

Equilibrium Consequences of Vouchers Under Simultaneous Extensive and Intensive Margins Competition

Cristián Sánchez*

April 13, 2026

Abstract

I study supply-side responses to a targeted voucher program with voluntary school participation, combining administrative data from Chile with an equilibrium model of school demand and competition. The policy induces market segmentation: lower-quality schools opt in and reduce fees, while higher-quality schools opt out and raise them. Low-income students face lower prices and modest quality gains; higher-income students shift toward non-participants. Budget-neutral redesigns improve the quality mix of participating schools but yield limited gains for low-income students, given their low willingness to pay for school attributes. Only mandated participation delivers appreciable gains, at higher fiscal cost.

Keywords: Education, vouchers, market segmentation, discrete-continuous games.

*Financial Research Unit, Central Bank of Chile. Agustinas 1180, Santiago, Chile. E-mail: csanchez@bcentral.cl. I would like to thank Andrew Sweeting and Sergio Urzúa for their invaluable guidance and support during the early stages of this project. I am grateful for comments from Jean Francois Fournel, Judy Hellerstein, Yuhta Ishii, Guido Kuersteiner, Horacio Larreguy, Andrés Luengo, Shaun McRae, Chris Neilson, Yongjoon Paek, Fernando Saltiel, Pedro Sant'Anna, and Lesley Turner, as well as seminar and conference participants at various institutions. Gabriel Cañedo Riedel provided excellent research assistance. I thank the Ministry of Education and the Agencia de Calidad de la Educación in Chile for providing the data for this paper. Such data was obtained during my time at the University of Maryland, and was not accessed via any of the confidential data agreements the Central Bank of Chile has signed with public institutions and regulating agencies. Most of this paper's work was carried out during my time as an assistant professor at ITAM. I thank ITAM for giving me access to its servers after I left the institution. The views expressed are those of the author and do not necessarily represent the views of the Central Bank of Chile or its board members. All errors are my own.

1 Introduction

Much of the impact of policies with voluntary participation from suppliers ultimately depends on these suppliers' decision to join the program. The number of participants determines policy coverage, while the identity of participants shapes the quality of the product or service offered. Moreover, extensive margin responses also affect intensive margin equilibria—such as prices and quality. As a result, policies with voluntary supplier participation not only shift demand; they also reshape market structure. Ultimately, the interplay between participation costs and policy design determines whether regulation improves welfare or backfires. Ignoring such supply-side responses risks misestimating effects and misguiding policy design.

In this paper, I study primary schools' participation responses in Chile's targeted voucher program, and their consequences on market equilibrium outcomes—such as prices and the quality composition of participating schools—as well as on students' outcomes and welfare, including fees paid and the quality of schools they access. I do so by combining rich administrative data with an empirical model of school demand and competition, in which I focus on schools' decision to join the program and on the fees they charge in equilibrium.

Chile introduced the targeted voucher in 2008, in addition to the long-standing universal voucher, to inject additional resources into schools that enroll socioeconomically vulnerable students and to expand low-income students' access to subsidized private schools that charge top-up fees. The program is voluntary for subsidized providers: schools that opt in receive an additional per-pupil subsidy for each eligible student they enroll, but must charge zero top-up fees to those students. Participating private-voucher schools may still charge fees to higher-income students, but do not receive the targeted subsidy for them. Schools that opt out continue to receive only the universal voucher and can charge top-up fees to all students.

Descriptives show that by 2013, virtually all public schools had joined the program and about 72% of private-voucher schools had opted in. Participation is strongly related to pre-reform fees: schools charging less than \$300 (annual) were very likely to join, whereas those charging \$750 or more were unlikely to have joined by 2013, with a negative relationship between fees and participation in the middle of the distribution (\$300–\$750). Participation also increases with the local share of eligible students, with a positive slope for shares between 0 and 60%, and near-full participation above that threshold.

In line with its motivating goals, the reform expanded access to previously higher-fee subsidized private schools for some low-income students. Relative to the pre-reform period, vulnerable students became more likely to enroll in subsidized private schools and to attend schools with positive fees. Notably, the public-to-private enrollment split for low-income students shifted from 58%/42% in 2007 to 48%/52% six years after the reform.

Motivated by the descriptive evidence, I develop and estimate an equilibrium model of school demand and supply, where the demand side has atomistic families choosing schools that are vertically and horizontally differentiated, and that compete by simultaneously choosing whether they participate in the targeted voucher program and the fees they charge on top of the subsidies in each participation regime.

On the demand side, I document that both low- and high-income households value school quality and proximity, but differ markedly in price sensitivity and in their valuation of program participation. Low-income households are very price-sensitive and value free or heavily subsidized options, but have a relatively modest willingness to pay for improvements in observed school quality. High-income households are less price-sensitive, yet place substantial weight on avoiding schools that participate in the targeted voucher program, generating a sizable “stigma” effect attached to participation. For example, high-income families with highly educated mothers are willing to pay over \$18,000 to avoid enrolling in a participating school, while low-income families would pay less than \$70 for an additional standard deviation in school quality.

On the supply side, I estimate schools’ cost structure within a model in which private-voucher schools make endogenous decisions along both the extensive margin (whether to participate in the targeted program) and the intensive margin (top-up fees), and find substantial heterogeneity in costs. The fixed cost of joining the targeted program averages \$23,600 but rises steeply with school quality, reaching over \$32,000 for high-quality schools—explaining why many of them opt out. Marginal costs differ by student socioeconomic group: educating a low-income student costs about \$395 on average, well below the targeted voucher amount (\$706), while educating a high-income student costs roughly \$1,000. Schools also place significant weight on enrollment in their objective function, valuing an additional student similarly to a nontrivial increase in operating margins.

