subsubsection 5.4.1, these terms may reflect persistent differences in both input costs and schools’ objective functions; they should not be interpreted as deep technological parameters. Panels B and C show that higher values of $\hat{c}_j$ are associated with greater expenditure per student and teachers with stronger observed credentials.¹⁶ These descriptive relationships are consistent with an input-cost component in $\hat{c}_j$, but they do not separately identify production technology and schools’ objective functions.

6. EQUILIBRIUM EFFECTS

6.1. Counterfactual Policy Simulations

We construct the at-scale counterfactual in three layers. The first holds school characteristics fixed, the second introduces capacity constraints through centralized assignment, and the third allows schools to adjust quality. This sequence isolates the direction and magnitude of each equilibrium force.

Formally, let $\tau_{nk} \in [0, 1]$ denote the fraction of type- k families at residential location n who receive the treatment, following the spatial structure of Neilson (2026). Each family’s school choice probabilities are a mixture of the treated and control predictions from the demand model. Aggregating these mixed shares across locations yields market-level demand by type; the average quality attended by each type is the corresponding share-weighted mean over schools.

Model Fit and At-Scale Predictions Before exploring the layered counterfactuals, we verify that the structural model reproduces the experimental evidence and then examine how the implied treatment effect changes when extrapolated to the broader population. Table V reports three objects for each family group. Column 1 is the experimental RCT average treatment ef-

¹⁶Per-student expenditure comes from MINEDUC’s accounting data (*Gasto Rendido*) and includes both personnel and operational costs.

TABLE V
SUMMARY OF DEMAND-SIDE POLICY EFFECTS

	<i>Panel A: Treatment Effects on Quality</i>			<i>Panel B: Quality Gap</i>	
	Experiment	Model (demand only)		No treatment	−0.214
	RCT ATE	Exp. sample	All mkts	Full treatment	−0.160
SEP families	0.220 (0.087)	0.221 (0.081)	0.191 (0.074)	Narrowing	0.054 (25%)
Non-SEP families	0.114 (0.104)	0.151 (0.099)	0.136 (0.085)		

Notes: Panel A reports treatment effects on expected school quality. Column 1 is the experimental estimate; Column 2 is the in-sample model prediction; and Column 3 is the demand-only prediction across all 53 markets. Panel B reports the SEP–non-SEP quality gap under no treatment and full treatment. Quality is measured in standard deviations of normalized value added. Standard errors are in parentheses. Column (1) uses the preschool-clustered experimental standard error; Columns (2)–(3) use the numerical delta method for the estimated treatment parameters. Column (2) uses the unrestricted eight-parameter specification, while Column (3) uses the seven-parameter benchmark. See Online Appendix O-10 for inference details.

fect on school quality from Table III. Column 2 uses the unrestricted treatment estimates to compute the model-predicted effect over the same experiment-sample markets—an in-sample fit check confirming that the structural parameters reproduce the targeted moments when aggregated through the counterfactual machinery. Column 3 uses the restricted treatment specification, imposing $\psi_{SEP}^{op} = 0$, to compute the effect over all 53 urban markets in the model population. The comparison therefore reflects both the population change and the treatment-parameter restriction: for SEP families, the population effect (0.191 SD) is somewhat smaller than the experiment-sample prediction (0.221 SD). This is consistent with, among other things, the experiment operating in markets where low-income families face larger quality gaps and therefore have more room to gain from information. The population extrapolation is the relevant benchmark for the policy simulations that follow.

At-Scale Demand Responses We apply the treatment to eligible families (types 1–4) at full take-up ($\tau = 1$) while holding school quality and prices fixed. The demand shift produces a substantial increase in expected school quality for SEP families, reflecting the large quality weight increase estimated in Table IV.

Value added is measured at observed student composition and is held fixed as students re-sort, so the benchmark does not let peer composition feed back into measured quality; the Online Appendix reports a sensitivity in which it does. In the historical supply estimation, composi-

tion affects schools' incentives through SEP revenue; in the benchmark quality counterfactual each school's competitive-quality term is held at its observed value, so composition enters only through assignment and the quality markdown. Neither families nor schools have a separate preference over peer attributes.

Capacity-Constrained Counterfactual The demand-only benchmark assumes every family can attend their most-preferred school. In practice, popular schools face enrollment limits. To assess how congestion attenuates treatment effects, we simulate a student-proposing Deferred Acceptance (DA) algorithm in each market, where families submit rank-ordered preference lists derived from their logit choice probabilities and schools rank applicants by random lottery priority, as in Chile's centralized system.¹⁷

School capacities are imputed from administrative records. For public and voucher schools we use reported enrollment slots (*cupos totales*) from Chile's 2020 school-level enrollment data (SAE), adjusted to the estimation period using historical classroom counts (see Online Appendix O-15.2 for details). For unmatched and private schools we use alternative administrative capacity measures and market-level medians as fallbacks. The resulting enrollment-to-capacity ratio averages 0.73, indicating meaningful slack in most markets. The DA exercise should therefore be interpreted as a short-run congestion benchmark, not as evidence that household choices in the baseline estimation sample were broadly rationed.

Table VI decomposes the equilibrium effects by reporting *mean quality levels* across the population under each successive layer—the treatment effects in Panel B are the implied differences relative to the no-treatment baseline. Column (2) corresponds to the demand-only population benchmark from Table V, Column 3; the remaining columns progressively add capacity constraints and supply responses. Congestion matters: roughly three-fifths of the demand-only SEP effect survives capacity constraints (Column 3). Untreated higher-education families (types 5–6) experience a small negative displacement as treated families crowd into higher-quality schools—a spillover that the supply response in the next section can reverse.

6.2. Supply Responses

The capacity-constrained simulations show that congestion substantially attenuates the demand-side gains when the policy is scaled up. The central question is whether supply-side

¹⁷The simulation uses a common lottery number across schools and imposes no SEP or other priority classes; see Online Appendix O-10.