My results show that the targeted voucher program meaningfully changes schools’ equilibrium strategies relative to a system with only a universal voucher. A sizable share of private-voucher schools opt in to the program, but participation is strongly negatively correlated with school quality: higher-quality schools face substantially higher fixed participation costs and therefore are less likely to join. Participating schools reduce top-up fees—particularly for high-income students—while non-participants raise them to screen families averse to participation, leading to a segmented market in which lower-quality schools tend to be low-price participants and higher-quality schools tend to be high-price non-participants.

These equilibrium responses translate into heterogeneous effects for students. When the program is introduced, low-income students benefit from sharply lower effective fees and modestly higher average school quality. High-income students reallocate toward non-participating, higher-quality schools and face higher effective prices on average.

I use the estimated model to evaluate policy counterfactuals. I find that modest, budget-neutral redesigns of the targeted voucher—such as quality-differentiated per-student subsidies or a fixed participation incentive paired with a lower per-student voucher—can raise the average quality of participating and tuition-free schools and reduce effective prices. However, these redesigns have limited effects on the quality of schools actually attended by low-income students, because low-income households exhibit low willingness to pay for observed school attributes and therefore respond primarily to large changes in prices or school characteristics. In contrast, a full expansion that mandates participation increases access (all private-voucher schools become free for eligible students) and raises the quality of tuition-free options, but at a nontrivial fiscal cost and with notable losses for high-income families who lose access to high-fee, high-quality non-participant schools.

Methodologically, the estimation of the model poses several econometric challenges. First, the model may admit multiple equilibria. Second, the data include many schools charging zero top-up fees, which are the result of corner solutions in schools’ best-response functions. Third, since the educational markets in my data consist of tens—and sometimes hundreds—of schools, solving the discrete-continuous game on the supply side becomes rapidly computationally expensive. I estimate the model using a GMM-MPEC algorithm that minimizes a GMM objective function subject to market equilibrium conditions.¹ As is well known, MPEC estimation is robust to multiple equilibria, since equilibrium constraints need to hold only at the optimum, not at intermediate iterations. I further implement censored regression models to allow for potential corner solutions in schools’ best-response functions for top-up fees. Lastly, I adopt a parsimonious and well-established equilibrium concept for the supply-side game: *fully cursed equilibrium* (Eyster and Rabin, 2005), which simplifies the behavioral perception of schools, while maintaining economically sound predictions and goodness of fit. Combining fully cursed equilibrium with the MPEC algorithm ensures computational tractability.

In this setting, fully cursed equilibrium implies that each school correctly predicts the distribution of other schools’ actions, but ignores the correlation between these actions and types—where a school’s type is the unobserved part of its cost structure. My empirical application of fully cursed equilibrium adds important flexibility to the theory, motivated by the context, the data, and my own conversations with school principals and managers. While the fully cursed equilibrium theory assumes no correlation between other schools’ actions and types and characteristics, I allow schools to partially predict their competitors’ actions based on a set of observable characteristics. School heterogeneity is thus captured by observables, and strong/weak competitors can be iden-

¹GMM-MPEC stands for Generalized Method of Moments-Mathematical Programming with Equilibrium Constraints. See Dubé et al. (2012) and Su and Judd (2012) for the theory underlying MPEC, and early applications in economics.

tified. More specifically, schools are allowed to take into account the correlation between other schools' actions and the observable part of their cost structure. Fully cursed equilibrium is then applied to the unobserved determinants of actions, after controlling for observable characteristics.

This paper builds on a rich prior literature. Neilson (2025) is possibly my closest predecessor. Neilson (2025) develops a demand and supply model of school competition to study the competitive and quality distribution consequences of the introduction of the targeted voucher program in Chile. I closely follow Neilson (2025) in my demand modeling. However, in contrast to Neilson (2025), who uses observed equilibria in the data and focuses on the quality effects of the policy, I explicitly solve and estimate a supply-side game, and focus on the participation decision of schools, a margin that has been largely overlooked.

This paper also builds on previous studies that develop supply-side equilibrium models to examine the effects of student aid policy in education. Allende (2019), Allende et al. (2019), and Bau (2022) estimate continuous strategy supply-side games where schools compete in quality and/or price to attract students. Similarly, Ferreyra and Kosenok (2018), Singleton (2019), and Dinerstein and Smith (2021) estimate discrete strategy games, where schools choose entry location and/or exit to maximize a profit-based objective function. This paper is part of a growing set of studies that unite these two strands of the literature by allowing schools to simultaneously choose discrete and continuous strategies in a static supply-side game (Barahona et al., 2025a; Pal, 2025).

More generally, this paper contributes by offering an equilibrium framework that combines a rich demand system with a supply model that endogenizes both extensive and intensive margins, which is essential for evaluating policies with voluntary supplier participation. Such policies are common across the education sector and beyond. For instance, Abdulkadiroglu et al. (2018) attribute the large negative achievement effects of the Louisiana Scholarship Program to the program's inability to attract high-quality schools. In Florida, Singleton (2019)'s findings underscore the role of school entry and exit decisions in mediating equity effects of charter school policies: cost-adjusted—as opposed to flat per-pupil—funding formulae incentivize school entry in poor neighborhoods, which in turn increases the share of disadvantaged students served. Also from the US, Dinerstein and Smith (2021) show how the New York City's Fair Student Funding reform triggered significant entry and exit responses in the private sector, that played a critical role in shaping the program's enrollment and welfare effects. In Chile, a teacher recruitment program incentivized college participation via generous scholarships, but it simultaneously reduced participation with strict admission cutoffs (Neilson et al., 2022; Pal, 2025). In Brazil, Barahona et al. (2025b) document the extent to which the effectiveness of the federal student loan program depends not only on how many students qualify for loans but also on which colleges choose to participate, which is determined by the set of rules colleges must comply with. Similar are the

takeaways from Chile’s implementation of a free college policy (Bucarey, 2018; Johnson, 2023).