TABLE VI
COUNTERFACTUAL SIMULATIONS: MEAN QUALITY AND TREATMENT EFFECTS (POPULATION)

	(1)	(2)	Equilibrium (DA + Capacity)		
			(3)	(4)	(5)
	Baseline	Demand Only	No Supply Response	Supply Response	Supply Resp. (+10% MC)
Panel A: Mean Expected Quality					
All (types 1–6)	0.558	0.678	0.621	0.710	0.602
SEP (types 1, 3)	0.298	0.488	0.413	0.524	0.442
Non-SEP (types 2, 4)	0.512	0.648	0.586	0.681	0.579
Panel B: Treatment Effects (relative to Baseline)					
All (types 1–6)		0.120	0.063	0.152	0.044
		(0.045)	(0.024)	(0.041)	
SEP (types 1, 3)		0.191	0.115	0.226	0.144
		(0.074)	(0.057)	(0.061)	
Non-SEP (types 2, 4)		0.136	0.074	0.170	0.067
		(0.085)	(0.062)	(0.075)	
Panel C: Differences Across Equilibrium Concepts					
	All	SEP	Non-SEP		
Congestion attenuation, (2) – (3)	0.057	0.076	0.062		
	(0.024)	(0.030)	(0.029)		
Supply response, (4) – (3)	0.089	0.111	0.095		
	(0.037)	(0.039)	(0.040)		
Net equilibrium correction, (4) – (2)	0.032	0.036	0.033		
	(0.032)	(0.034)	(0.036)		

Notes: Panel A reports population-weighted mean expected quality; Panel B reports changes relative to baseline; Panel C reports differences across counterfactuals. All-family averages include types 1–6, while treatment applies to types 1–4. Columns (3)–(5) impose deferred acceptance with administrative capacities. Private-school quality adjusts according to the anchored quality FOC in Columns (4)–(5); public-school quality remains fixed. Column (5) raises private schools’ marginal quality costs by 10% and reports point estimates only. Standard errors in parentheses use the numerical delta method for the seven estimated treatment parameters and are conditional on the remaining model inputs. See Online Appendix Sections O-15.2 and O-10 for capacity, simulation, and inference details.

responses—schools adjusting quality in response to stronger competitive pressure—can offset this attenuation. We interpret the exercise as a medium-run approximation: schools adjust quality given the shifted demand system and observed capacity, but we abstract from entry, exit, and capacity investment.

The information treatment changes competitive incentives because it makes demand more responsive to quality. In the supply first-order condition, this compresses the quality markdown for affected schools and raises the return to improving quality. The policy therefore does not just reallocate families across existing schools; it changes the competitive environment in which schools operate.

These changes in competitive incentives motivate the equilibrium exercise. The historical quality FOC in Equation 10 can be written as

$$q_j + \text{MD}_j = A_j,$$

where A_j is the school-specific competitive-quality term. We anchor this term at the observed equilibrium:

$$\widehat{A}_j = q_j^0 + \widetilde{\text{MD}}_j(\mathbf{q}^0, 0),$$

where $\widetilde{\text{MD}}_j$ is the quality markdown after applying the same cap before and after treatment.¹⁸

Given the demand system under full treatment, private-school equilibrium quality satisfies

$$q_j^1 = \text{clip}_{[q_j^0-0.5, q_j^0+0.5]} \left\{ \widehat{A}_j - \widetilde{\text{MD}}_j(\mathbf{q}^1, 1) \right\}. \quad (14)$$

Because each school's markdown depends on the full vector $\mathbf{q}^1$, these equations jointly determine the private-school quality vector. Schools facing a larger treatment-induced markdown reduction increase quality more, subject to the half-standard-deviation bound. Public-school quality remains fixed. The resulting qualities enter the capacity-constrained assignment, but capacity does not switch private schools to a different quality equation. We interpret this anchored-FOC system as a medium-run equilibrium approximation; the two-regime FOCs continue to govern the historical supply estimation.

For the cost sensitivity, let $\lambda_c \geq 1$ denote the common multiplier on private schools' marginal quality costs. The corresponding equilibrium condition is

$$q_j^1(\lambda_c) = \text{clip}_{[q_j^0-0.5, q_j^0+0.5]} \left\{ \frac{\widehat{A}_j}{\lambda_c} - \widetilde{\text{MD}}_j(\mathbf{q}^1(\lambda_c), 1) \right\}. \quad (15)$$

The baseline equilibrium in Equation 14 corresponds to $\lambda_c = 1$.¹⁹

¹⁸Let $\eta_j(\mathbf{q}, \tau) = (1/s_j)\partial s_j/\partial q_j$ denote the own-quality semi-elasticity. We use $\widetilde{\text{MD}}_j = [\max\{\eta_j, 1/3.343\}]^{-1}$ to prevent near-zero demand derivatives from generating extreme markdowns; Online Appendix O-10 gives implementation details.

¹⁹We trace the fixed point by deterministic continuation from no treatment to full treatment and verify the equilibrium residual at every step. Online Appendix O-10 reports the continuation increments, path-stability checks, and qualification criterion.

The supply response roughly offsets the attenuation from capacity constraints, bringing the equilibrium effect back to, and slightly above, the demand-only benchmark. For SEP families in the population, the demand-only effect is 0.191 SD, falls to 0.115 SD with capacity constraints, and rises to 0.226 SD when private schools adjust quality (Table VI). The supply response relative to the capacity-constrained allocation is 0.111 SD (SE 0.039), while the final equilibrium effect is only 0.036 SD (SE 0.034) above the demand-only benchmark. Thus supply significantly offsets congestion, but the final effect is not statistically distinguishable from the demand-only effect. The same qualitative pattern holds for non-SEP families: congestion reduces the effect from 0.136 to 0.074 SD, and the supply response raises it to 0.170 SD. The 0.095-SD supply response for this group is statistically significant (SE 0.040). Although the point estimates are larger for SEP families, we do not interpret this difference as evidence of statistically significant heterogeneity. As stated in the table notes, these standard errors propagate uncertainty in the experimental treatment parameters and are conditional on the baseline demand parameters, historical supply estimates, administrative capacities, and the anchored quality FOC.

We also examine sensitivity to unmodeled input-price pressure. Column (5) of Table VI and the Online Appendix raise the composite marginal cost of quality uniformly for all private schools and re-solve the equilibrium using Equation 15. A 5% increase reduces the SEP effect from 0.226 to 0.184 SD, and a 10% increase reduces it to 0.144 SD. The all-family effect falls from 0.152 to 0.044 SD at 10%. Chile's regulated municipal-school pay schedules may anchor wages in the broader teacher market, although private voucher schools retain flexibility over supplemental pay and can recruit across sectors (Bertoni et al., 2018). We interpret these calculations as sensitivity exercises rather than estimates of an endogenous wage response.

The central lesson is that partial-equilibrium demand calculations can mislead in opposite directions depending on which equilibrium margins are omitted. A naive extrapolation that ignores congestion would overstate the short-run gains. But an analysis that accounts for congestion without allowing supply responses would understate the medium-run gains. In the supply-adjusted equilibrium, the competitive channel—schools facing stronger incentives to improve quality—roughly offsets the congestion channel, returning the effect close to the demand-only benchmark.