Policies triggering extensive margin responses from suppliers are not exclusive of the education sector. The case studied in Wollmann (2018) is stark. The paper shows how product repositioning decisions are pivotal for evaluating antitrust and bailout policies. Such decisions determine whether consolidation leads to persistent market power or whether competitive forces restore variety and discipline prices. In fact, accounting for product entry and exit flips the actual policy conclusion: the 2009 auto bailout looks necessary if entry and exit are ignored but unnecessary if they are considered. In the pharmaceutical market, Atal et al. (2025) show that introducing a minimum quality standard in Chile forced low-quality generics to leave the market, and incentivized the entry of high-quality drugs, producing a clear quantity-quality trade-off: reduced market size increases prices, while a larger pool of high-quality suppliers enhances consumer welfare. In Lagos, the rollout of a public bus network triggered extensive margin adjustments by private minibuses operators, which significantly shaped commuter welfare: treated routes saw a reduction in minibus departures, while private drivers’ rerouting resulted in lower fares in untreated but connected routes (Bjorkegren et al., 2025). In credit markets, the perceived success of Chile’s government-guaranteed loans program during the COVID-19-induced crisis can be attributed to banks’ strategic screening and allocation decisions (Chittaro and Sánchez, 2025; de Elejalde and Sánchez, 2025) that resulted in high coverage with low aggregate risk (Huneus et al., 2025).

Finally, this paper’s empirical results contribute to the evidence on the effects of voucher programs, reviewed by Epple et al. (2017).

The remainder of the paper is organized as follows. Section 2 describes Chile’s educational system and the targeted voucher program. Section 3 presents the model. Section 4 details the data and market construction. Section 5 discusses estimation and identification. Section 6 reports results and policy simulations. Section 7 concludes.

2 The Educational System and Vouchers in Chile

Since 1981, Chile’s primary and secondary education system has operated under a nationwide school-choice voucher program. Families are free to choose schools, subject only to travel costs and any top-up fees charged by the school. Schools, in turn, receive government funding through per-student vouchers, with the exact rules depending on their administrative category.

Schools can be grouped into three administrative categories, based on ownership and management and eligibility to receive subsidies. *Public* schools are owned and managed by local municipalities, are funded through vouchers, and are tuition free. *Private-voucher* schools are privately owned and managed, receive vouchers, and may charge top-up fees above the voucher subsidy. About 53% of these schools charged positive top-up fees in 2013, which averaged ~\$500

per year. Top-up fees are capped at 4 Educational Subsidy Units per month (USE, for its acronym in Spanish), or about \$1,780 annually.² The third group is *private-non-voucher* schools, which are privately owned and managed, receive no subsidies, and are financed entirely through tuition. On average, these schools charged ~\$6,400 in 2013.

The vast majority of primary-school students attend a subsidized school. In 2013, 40% attended a public school, 52% attended a private-voucher school, and only 7% attended a private-non-voucher school. These figures highlight both the sizable role of private education—especially the subsidized private sector—and the broad reach of vouchers: more than 90% of students use them to attend either a public or a private-voucher school.

The government subsidizes schools through two voucher programs. The first is the *universal* voucher: a per-student subsidy paid to all public and private-voucher schools based on enrollment. In 2013, the universal voucher was about \$1,202 per student per year.³ The second program is the *targeted* voucher, introduced in 2008 to provide additional funding for students from disadvantaged backgrounds. A central motivation was to expand the school choice set of low-income students, who were thought to be budget-constrained when considering high-quality, high-fee subsidized private schools. The targeted voucher requires participating schools to charge zero top-up fees to eligible low-income students. In 2013, the targeted subsidy was about \$706.

A distinctive feature of the targeted voucher program is that participation is voluntary. Schools that opt in receive the targeted subsidy for each eligible low-income student, and must charge zero top-up fees to those students. Participating schools may still charge top-up fees to higher-income students, but they do not receive the targeted subsidy for these non-eligible students. In addition, upon joining, participating schools must prepare an improvement plan with ambitious but realistic learning goals, subject to medium-term evaluation by the Ministry of Education. In practice, many schools hire specialized education consulting firms called ATE (*Asistencia Técnica Educativa*) to help prepare these plans. The targeted voucher is paid based on student attendance, which must be reported to the government on a monthly basis; as a result, participation in the program translates into an increase in the day-to-day administrative burden on schools. Schools that do not join continue to receive only the universal voucher and may charge top-up fees to all students. By 2013, virtually all public schools had joined the targeted program, and ~72% of private-voucher schools had opted in.

Eligibility for the targeted voucher is primarily income-based. Children from families in the first tercile of the income distribution are eligible, as are children from families participating in the *Chile Solidario* social program. In 2013, 52% of primary-school students were eligible for the

²In practice, this cap is not binding for the vast majority of schools. In 2013, only four private-voucher schools charged 4 USE or more per month.


³This amount slightly varies by grade and by urban status of the school.

targeted voucher.

Consistent with its stated goals, the targeted voucher expanded access to private schools that previously charged fees for some low-income students. Before the reform, more than 75% of low-income students attended zero-fee schools. By 2013, the student at the 75th percentile of the distribution of low-income students—ranked by the pre-reform fees of the school they attend—was enrolled in a school that had charged \$103 prior to the reform. Moreover, whereas 58% and 42% of low-income students attended public and subsidized private schools in 2007, respectively, six years after the reform the distribution had shifted toward the subsidized private sector: 48% attended public schools and 52% attended private-voucher schools.

On the supply side, private-voucher schools' participation decisions are strongly correlated with both pre-reform top-up fees and the local share of eligible low-income students. Panel A of Figure 1 shows that schools charging very low fees before the reform were much more likely to participate. Conversely, schools with very high pre-reform fees were unlikely to have joined by 2013; among schools in the \$300–\$750 range, participation declines with fees. Panel B shows a similar pattern with respect to local demographics: schools in municipalities with a higher pre-reform share of eligible students participate at higher rates, and in municipalities where 60% or more of students are eligible, almost all private-voucher schools participate. Panel C combines these relationships in a contour plot of participation probabilities as a function of both fees and the local share of eligible students (bluer colors indicate lower participation, greener colors indicate higher participation). Three patterns stand out. First, regardless of the share of eligible students, almost all schools with fees below \$300 join the program. Second, schools charging \$750 or more have close-to-zero participation. Third, in the middle of the fee distribution, participation is partly determined by the prevalence of low-income students, conditional on fees. Panel C also shows that top-up fees are lower in municipalities with higher shares of vulnerable students, underscoring a tight link between locality characteristics and equilibrium outcomes.

Figure 1: Private-Voucher Schools' Program Participation



Notes: Panel A plots a nonparametric estimate of private-voucher schools' participation in the targeted program against the school's pre-reform top-up fees. Panel B plots the analogous relationship against the pre-reform share of eligible low-income students in the school's municipality. Panel C shows a contour plot of the predicted probability that a private-voucher school has joined the program by 2013 as a function of pre-reform fees and the pre-reform local share of eligible students.

The patterns in Figure 1 are consistent with private-voucher schools behaving as profit seekers (or, more generally, as having an objective function that places substantial weight on profits). To see why, consider the revenues of a subsidized school that does not participate in the targeted program. For each student, the school receives the universal voucher plus any top-up fees it charges. If the school instead participates, its revenue from non-vulnerable students is unchanged; however, for each low-income student it receives the universal voucher plus the targeted voucher and must charge zero fees. Participation therefore trades off fee revenue from low-income students against the targeted subsidy. Under this logic, it is not surprising that low-fee schools tend to opt in, whereas high-fee schools tend to opt out. The incentive is stronger in areas where a larger share of demand comes from low-income students.

This intuition motivates the simple exercise reported in Table 1. For each private-voucher school, I compute an ex-ante indicator of whether participation is profitable, using information available one year before the reform: the school's posted fees, its share of low-income students, and the targeted voucher value. Specifically, the indicator is:

$$\mathbb{1} \left\{ \# \text{ low-income students} \times (\text{targeted voucher} - \text{top-up fees}) > 0 \right\}.$$

Table 1 then counts schools with strictly positive versus non-positive net benefits and cross-tabulates them by actual participation. The results align closely with the incentive logic: 83% of schools for which participation is ex-ante profitable opt in, while 84% of schools for which

participation is not profitable opt out.

Table 1: Program Participation by Ex-Ante Net Benefit of Participating

	in program	not in program	total
ex-ante profit > 0	2,116	442	2,558
ex-ante profit ≤ 0	53	287	340
total	2,169	729	2,898

Notes: This table displays private-voucher schools' participation decisions as a function of the following measure of ex-ante net benefit of joining the program: $\mathbb{1}\{\#\text{ low-income students} \times (\text{targeted voucher} - \text{top-up fees}) > 0\}$, where low-income students and top-up fees are measured in the year preceding the introduction of the reform. Only private-voucher schools are considered.

To complement these figures, I relate each school's participation decision to the ex-ante profit indicator in a simple linear regression. The regression has an R^2 of 0.25, indicating that this rule-of-thumb measure alone explains a sizable share of the variation in participation. Taken together, the descriptive patterns and this regression evidence suggest that managers at private-voucher schools made participation decisions that were strongly shaped by the underlying profit trade-off, and that the simple ex-ante indicator provides a useful approximation to that trade-off.⁴

I close this section by examining the quality dimension of schools' sorting into the program. If educational quality is costly to produce, then part of the fees schools charge should reflect these quality-related costs; as a result, fees and educational quality are likely to be positively correlated. As shown above, it was mostly low-fee schools that chose to join the targeted program. A simple—though admittedly noisy—proxy for school quality is students' average test scores. Private-voucher schools that had joined the program by 2013 have a pre-reform average test score of -0.16 standard deviations (s.d.), whereas schools that opted out have a much higher pre-reform average test score, of about 0.34 s.d. Turning to the schools attended by low-income students, one year before the reform the median low-income student attended a school with an average test score of -0.19 s.d. Six years later, the median low-income student was enrolled in a school with a pre-reform average test score of -0.11 s.d. Taken together, the strong sorting by schools on this quality measure, and the seemingly modest improvement in the quality of schools attended by low-income students, raise doubts about whether the reform met its stated goal: expanding low-income students' access to high-quality private schools. Instead, the reform may have primarily increased access to private but low-quality schools. Lessons from the Louisiana Scholarship Program suggest that this latter outcome is plausible, with undesirable consequences

⁴In a context where supply-side agents consider discrete-continuous strategies (product entry/exit and pricing), Wollmann (2018) describes a similar rule-of-thumb threshold explaining commercial vehicle companies' decisions to enter or exit a model.

for students’ learning (Abdulkadiroglu et al., 2018).