6.3. *Robustness and Discussion*

The equilibrium exercises rest on several modeling choices and institutional assumptions. The Online Appendix reports the higher-cost exercise summarized above and a separate accounting sensitivity for peer composition.

On the household side, the demand system is static: it represents the choice problem faced by one entering cohort and holds prior sibling placements and family composition fixed. It therefore does not separately model the sequence of school choices across children within a family. The baseline demand parameters are estimated from realized population choices and thus reflect the average influence of persistent family constraints, although sibling-school preferences are not separately identified. Three facts limit the concern that sibling constraints distort the transported demand shift: older-sibling prevalence is similar across the experimental timing groups, the choice response is not detectably attenuated among sibling families, and adding sibling status leaves the experimental moments essentially unchanged. At first appearance in the eventual primary school, 31.9% of the structural sample and 37.1% of comparable families in the model markets have an older sibling in grades 1–8; among them, 72.5% and 75.6%, respectively, enter the sibling’s school (see the Online Appendix). The counterfactuals should nevertheless be read as effects for a single entering cohort, not as a model of how the intervention changes younger siblings’ future choices.

Over a longer horizon, entry, exit, or capacity investment could in principle respond to a sustained demand shift. In practice, prior work in the Chilean context finds that these margins have not been large drivers of market-level change, given the existing market structure and apparent excess capacity (Neilson, 2026). As in Wollmann (2018), our counterfactual focuses on the adjustment of product characteristics by incumbent schools. We do not model the direction or magnitude of longer-run entry, exit, and capacity responses.

The counterfactual also leans on two institutional features of the Chilean environment. First, restrictions on school selection predate the current assignment mechanism: SEP prohibited participating schools from screening applicants, and the 2009 General Education Law extended the prohibition more broadly across subsidized elementary schools. The later centralized deferred-acceptance system (section 3) changes how oversubscribed seats are allocated, not whether schools may select. Student composition can still change when the information policy changes applications and the resulting assignments. Second, voucher-school copayments were heavily regulated during our estimation period; the 2015 Inclusion Law subsequently capped existing

copayments and began phasing out shared financing.²⁰ These institutions motivate a medium-run benchmark with quality as the school adjustment margin. We hold observed tuition fixed in all sectors; for private non-voucher schools, this is a maintained modeling assumption.

7. DESIGNING A LARGE-SCALE RCT

The preceding sections established that equilibrium forces—congestion and supply-side quality adjustment—are quantitatively important for evaluating the information policy at scale, and that an individual-level RCT cannot capture them. If treating one family shifts demand at a school, nearby control families may also be affected. Cluster-randomized trials, which assign entire geographic areas to treatment or control, can measure such effects by comparing outcomes across differentially saturated clusters (e.g., Miguel and Kremer, 2004, Muralidharan and Sundararaman, 2015). In school choice, the natural cluster would be the market, but with only 53 urban markets available, market-level randomization may provide too few independent units.

The model suggests a middle ground. Competitive pressure is highly location-specific: a school's demand depends primarily on nearby families, and its supply response depends on local markdown compression. This locality makes it possible to define geographic clusters within a market. Smaller clusters provide more randomized units but leave treated and control areas connected through shared schools. Larger clusters contain more of these local interactions but provide fewer independent units. The design problem is to balance contamination against precision.

We use Santiago to illustrate how the model can inform this trade-off before a larger-scale experiment is fielded. We examine how cluster size, the provider core, and buffer zones affect the contrast in private-school quality responses recovered by the experiment.

Design. We tile Santiago with a regular hexagonal grid of circumradius r kilometers (Figure 3). Hexagons approximate circles while tiling the plane without gaps. Eligible hexagons are assigned to treatment or control through stratified randomization, balanced within strata defined by distance from the market centroid and angular quadrant.

The design separates the area in which families may receive treatment from the provider sample used for estimation. Schools enter the regression if they lie in a core of radius I around

²⁰Even before the 2015 law, the shared-financing rules (DFL No. 2 of 1998) capped copayments and reduced the general subsidy as fees rose; SEP required participating schools to charge eligible students zero. See Sanchez (2017) for related discussion.

the center of an eligible hexagon. The treatment region extends to radial distance $r - B$, where B is a buffer. Increasing B contracts the treatment region toward the center while leaving the provider core fixed. A hexagon is eligible only if its core contains residential demand and at least one voucher or private school included in the provider-response model.

Let Z_h indicate assignment of eligible hexagon h . Within the treatment region of an assigned hexagon, one quarter of intervention-eligible families receive the intervention ($\tau = 0.25$).²¹ The calculation represents this saturation as fractional family mass rather than drawing individual treatment realizations. Families in control hexagons or outside the radial treatment cutoff remain untreated.

Endpoint-exposure approximation. Let q_j^0 and q_j^1 denote school j 's quality in the solved no-treatment and full-treatment equilibria. For eligible family cell i , let p_{ij}^0 denote the baseline probability of choosing school j , m_i its population weight, and $s_i(Z)$ its treatment saturation under assignment Z . School j 's exposure to treated eligible demand is

$$E_j(Z) = \frac{\sum_i m_i p_{ij}^0 s_i(Z)}{\sum_i m_i p_{ij}^0}. \quad (16)$$

We map this exposure into provider quality according to

$$q_j^g(Z) = q_j^0 + g(E_j(Z)) (q_j^1 - q_j^0). \quad (17)$$

The benchmark uses the linear mapping $g(E) = E$. We also calculate a concave mapping, $g(E) = \sqrt{E}$, and a convex mapping, $g(E) = E^2$. All three are anchored to the stored no-treatment and full-treatment endpoints by construction.

This is the simplifying assumption that makes the complete prospective-design exercise feasible. Exposure is calculated from baseline demand, the resulting quality vector is not fed back into demand, and the provider problem is not re-solved under each randomized assignment. The approximation is therefore transparent about equilibrium selection: it does not select among multiple intermediate roots. Instead, it traces a deterministic path between two solved equilibria and tests whether the design lessons survive alternative responses along that path. In

²¹ Intervention-eligible families correspond to model types 1–4; the remaining two family types are outside the intervention target.

particular, a provider’s intermediate response depends on its own baseline-demand exposure; competitor responses do not feed back through the partial-treatment quality vector. Because within-cluster saturation is 0.25, exposure in the design grid never approaches one. The square-root and square mappings should therefore be interpreted primarily as sensitivity to the level of the partial-saturation response, not as direct evidence about its curvature near full treatment.