3 The Model

I develop a structural model of demand for and supply of schools motivated by the voucher programs for primary education in Chile. The model is static. There are several education markets that are geographically separated from one another. Each market is populated by households that live in different locations within the market, each with one child attending primary education. Given its budget constraint, each household chooses among the schools available in the market.

Schools are (exogenously) distributed within the market’s area. Schools fall into three categories: public, private-voucher, and private-non-voucher. Public schools are tuition-free, while both private-voucher and private-non-voucher schools are allowed to charge fees. Public and private-voucher schools receive a per-student flat subsidy, the universal voucher. In addition, a complementary subsidy program is available for public and private-voucher schools: a targeted voucher for economically disadvantaged students. Participation in this targeted program is optional for schools. The targeted voucher program adds extra per-pupil funds over the universal voucher for every eligible low-income student that the school enrolls, requiring the school to charge zero fees to those students. Participating private-voucher schools can still, however, charge top-up fees to higher-income students. Non-participating private-voucher schools receive the universal voucher and their posted fees for every student. Private-non-voucher schools do not receive any subsidy, and charge a uniform fee to all students.

3.1 Demand

Students have heterogeneous preferences over schools’ fees, quality (i.e., a measure of how much the school increases students’ test scores), program participation status, geographical proximity, and a set of school observable and unobservable (to the econometrician) characteristics. I capture heterogeneity in preferences with a set of coefficients that vary with students’ observed demographic characteristics. Formally, in each market $m \in \{1, \dots, M\}$, student $i \in \{1, \dots, I\}$ chooses the school $j \in \{1, \dots, J\}$ that maximizes her utility. I specify the student’s conditional indirect utility as:⁵

$$U_{ij} = \beta_{1i} p_j^\zeta + \beta_{2i} d_{ij} + \beta'_{3i} X_j^\beta + \beta_4^\zeta q_j + \beta_{5i} \tau_j + \xi_j^\zeta + \varepsilon_{ij} \quad (1)$$

⁵I suppress the market subscript m for ease of exposition.

where the superscript $\zeta \in \{L, H\}$ refers to the eligibility status of the student, i.e., low-income or high-income. Thus, p_j^ζ is school j 's fees charged to students of eligibility group ζ , d_{ij} is distance from student i 's home to school j , X_j is a vector of school j 's observable characteristics, q_j is school j 's quality measure, τ_j is an indicator for school j 's participation in the targeted voucher program, ξ_j^ζ is the common preference that students of group ζ have for school j 's unobservable (to the econometrician) characteristics, and ε_{ij} is an i.i.d. preference shock. Also, for any $\beta^\zeta \in \{\beta_4^\zeta, \xi_j^\zeta\}$, let $\beta^\zeta = D_i\beta^L + (1 - D_i)\beta^H$, where $D_i = \mathbb{1}[i \text{ is low-income}]$. Similarly, for $k = 1, 2, 3, 5$, $\beta_{ki} = D_i\beta_{ki}^L + (1 - D_i)\beta_{ki}^H$, where $\beta_{ki}^L = \beta_k^L + \sum_r z_{ir}\beta_{kr}^L$ and $\beta_{ki}^H = \beta_k^H + \sum_r z_{ir}\beta_{kr}^H$, with z_{ir} being student i 's demographic characteristic r .

Note that the fees school j charges student i , p_j^ζ , depend on whether the student is low-income, and on whether the school participates in the targeted voucher program. Specifically,

$$p_j^\zeta = \tau_j(1 - D_i)p_j^1 + (1 - \tau_j)p_j^0,$$

where p_j^1 is school j 's counterfactual fees when the school participates in the targeted program, and p_j^0 is school j 's counterfactual fees when the school does not participate in the targeted program.

Let $V_{ij} = \beta_{1i}p_j^\zeta + \beta_{2i}d_{ij} + \beta'_{3i}X_j^\beta + \beta_4^\zeta q_j + \beta_{5i}\tau_j + \xi_j^\zeta$. Thus, $U_{ij} = V_{ij} + \varepsilon_{ij}$. Assuming $\varepsilon_{ij} \sim$ Type I Extreme Value, the probability that student i chooses school j is logistic:

$$P_{ij} = \frac{e^{V_{ij}}}{\sum_k e^{V_{ik}}}.$$

While other studies explicitly model students' preferences for peers (Ferreyra and Kosenok, 2018; Allende, 2019; Abdulkadiroglu et al., 2020; Crema, 2024; Li et al., 2025), I incorporate this feature by including the school's decision to participate in the program directly in the indirect utility function. This modeling choice captures students' taste for peers implicitly, without requiring an aggregate game on the demand side. For instance, if high-income students prefer other high-income peers and a school attracts many low-income students by joining the program, then the school would lose high-income enrollment through peer preferences. In my model, this effect would show up as a negative coefficient for high-income households on the school's participation decision, making the model sensitive to this feature of the data. Moreover, this approach has attractive computational implications—it avoids solving a game on the demand side—and allows me to define a rich and realistic supply-side school competition game, which is the focus of this paper.

3.2 Supply

3.2.1 Schools' Problem.

I model strategic decisions only for private-voucher schools. I do so because the variation in the data needed for econometric identification is available only for these schools. For instance, the data show that virtually all public schools have joined the targeted program, leaving insufficient variation to identify the participation margin for these schools. Additionally, public schools are required by law to charge no fees. Other related studies adopt similar modeling decisions (see, e.g., Allende et al., 2019).