Estimand and estimator. The design target is the average quality response of responding core schools when every eligible cluster is assigned at saturation $\tau = 0.25$, relative to assignment of no clusters. It is distinct from q^1 , which corresponds to full exposure and is used only to anchor the interpolation. For a randomized assignment, the modeled quality response is $\Delta Q_j^g(Z) = q_j^g(Z) - q_j^0$. We estimate


$$\Delta Q_j^g = \alpha + \beta Z_{h(j)} + \gamma Q_{j,\text{baseline}} + \epsilon_j, \quad (18)$$

giving equal weight to responding voucher and private schools in eligible cores and clustering standard errors at the hexagon level. Public schools are excluded because their quality is held fixed in the provider counterfactual.

Bias–precision trade-off. Interference prevents assignment from generating a sharp treated-versus-control exposure contrast. A provider in a treated core may draw demand from control areas, while a provider in a control core may draw demand from treated areas. Larger clusters can contain more of these local interactions but leave fewer randomized units. A buffer can improve separation between treated and control providers but removes treated families near cluster boundaries and therefore changes the target. Narrowing the provider core can exclude especially exposed boundary observations but reduces the estimation sample. The relevant comparison must therefore consider bias and precision together rather than treating any single geometric margin as uniformly beneficial.

Simulation. We evaluate 96 candidate combinations of hexagon radius, provider-core radius, and buffer. For each geometry, 40 independent assignments determine the centered reference and 200 additional assignments evaluate the design. Each assignment uses a random grid origin and stratified treatment assignment. Figure 3 presents the benchmark linear mapping for the no-buffer designs. We use the remaining buffers and the concave and convex mappings as sensitivity checks.

FIGURE 3.—Scale-Up RCT Design



Notes: Panel (a) illustrates the Santiago geometry: teal hexagons are assigned to treatment, gray hexagons to control, and darker inner areas are provider cores. Panel (b) reports RMSE relative to the all-assigned endpoint-interpolation target and the share of that target recovered by the assignment contrast under the benchmark linear mapping with $B = 0$. The burgundy outline marks the minimum-RMSE displayed design recovering at least 90 percent of its design-specific target ($r = 4.0, I = 2.0$ km). Gray cells are geometrically infeasible. The calculation holds fitted inputs fixed and does not solve a partial-treatment equilibrium for each assignment.

Results. Cluster geometry materially changes attenuation and RMSE. Among the designs with defined regressions, larger-radius configurations with moderate provider cores tend to produce smaller RMSE under the linear benchmark than small-radius configurations with broad cores. Buffers often reduce absolute bias and RMSE within a fixed radius and core, but they also reduce the all-assigned target by leaving additional boundary families untreated. The target also changes with radius and core because these dimensions alter cluster eligibility and the provider sample. Absolute RMSE should therefore not be used alone to rank designs anywhere in the grid. The lower panel makes the corresponding attenuation explicit by reporting the estimated contrast as a share of its own design-specific target.

The highlighted benchmark design, $r = 4.0$ and $I = 2.0$ km, has RMSE 0.0046 and recovers 91 percent of its target under the linear no-buffer mapping. This is a benchmark-conditional recommendation, not an unconditional structural optimum. Under the same 90-percent-recovery criterion, the concave and convex exposure mappings also select $r = 4.0$ km but select different

core radii. The preference for a large hex radius is therefore more stable than the exact core choice.

Thirteen of the 96 geometries contain at least one random assignment for which the provider regression is undefined, generally because a very small core inside a large hexagon leaves too few usable treated or control providers. These failures are part of the design outcome: the corresponding cells remain marked in the sensitivity figures rather than dropping assignments or replacing them.

Conditional rejection of zero, retained as a separate diagnostic, is high for most feasible cells, but it is not field power. The calculation observes modeled provider quality without measurement error and conditions on the fitted endpoints; several sparse designs also show nontrivial centered-null rejection.

The model nevertheless provides concrete guidance. It identifies the margins governing the bias–precision trade-off, shows which geometries are infeasible, and describes how much of the all-assigned quality response a clustered experiment would recover under transparent alternative exposure mappings. A future implementation can choose among operationally feasible designs using preregistered contamination and precision criteria, then use pilot data to calibrate the observation process and assess the endpoint-exposure approximation.

8. DISCUSSION

This paper studies how a low-cost information policy affects school choice and what its effects would be at scale. In the experiment, treated families—especially those reached before they had enrolled—choose schools with higher test scores and value added, and their children score higher in fourth grade. Estimates on later assessments remain positive, although they are observed for smaller samples selected by test participation. Together, these results show that a simple intervention delivered through existing preschool infrastructure can shift actual choices toward higher-quality schools and improve academic outcomes.

The policy question, however, is not only whether the intervention works in a small-scale trial, but what happens once it is implemented broadly enough to change market conditions. Our counterfactuals show that this extrapolation is not mechanical. Congestion attenuates the demand-side gains when more families seek access to the same high-quality schools. At the same time, the increase in demand for quality changes schools' incentives, and in our preferred specification the resulting supply-side quality response roughly offsets the attenuation from

congestion. The quantitative effects remain conditional on the assignment environment, the estimated cost structure, and our scope restriction that public-school quality is held fixed.

The methodological contribution comes from the complementary roles of the experiment and the model. We do not rely on the structural model alone to infer the demand shift induced by the policy; the randomized trial provides direct experimental evidence on that shift. But an experiment that treats a small share of families in connected urban markets cannot identify the equilibrium effects that arise only when enough families are treated to change congestion, assignment outcomes, and schools' quality incentives. The model provides that equilibrium extension and makes the maintained assumptions behind it explicit. In that sense, the framework turns a small-scale experiment into a tool for prospective policy analysis rather than treating it as a stand-alone estimate of policy impact.

The same framework is useful for experimental design. When a policy is rolled out gradually or across only part of a set of urban markets, the design of the larger-scale evaluation determines what can be learned about equilibrium effects. Our scale-up exercise shows how the estimated model can be used to choose clusters, provider cores, and buffer zones in a way that balances precision against contamination from spillovers, while making explicit the approximation used between solved equilibrium endpoints. More broadly, the paper illustrates how a credible small-scale experiment, combined with enough economic structure to analyze market adjustment, can provide disciplined guidance about both the likely equilibrium effects of a new policy and the design of the larger-scale experiments needed to learn more about it.

REFERENCES

- ANDRABI, TAHIR, JISHNU DAS, AND ASIM IJAZ KHAWAJA (2017): "Report cards: The impact of providing school and child test scores on educational markets," *American Economic Review*, 107 (6), 1535–63. [1, 3]
- BERRY, STEVEN, JAMES LEVINSOHN, AND ARIEL PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890. [24]
- (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Vehicle Market," *Journal of Political Economy*, 112 (1), 68–105. [55]
- BERTONI, ELEONORA, GREGORY ELACQUA, ANALÍA JAIMOVICH, JULIO RODRÍGUEZ, AND HUMBERTO SANTOS (2018): "Teacher Policies, Incentives, and Labor Markets in Chile, Colombia, and Peru: Implications for Equity," Tech. rep., Inter-American Development Bank. [37]
- BOBBA, MATTEO, VERÓNICA FRISANCHO, AND MARCO PARIGUANA (2026): "Perceived Ability and School Choices: Experimental Evidence and Scale-up Effects," *Journal of Political Economy Microeconomics*, forthcoming. [4]

- BORGER, MICHAEL, GREGORY ELACQUA, ISABEL JACAS, CHRISTOPHER NEILSON, AND ANNE SOFIE WESTH OLSEN (2024): “Report cards: Parental preferences, information and school choice in Haiti,” *Economics of Education Review*, 102, 102560. [1]
- BROWN, ZACH Y. AND JIHYE JEON (2024): “Endogenous Information and Simplifying Insurance Choice,” *Econometrica*, 92 (3), 881–911. [2, 16, 18, 19, 31, 49]
- CAMPOS, CHRISTOPHER (2024): “Social Interactions, Information, and Preferences for Schools: Experimental Evidence from Los Angeles,” NBER Working Paper No. 33010. [1]
- CHUMACERO, ROMULO A., DANIEL GOMEZ, AND RICARDO D. PAREDES (2011): “I would walk 500 miles (if it paid): Vouchers and School Choice in Chile,” *Economics of Education Review*, 30 (5), 1103 – 1114. [2]
- CORRADINI, VIOLA (2024): “Information and Access in School Choice Systems: Evidence from New York City,” Working Paper. [1]
- DINERSTEIN, MICHAEL AND TROY SMITH (2015): “Quantifying the Supply Response of Private Schools to Public Policies,” Tech. rep., Stanford Institute for Economic Policy Research. [2, 4]
- GALLEGO, FRANCISCO (2013): “When Does Inter-School Competition Matter? Evidence from the Chilean “Voucher” System,” *The BE Journal of Economic Analysis & Policy*, 13 (2), 525–562. [23]
- GALLEGO, FRANCISCO AND ANDRES E. HERNANDO (2008): “On the Determinants and Implications of School Choice: Semi-Structural Simulations for Chile,” *Journal of LACEA Economia*. [2]
- HANSEN, LARS PETER (1982): “Large sample properties of generalized method of moments estimators,” *Econometrica: Journal of the Econometric Society*, 1029–1054. [55]
- HASTINGS, JUSTINE S. AND JEFFREY M. WEINSTEIN (2008): “Information, School Choice, and Academic Achievement: Evidence from Two Experiments,” *Quarterly Journal of Economics*, 123 (4), 1373–1414. [1]
- HECKMAN, JAMES J, LANCE LOCHNER, AND CHRISTOPHER TABER (1998): “General-Equilibrium Treatment Effects: A Study of Tuition Policy,” *American Economic Review*, 88 (2), 381–386. [1]
- HOXBY, CAROLINE M. (2003): “School Choice and School Productivity: Could School Choice Be a Tide That Lifts All Boats?” in *The Economics of School Choice*, ed. by Caroline M. Hoxby, University of Chicago Press, 287–342. [22]
- KAPOR, ADAM, CHRISTOPHER A. NEILSON, AND SETH D. ZIMMERMAN (2020): “Heterogeneous Beliefs and School Choice Mechanisms,” *American Economic Review*, 110 (5), 1274–1315. [1, 2]
- LISE, JEREMY, SHANNON SEITZ, AND JEFFREY SMITH (2015): “Evaluating search and matching models using experimental data,” *IZA Journal of Labor Economics*, 4 (1), 16. [4]
- MIGUEL, EDWARD AND MICHAEL KREMER (2004): “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities,” *Econometrica*, 72 (1), 159–217. [39]
- MURALIDHARAN, KARTHIK AND VENKATESH SUNDARARAMAN (2015): “The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India,” *The Quarterly Journal of Economics*, 130 (3), 1011–1066. [39]
- NEILSON, CHRISTOPHER A. (2026): “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students,” *Econometrica*, conditionally Accepted. [2, 4, 6, 8, 10, 11, 13, 16, 17, 19, 20, 21, 23, 24, 26, 29, 30, 31, 32, 38, 49, 50, 51, 54]
- OECD (2016): *Starting strong IV: Monitoring quality in early childhood education and care - Chile*, OECD Publishing. [8]

- PETRIN, AMIL (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110 (4), pp. 705–729. [25]
- SANCHEZ, CRISTIAN (2017): “Understanding School Competition Under Voucher Regimes,” Job market paper, phd dissertation, ITAM. [21, 39]
- TODD, PETRA E. AND KENNETH I. WOLPIN (2006): “Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility,” *American Economic Review*, 96 (5), 1384–1417. [4]
- (2023): “The Best of Both Worlds: Combining Randomized Controlled Trials with Structural Modeling,” *Journal of Economic Literature*, 61 (1), 41–85. [4]
- TRAIN, KENNETH E. (2009): *Discrete Choice Methods with Simulation*, Cambridge University Press, 2nd ed. [19]
- WALTERS, CHRISTOPHER R. (2018): “The Demand for Effective Charter Schools,” *Journal of Political Economy*, 126 (6), 2179–2223. [2, 4]
- WOLLMANN, THOMAS G. (2018): “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles,” *American Economic Review*, 108 (6), 1364–1406. [38]

APPENDIX

A1. *Reduced-Form Regression Specifications*

Tables 1 and 2 report intent-to-treat coefficients from reduced-form regressions of the form

$$Y_i = \alpha + \beta T_p + Z_p' \delta + X_i' \gamma + \mu_c + \varepsilon_i,$$

where T_p is treatment assignment at the preschool-center level, μ_c are preschool-municipality fixed effects, and standard errors are clustered at the preschool-center level.