I assume private-voucher schools make their strategic decisions in the context of a static game of incomplete information. Let $a_j = (\tau_j, p_j^1, p_j^0)$ denote the actions available to school j , where $\tau_j \in \{0, 1\}$ is school j 's decision to participate in the targeted program, $p_j^1 \geq 0$ is the price the school sets if it joins the program, and $p_j^0 \geq 0$ is the price it sets if it does not join the program.⁶ Define, too, school j 's conditional (on program participation) vectors of actions, $a_j^1 = (\tau_j = 1, p_j^1, p_j^0)$ and $a_j^0 = (\tau_j = 0, p_j^1, p_j^0)$. Let κ_j be school j 's fixed cost of participating in the program. This cost reflects the additional administrative requirements and government monitoring that participation entails, as well as how closely the program's objectives are aligned with the school's own mission. In addition, let (c_j^L, c_j^H) be school j 's marginal costs of educating low- and high-income students. I collect all cost terms in $t_j = (\kappa_j, c_j^L, c_j^H)$. I assume t_j is private information; i.e., t_j is school j 's type. Finally, denote school j 's pure strategy as $s_j(t_j) = (\tau_j(t_j), p_j^1(t_j), p_j^0(t_j))$.

Private-voucher school j of type t_j chooses its actions to maximize an objective function that combines expected profits and enrollment. This class of objective functions for service providers in the education and health sectors, which combine profits with a measure of quantity or quality, is well motivated (Lakdawala and Phillipson, 1998; Gaynor and Town, 2011) and is commonly adopted in the empirical literature (Hackmann, 2019; Singleton, 2019).⁷ Specifically,

$$\max_{\tau_j \in \{0,1\}, p_j^1 \geq 0, p_j^0 \geq 0} E_{t_{-j}} [\tau_j (\mathcal{U}_j^1(a_j^1, s_{-j}(t_{-j}); t) - \kappa_j) + (1 - \tau_j) \mathcal{U}_j^0(a_j^0, s_{-j}(t_{-j}); t)],$$

⁶As described in Section 2, in Chile a maximum cap on the top-up fees private-voucher schools can charge to families exists. However, this cap is rarely binding, with only four schools exceeding the cap, according to the administrative data. Moreover, these four schools' fees are observed to be higher than the cap, suggesting that the upper limit is in practice not enforced for these schools. I therefore decide not to explicitly model the maximum cap on top-up fees.

⁷See also Blanchard et al. (2025) for an application to pension fund administrators.

where,

$$\begin{aligned}
\mathcal{U}_j^1(a_j^1, s_{-j}(t_{-j}); t) &= (p_j^1 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^1, s_{-j}(t_{-j}); t) \\
&\quad + (v^t + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^1, s_{-j}(t_{-j}); t) + \psi P_{ij}(a_j^1, s_{-j}(t_{-j}); t), \\
\mathcal{U}_j^0(a_j^0, s_{-j}(t_{-j}); t) &= (p_j^0 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \\
&\quad + (p_j^0 + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^0, s_{-j}(t_{-j}); t) + \psi P_{ij}(a_j^0, s_{-j}(t_{-j}); t).
\end{aligned}$$

The expression for school j 's objective function consists of two parts. The first part, $(\mathcal{U}_j^1 - \kappa_j)$, is a linear combination of profits and enrollment associated with joining the program, net of participation cost κ_j . In this regime, per-student profits vary depending on the socioeconomic background of the enrolled student. For each high-income student, the school receives the fees it charges, p_j^1 , the universal voucher, v^u , and incurs a marginal cost of c_j^H . For each low-income student, the school receives the targeted and universal vouchers, $v^t + v^u$, and incurs a marginal cost of c_j^L , which is allowed to differ from c_j^H . Added to profits is expected enrollment, whose relative weight is captured by the parameter $\psi > 0$. The second part of school j 's objective function, \mathcal{U}_j^0 , is a linear combination of profits and enrollment associated with not joining the targeted program. In this case, per-student revenue consists of the fees it charges (p_j^0) plus the universal voucher (v^u), regardless of students' socioeconomic status. Marginal costs are c_j^H for high-income students, and c_j^L for low-income students. Enrollment is weighted by ψ .

3.2.2 Motivations for Equilibrium Definition

The standard equilibrium notion for this class of games is Bayesian Nash equilibrium (BNE). This equilibrium concept assumes that schools do not observe the types of their opponents but know the probability distribution from which each type is drawn. Moreover, schools expect that their opponents choose their actions according to their types. That is, schools anticipate a correlation between other schools' actions and types.

In large discrete-continuous games, BNE involves the computation of many counterfactual equilibria in the continuous strategies to form the expected utility function of each player. This procedure quickly becomes computationally expensive and demanding in terms of players' cognition.

For illustration, note that in a BNE framework, school j 's belief about the expected payoffs

in the not-in-program regime is,

$$\begin{aligned}
E_{t_{-j}}^{BN} [\mathcal{U}_j^0(a_j^0, s_{-j}(t_{-j}); t)] &= \int_{\mathcal{S}_{t_{-j}}} \left[(p_j^0 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \right. \\
&\quad + (p_j^0 + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \\
&\quad \left. + \psi P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \right] dF(t_{-j}|t_j),
\end{aligned}$$

where $\mathcal{S}_{t_{-j}}$ is the space of other schools' types, and $F(t_{-j}|t_j)$ is the probability distribution of other schools' types from the perspective of school j of type t_j . Using iterated expectations, school j 's expected profits in the not-in-program regime can be rewritten as,

$$\begin{aligned}
E_{t_{-j}}^{BN} [\mathcal{U}_j^0(a_j^0, s_{-j}(t_{-j}); t)] &= \sum_{\tau_{-j}} \left\{ \int_{\mathcal{S}_{c_{-j}}} \left[(p_j^0 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \right. \right. \\
&\quad + (p_j^0 + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \\
&\quad \left. \left. + \psi P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \right] dG(c_{-j}|\tau_{-j}; t_j) \right\} Pr(\tau_{-j}|t_j),
\end{aligned}$$

where $\mathcal{S}_{c_{-j}}$ is the space of other schools' marginal costs, and $G(c_{-j}|\tau_{-j}; t_j)$ is the probability distribution of other schools' marginal costs conditional on other schools' program participation decisions, τ_{-j} , from the perspective of school j of type t_j . The term inside the brackets is therefore the conditional expected payoffs in the not-in-program regime for a particular value of τ_{-j} , and the summation takes the expectation of those payoffs over all possible values of τ_{-j} . Note that, for a given τ_{-j} , conditional expected payoffs are a function of equilibrium counterfactual fees, $(p^1(\tau_j = 0, \tau_{-j}), p^0(\tau_j = 0, \tau_{-j}))$. Thus, to compute the expected payoffs in the not-in-program regime, we first need to obtain the equilibrium fees and the corresponding conditional expected payoffs for each τ_{-j} . Then, we take the weighted sum of all conditional expected payoffs, where the weights are the probabilities of each τ_{-j} . This is a demanding task in large games. For instance, in a market with $J - 1 = 10$, 1,024 different sets of equilibrium fees need to be obtained. In a market with $J - 1 = 20$, 1,048,576 different sets of equilibrium fees need to be obtained. And in a market more typical of the setting I study in this paper, with, say, $J - 1 = 40$, $5.498e^{11}$ different sets of equilibrium fees need to be obtained.⁸ This procedure becomes computationally intractable very fast.

A number of alternative equilibrium concepts for static games of incomplete information provide equilibrium solving procedures that are computationally lighter than BNE. Most of these concepts rely on simplifications of agents' perceptual assumptions. Models with agents that

⁸In general, 2^{J-1} equilibrium fees and associated payoffs need to be computed.

overlook the correlation between others' actions and private types are both theoretically and computationally attractive. One example is Eyster and Rabin (2005)'s concept of fully cursed equilibrium.⁹ Fully cursed agents have correct beliefs about the distribution of others' types, but fail to account for the correlation between other players' types and actions. More practically, fully cursed agents best-respond to their opponents' (ex-ante) average strategy. Fully cursed equilibrium is a limit case of Eyster and Rabin (2005)'s more general concept of (partial) cursed equilibrium, which assumes that each player best-responds to a convex combination of others' actual strategies and the aggregate distribution. A disadvantage of this intermediate case of cursed equilibrium is that it is hard to imagine a learning process that leads to such a concept (Eyster and Rabin, 2005; Fudenberg, 2006). By contrast, fully cursed equilibrium is founded on a well-defined learning process, in which agents observe others' actions but neither their types nor their payoffs.

Fully cursed equilibrium is related to other equilibrium concepts of games with agents that fail to account for the informational content of others' play. Self-confirming equilibrium (Fudenberg and Levine, 1993; Dekel et al., 2004; Battigalli et al., 2015) does not assume that players have correct beliefs about the distribution of opponents' play, but only that the beliefs are consistent with what players observe when the game is played—i.e., beliefs are correct only along the equilibrium path. Fully cursed equilibrium corresponds to a self-confirming equilibrium where players only observe the aggregate distribution of others' actions but not their types. Relatedly, Esponda (2008) combines self-confirming equilibrium with some monotonicity restrictions to study the impact of naive agents on equilibrium play in adverse selection models. Jehiel and Koessler (2008)'s analogy-based expectation equilibrium (ABEE) assumes that players bundle states into analogy classes and best-respond to opponents' average strategy in those analogy classes.¹⁰ Analogy classes are simply a coarse partition of the distribution of states and types. Fully cursed equilibrium is a special case of ABEE, in which all states are bundled into a common analogy class. Moreover, ABEE can be viewed as a natural selection of self-confirming equilibrium, in which the signals received by players after each round of play correspond to the average play in each analogy class. In the same vein, Esponda and Pouzo (2016)'s Berk-Nash equilibrium presents a unifying framework of equilibrium models of games with agents with misspecified views of their environment, which includes Bayesian Nash equilibrium, self-confirming equilibrium, ABEE, and fully cursed equilibrium as special cases.

All of the above-mentioned theories constitute a response to a body of evidence on bounded rationality that is both economically significant and regular enough to be modeled (Fudenberg,

⁹For early developments of the concept of fully cursed equilibrium, see Kagel and Levin (1986), and Holt and Sherman (1994).

¹⁰See, also, Jehiel (2005), and Spiegel (2011).