In the main specification, the market controls are nine market-level variables: the average SIMCE score used in the report card, the number of schools in the market, the 10th, 25th, 50th, 75th, and 90th percentiles of the SIMCE distribution, the interquartile range of SIMCE scores, and the standard deviation of SIMCE scores. The expanded family controls are child sex, SEP priority eligibility, indicators for the mother having completed high school or higher education, maternal marital status, household size, an index of durable-goods ownership, and home ownership. Missing-value indicators are included for every family covariate and missing values are set to zero, so the controls do not alter the estimation samples.

The Online Appendix reports estimates with no family controls, maternal education only, alternative birth-record controls, and richer combinations of demographic and survey controls. These controls are introduced as complete blocks rather than selected separately by outcome. It also reports balance and coverage for the controls used in the main specification. The household variables are observed for every child in Panels B and C, apart from one missing home-ownership value in Panel C; missing-indicator coding preserves the larger pooled sample.

Panel A uses the full experimental sample. Panel B restricts to families who had already enrolled at the time of the Pre-K visit, and Panel C restricts to families who had not yet enrolled. For Table 1, column 1 restricts the distance outcome to observations with distance below 4 km. Columns 2–6 use the subsample with non-missing school value added so that the school-choice outcomes are compared on a common sample. For Table 2, columns 7–8 are conditional on taking the college entrance exam, and the reported scores pool the 2023–2025 cohorts standardized within cohort.

A2. *Experiment Moments for Structural Estimation*

Table 3 reports the experimental moments used in the second-step demand estimation. These regressions use the model-overlap sample described in Section 5.5.3: families who had not yet enrolled at the time of treatment, can be matched to the six markets covered by the baseline demand model, pass the geographic consistency filter, and have non-missing school characteristics used in estimation.

The structural moment regressions include preschool-municipality fixed effects, the nine market-level SIMCE controls listed in Appendix A1, indicators for the mother having completed high school or higher education, and a missing-education flag. Standard errors are clustered at the preschool-center level. This narrow maternal-education block differs from the expanded family controls in the reduced-form analysis. The GMM vector stacks effects on five chosen-school outcomes—quality (normalized value added Q_2), sticker price, out-of-pocket price, Google Maps driving distance restricted to observations within 4 km, and voucher-school attendance—separately for SEP and non-SEP families. These ten entries are followed by two residualized mean-utility entries, one per group, giving twelve entries in total. Pooled estimates are descriptive and do not add entries to the GMM objective.

A3. *Scaled-Logit Derivation*

This subsection gives a simplified type-level illustration of the scaled-logit interpretation summarized in Section 5.3, suppressing within-type random quality tastes for exposition. The purpose is interpretive: it shows why imperfect information about school characteristics can motivate both characteristic-specific treatment effects and a common scale shift in the choice model. The demand estimation and counterfactual exercises in the main text are implemented directly in terms of the effective parameters, so none of the policy results requires separately identifying the underlying signal process.

Suppose that for each characteristic $c \in \{q, \text{op}, d\}$, family i of type k observes a noisy signal

$$\tilde{s}_{ij}^c = c_j + v_{ij}^c,$$

where $v_{ij}^c \sim N(0, \sigma_{v,c,k}^2)$ is independent noise. Given a prior $c_j \sim N(0, \sigma_c^2)$, the posterior mean is

$$E[c_j \mid \tilde{s}_{ij}^c] = \rho_k^c \cdot \tilde{s}_{ij}^c, \quad \rho_k^c = \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{v,c,k}^2} \in (0, 1), \quad (19)$$

where ρ_k^c is the signal-to-noise ratio for characteristic c and type k .

Under this simplification, substituting the posterior means into utility yields expected utility of the form

$$U_{ij}^{\mathcal{E}} = \phi_k^q q_j + \phi_k^{op} \text{op}_{k(i),j} + \phi_k^d d_{\text{loc}(i),j} + \bar{\beta} x_j + \xi_j + \tilde{\epsilon}_{ij}, \quad (20)$$

where $\phi_k^q = \beta_k \rho_k^q$, $\phi_k^{op} = \alpha_k \rho_k^{op}$, and $\phi_k^d = \lambda_k \rho_k^d$ are products of structural preferences and signal precision. Thus, noisier signals attenuate the effective weight placed on the corresponding characteristic. The composite error is

$$\tilde{\epsilon}_{ij} = \phi_k^q v_{ij}^q + \phi_k^{op} v_{ij}^{op} + \phi_k^d v_{ij}^d + \epsilon_{ij}. \quad (21)$$

The variance of this composite error is

$$\text{Var}(\tilde{\epsilon}_{ij} \mid k) = (\phi_k^q)^2 \sigma_{v,q,k}^2 + (\phi_k^{op})^2 \sigma_{v,op,k}^2 + (\phi_k^d)^2 \sigma_{v,d,k}^2 + \frac{\pi^2}{6}. \quad (22)$$

As in Brown and Jeon (2024), improving signal precision can therefore affect not only the effective coefficients on informed characteristics but also the overall error variance of the choice model.

We approximate this environment with a scaled logit:

$$\tilde{\epsilon}_{ij} \approx \mu_k \cdot \epsilon_{ij}^*,$$

where ϵ_{ij}^* is i.i.d. EV1 and the type-specific scale is chosen to match the variance of the composite error:

$$\mu_k = \sqrt{1 + \frac{6 \sigma_{\text{noise},k}^2}{\pi^2}} \geq 1, \quad (23)$$

with

$$\sigma_{\text{noise},k}^2 \equiv (\phi_k^q)^2 \sigma_{v,q,k}^2 + (\phi_k^{op})^2 \sigma_{v,op,k}^2 + (\phi_k^d)^2 \sigma_{v,d,k}^2.$$

Dividing expected utility by μ_k then yields standard logit choice probabilities with effective parameters equal to structural parameters divided by the error scale.

The baseline demand estimation in Neilson (2026) identifies these effective parameters under the normalization $\mu_k^C = 1$. If treatment changes signal precision, the treated and control scales

need not coincide. The common scale parameter in Section 5.2 is therefore

$$\kappa_k \equiv \frac{\mu_k^C}{\mu_k^T}. \quad (24)$$

When improved signal precision reduces composite error variance, $\mu_k^T < \mu_k^C$ and hence $\kappa_k > 1$: choice behavior becomes sharper even for school attributes about which the intervention provides no direct information. When $\kappa_k = 1$, the model collapses to the standard specification with only characteristic-specific treatment effects. In the main text we treat this signal model as one interpretation that motivates allowing a common scale shift; the paper's identification and counterfactual analysis proceed directly in terms of the effective parameters and do not require separately estimating the latent signal process.

A4. Estimation Specifics

We take the published baseline demand estimates (θ_1, θ_2) and recovered mean utilities δ from Neilson (2026) as fixed inputs rather than re-estimating them in this paper. For completeness, we summarize the MPEC problem used to obtain those inputs. This formulation treats δ as an optimization variable and exploits the sparsity of the Jacobian of the market-share equations. Let $m_{\text{micro}}(\delta, \theta_2)$ denote the vector of micro moments from school choice and let $m_{\text{IV}}(\delta, \theta_1, \theta_2)$ denote the vector of IV orthogonality conditions for the demand residual. Neilson's first-step problem can be written as

$$(\theta_1^*, \theta_2^*, \delta^*) = \arg \min_{\theta_1, \theta_2, \delta} \begin{bmatrix} m_{\text{micro}}(\delta, \theta_2) - \bar{m}_{\text{micro}} \\ m_{\text{IV}}(\delta, \theta_1, \theta_2) \end{bmatrix}' \begin{bmatrix} W_{MM} & 0 \\ 0 & W_{IV-D} \end{bmatrix} \begin{bmatrix} m_{\text{micro}}(\delta, \theta_2) - \bar{m}_{\text{micro}} \\ m_{\text{IV}}(\delta, \theta_1, \theta_2) \end{bmatrix}. \quad (25)$$

The constraints are:

$$\begin{aligned} \delta &= s^{-1}(\bar{S}, \theta_2) && \text{(BLP inversion matching observed shares)} \\ \xi_{jt}(\delta, \theta_1, \theta_2) &= \delta_{jt} - f_{jt}(\theta_1) && \text{(demand residual)} \\ \xi^{\text{norm}} &= 0 && \text{(normalization restrictions),} \end{aligned}$$

where $f_{jt}(\theta_1) = \sum_r \eta_k^r x_{jt}^r$ is the linear index in observed school characteristics. The normalization restriction means that, within each market-year, mean utilities are defined relative to

one designated normalization school, so δ_{jt} and ξ_{jt} are identified only up to a market-year-specific additive constant. Thus, $m_{\text{micro}}(\delta, \theta_2)$ corresponds to model-implied averages of the chosen school characteristics by family type, $\bar{m}_{\text{micro}}$ is the corresponding vector in the data, and $m_{\text{IV}}(\delta, \theta_1, \theta_2)$ stacks the sample analog of $\mathbb{E}[Z_{jt}\xi_{jt}] = 0$ for the excluded demand instruments Z_{jt} .

A4.1. Demand Moments and Excluded Instruments

We adopt the baseline demand system from Neilson (2026). For convenience, we summarize the moments and excluded instruments used in Neilson’s estimation, since they are central to the interpretation of our treatment and equilibrium exercises.

Moments. Neilson’s demand estimation combines three types of moments. First, *market share moments* match observed school-level shares and recover mean utilities through the BLP inversion. Second, *micro moments* match the average quality, out-of-pocket payment, and distance chosen by each family type in the micro-data. Third, *IV moments* impose orthogonality between the demand residual and excluded instruments. These moments play different roles. The market-share moments recover the school-level unobserved demand component, the micro moments discipline heterogeneity in preferences across family types, and the IV moments address the endogeneity of price and academic quality. Thus the micro moments complement, rather than replace, the market-share inversion and IV conditions. In Neilson (2026), school-level choice data are available for the full population in the modeled markets, while family-specific distance information is available for a smaller but still very large geocoded subsample; accordingly, the distance micro moments are constructed from that subsample, whereas the market-share and IV moments use the broader administrative panel. For our experimental second step, the distance moment is likewise built from family-specific home-to-chosen-school driving distance, while the counterfactual simulations use node-to-school driving distances for the full choice set once each family is mapped to a residential location (node) within a market.

Why instruments are needed. Price and academic quality are endogenous school-side choice variables: schools may choose them in response to unobserved demand components. Following Neilson (2026), let the estimated demand residual be $\tilde{\xi}_{jt} = \xi_j + \Delta\xi_{jt}$, where ξ_j is a fixed non-academic school attribute and $\Delta\xi_{jt}$ is a period-specific shock. The key maintained conditions are that ξ_j is orthogonal to the instruments and that $\Delta\xi_{jt}$ is independent of both the instruments and the school’s price and quality choices. Under these assumptions, excluded instruments shift

school incentives without directly moving household demand through unobserved school appeal.

Excluded demand instruments. The main excluded instruments are:

- *Policy-exposure instrument:* the timing of the SEP reform interacted with a school's exposure to SEP-eligible students in its surrounding neighborhood.
- *Simulated transfer instrument:* the change in government transfers implied by the SEP reform when school composition is held fixed at its pre-policy level.
- *Local labor-cost shifter:* a location-by-year proxy for the cost of hiring skilled workers, constructed from wage residuals outside the school sector.

The first two instruments primarily shift schools' revenue and pricing incentives. The local labor-cost shifter is the most direct excluded source of variation for academic quality, because quality provision in this setting is teacher-intensive and labor spending accounts for the large majority of school costs. Together, these instruments provide excluded variation for the endogenous price and quality choices in the baseline demand system.

In the second step, we estimate the treatment-effect parameter vector θ_2^T by matching treatment effects in the experimental data to treatment effects simulated by the model. Let

$$m_{\text{RCT}}(\theta_2^T)$$

denote the vector of second-step moment residuals. In the specification with separate scale parameters for SEP and non-SEP families, this vector stacks the twelve moments described in Section 5:

- treatment effects on chosen quality, sticker price, out-of-pocket payment, and distance, each by SEP status;
- treatment effects on voucher-school attendance, by SEP status; and
- treatment effects on residualized mean utility $\tilde{\delta}_j \equiv \delta_j - \bar{\beta}^q q_j$, by SEP status.

Each element of $m_{\text{RCT}}(\theta_2^T)$ is therefore a difference between an experimental treatment coefficient estimated in the data, $\hat{\beta}^{\text{RCT}}$, and the analogous coefficient computed from simulated choices, $\hat{\beta}^{\text{sim}}(\theta_2^T)$. The second-step problem is

$$\theta_2^{T*} = \arg \min_{\theta_2^T} m_{\text{RCT}}(\theta_2^T)' W_{\text{RCT}} m_{\text{RCT}}(\theta_2^T). \quad (26)$$

In the common- κ specification, the $\tilde{\delta}$ moments are omitted and $m_{\text{RCT}}(\theta_2^T)$ contains ten moments instead of twelve.

In the third step, we estimate supply-side parameters (θ_3) from the quality first-order conditions. For an unconstrained school-year, the implemented residual is

$$\Omega_{jt}^U(\theta_3) = \text{CompQ}_{jt} - \frac{R_{jt} - \gamma_{\text{SEP}} \text{SEP}_{jt} - \gamma_{mt}^{\text{type}(j)}}{c_j + \delta_{\text{type}(j)}(t - 2006)}, \quad \text{CompQ}_{jt} \equiv q_{jt} + s_{jt} \left[\frac{\partial s_{jt}}{\partial q_{jt}} \right]^{-1}. \quad (27)$$

For a capacity-constrained school-year, the residual instead imposes the composition-channel condition:

$$\Omega_{jt}^C(\theta_3) = \Delta R_{jt} \frac{\partial s_{jt}^{\text{SEP}} / \partial q_{jt}}{s_{jt}} - [c_j + \delta_{\text{type}(j)}(t - 2006)].$$

The estimation uses Ω_{jt}^U or Ω_{jt}^C according to the capacity classification described in Section 5.5.5. It concentrates out the school-specific quality cost c_j and the market-year-type effect $\gamma_{mt}^{\text{type}(j)}$ jointly by alternating projections. The freely estimated parameters are γ_{SEP} , δ_{voucher} , and δ_{private} .

Let Z_{jt}^S denote the supply instrument matrix. The supply moments are the sample analog of

$$\mathbb{E}[Z_{jt}^S \Omega_{jt}(\theta_3)] = 0,$$

and we stack them in a vector $m_S(\theta_3)$. The matrix contains the baseline demand instruments and predetermined school characteristics, interactions of that matrix with for-profit and nonprofit-voucher indicators, and a new cost-side instrument: the inverse-distance-weighted average of public-school teacher PSU scores in the same market-year. The excluded baseline instruments are simulated SEP transfer changes, reform timing interacted with pre-policy neighborhood SEP exposure, and local non-education earnings deviations. SEP participation itself is the structural cost shifter whose coefficient is estimated; it is not the excluded instrument. The third-step GMM problem is

$$\theta_3^* = \arg \min_{\theta_3} m_S(\theta_3)' W_{IV-S} m_S(\theta_3). \quad (28)$$

A4.2. Computational Details

Numerical integration. The demand model includes a single random coefficient on quality ($\nu_i \sim N(0, \sigma^2)$). All integration over this random coefficient—in the market share formula (Equation 2), in the simulated treatment effects (Equation 12), and in the supply-side markdowns—uses 5-point Gauss-Hermite quadrature, a deterministic rule that is exact for polynomials up to degree 9. No Monte Carlo simulation draws are used in the estimation.

Weighting matrices. For Step 1, we use the published demand weighting matrix from Neilson (2026):

$$W = \text{diag}(W_{MM}, W_{IV-D}).$$

The second-step (treatment effects) and third-step (supply) estimations each use two-step efficient GMM: the first GMM step uses the identity matrix $W^{(1)} = (1/n_m) \cdot I$, and the second uses the inverse of the estimated covariance matrix of the moments: $W_{\text{RCT}}^{(2)} = \hat{V}_{\text{CR}}^{-1}$ (clustering at the preschool center level) for Step 2, and $W_S^{(2)} = \hat{V}_S^{-1}$ for Step 3, where $\hat{V}_S$ is the heteroskedasticity-robust covariance of the school-year supply moments, computed without clustering; supply standard errors use the corresponding GMM sandwich formula.

Optimization algorithms and initialization. In Step 1, our code loads the published baseline demand estimates from Neilson (2026) rather than re-estimating them. Step 2 uses a constrained optimization algorithm with analytical gradients, a nonnegativity bound on the SEP quality-shift parameter, bounds $\kappa \in [0.5, 4.0]$, and a multi-start strategy. In the preferred two- κ specification reported in the paper, we use 7 starting points in each of the two GMM steps; all 7 converge to the same objective value and the same parameter estimates in the identity-weighted step, and all 7 converge to the same objective value and the same parameter estimates in the efficient-GMM step. The optimization tolerances are 10^{-8} (optimality) and 10^{-10} (step size). Step 3 concentrates out the school-specific quality costs c_j and market-year-type effects $\gamma_{mt}^{\text{type}(j)}$ via an alternating-projection algorithm. The three free parameters are γ_{SEP} , δ_{voucher} , and δ_{private} . The concentrated objective is minimized using a Nelder-Mead simplex algorithm with tolerances of 10^{-12} (function) and 10^{-10} (parameter).

A5. *Calculating Standard Errors*

The standard errors of the estimated parameters in each step of the estimation procedure are obtained from the variance-covariance matrix for the GMM estimator proposed by Hansen (1982). We will discuss how we calculate the standard error for a generic case (the parameters θ and moments $\mathbf{M}$), and then we discuss the case of each set of parameters more specifically. Each one of our GMM estimators is the result of an optimization problem in which the objective function has a quadratic form:

$$\min_{\theta} Q_{obj} = \mathbf{M}' W_m \mathbf{M}$$

For which the gradient is:

$$\frac{\partial Q_{obj}}{\partial \theta} = 2 \cdot \mathbf{J}'_M W_m \mathbf{M}$$

Where $\mathbf{J}_M$ is the Jacobian.

The variance-covariance matrix for a GMM estimator is calculated using the estimator proposed by Hansen (1982):

$$\text{cov}(\theta) = (\mathbf{J}'_M W_m \mathbf{J}_M)^{-1} \mathbf{J}'_M W_m \mathbf{V} W'_m \mathbf{J}_M (\mathbf{J}'_M W_m \mathbf{J}_M)^{-1}$$

Where V is the variance-covariance matrix of the moments. With deterministic Gauss-Hermite quadrature, there is no simulation error in the model predictions. For the parameters estimated from the experimental moments, the relevant source of uncertainty is therefore sampling error in the treatment-effect estimates. As discussed by Berry et al. (2004), when both sources are present they enter additively; in our current implementation only the sampling component remains. The RCT moments in Equation 12 take the difference between the estimated treatment effects and the model's predictions, so the variance of those moments is estimated from the cluster-robust variance-covariance matrix of the underlying treatment-effect regressions.