2006). The winner’s curse has been documented in experimental and non-experimental studies of auctions and trade with adverse selection (Kagel and Levin, 1986, 2002; Thaler, 1988; Holt and Sherman, 1994; Charness and Levin, 2009).¹¹ Individuals’ failure to account for the informational content of other people’s actions has also been documented empirically in contexts of voting in elections and juries (Converse, 2000; Esponda and Vespa, 2014). More importantly, Kagel and Levin (1986), and Eyster and Rabin (2005) find that these phenomena are more likely to arise in large games, highlighting the limits of human cognition and the difficulty of evaluating counterfactual situations as the number of players/types/states grows—similar to the storage and memory limits of a personal computer.¹²

To gain insight into the perceptions and behavior of schools in my sample, I interviewed a group of managers in the private-voucher sector across different markets in the country.¹³ The interviews were structured around the following questions:

- How does the school decide on its fees?
- What motivated the school to join or not the targeted voucher program?
- Does the school know who its closest competitors are?
- Does the school monitor its closest competitors’ strategies (e.g., fees and program participation) when making its own decisions?
- How well does the school know its local demand?

Across cases, a common pattern emerges: managers maintain detailed knowledge of their market and demand, but they do not track or react to the actions of individual rival schools when setting strategies. Instead, strategic decisions are anchored in internal objectives (e.g., mission, quality improvement) and in aggregate market signals (e.g., relative position/rank, realized and potential demand).

Three cross-cutting observations summarize the interviews:

- i. Rival-specific monitoring is limited. One manager emphasizes that his school does not identify a “direct competitor” among private-voucher schools and explicitly avoids conditioning strategy on others’ behavior; he frames competition as “with ourselves,” focused on continuous quality improvements. A second manager similarly reports making decisions with reference to the “market as a whole” rather than to any specific rival or small subset of

¹¹The winner’s curse is generally defined as the winner’s disappointment after the auction or trade takes place (Thaler, 1988).

¹²See related discussions in Jehiel (2005), Jehiel and Samet (2007), Jehiel and Koessler (2008), and Esponda and Vespa (2014).

¹³I talked to managers from Colegio Inmaculada Concepción in Puerto Varas, Colegio San Francisco Javier in Cerro Navia, Colegio Augusto Winter in Temuco, Colegio Santa Marta in La Unión, Colegio San Miguel in Calbuco, Colegio Sagrada Familia in Hornopirén, Liceo San Conrado in Futrono, and Colegio María Deogracía in Futrono.

rivals. Another school leader notes that pricing and program participation choices are not based on close monitoring of other schools.

- ii. Market observables are well known. Managers can list competing schools, their locations, religious status, program focus (e.g., technical, vocational, general), grade/level offerings, the presence of full day shift, and whether they participate in the targeted program or charge fees (though not necessarily exact price levels). One school manager highlights awareness of the school’s rank in standardized tests and of both potential and realized demand. Another stresses knowing which schools are more or less preferred (by families) locally and having granular knowledge of demand.
- iii. Decisions are conditioned on aggregate signals and mission. A manager describes strategic choices as guided by the institution’s mission and by how demand responds, effectively summarizing the competitive environment into a small set of market-level statistics (e.g., own rank, demand, rivals’ fee/voucher program prevalence) rather than a vector of rival-specific actions.

Taken together, the interviews point to a “coarse” representation of the competitive environment: school managers track comprehensive market characteristics and their own market position but refrain from rival-by-rival best-responses. This interpretation helps rationalize limited tactical price or policy responses to individual rivals and suggests that competitive externalities may operate primarily through market aggregates (rankings, demand shifts, policy participation rates) rather than pairwise strategic interactions.

3.2.3 Definition of Equilibrium

I depart from the standard Bayesian Nash equilibrium framework, and instead assume that schools are fully cursed (Eyster and Rabin, 2005). As discussed above, fully cursed equilibrium is a well-accepted equilibrium concept for Bayesian games with agents that fail to account for the informational content of other players’ actions. It is a special case of several other related models, and is founded on a realistic learning process. It is parsimonious, and therefore computationally tractable.

As a first step toward defining fully cursed equilibrium in this game, let $\bar{s}_{-j}(t_j) \equiv \int_{\mathcal{S}_{i_{-j}}} s_{-j}(t_{-j}) dF(t_{-j}|t_j)$. That is, $\bar{s}_{-j}(t_j)$ is the average strategy of other players, averaged over other players’ types, from the perspective of school j of type t_j .

When schools are fully cursed, they (mistakenly) believe that each type profile of the other players plays the same action profile, $\bar{s}_{-j}(t_j)$, whenever they play $s_{-j}(t_{-j})$. Consequently, fully

cursed school j 's beliefs about the expected payoffs in each regime are such that,

$$\begin{aligned}
E_{t_{-j}}^{FC} [\mathcal{U}_j^1(a_j^1, s_{-j}(t_{-j}); t)] &= \mathcal{U}_j^1(a_j^1, \bar{s}_{-j}(t_j); t_j) \\
&= (p_j^1 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^1, \bar{s}_{-j}(t_j); t_j) \\
&\quad + (v^t + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^1, \bar{s}_{-j}(t_j); t_j) + \psi P_{ij}(a_j^1, \bar{s}_{-j}(t_j); t_j), \\
E_{t_{-j}}^{FC} [\mathcal{U}_j^0(a_j^0, s_{-j}(t_{-j}); t)] &= \mathcal{U}_j^0(a_j^0, \bar{s}_{-j}(t_j); t_j) \\
&= (p_j^0 + v^u - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j) \\
&\quad + (p_j^0 + v^u - c_j^L) \sum_i D_i P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j) + \psi P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j).
\end{aligned}$$

Define $\mathcal{U}_j(a_j, \bar{s}_{-j}(t_j); t_j) \equiv \tau_j(\mathcal{U}_j^1(a_j^1, \bar{s}_{-j}(t_j); t_j) - \kappa_j) + (1 - \tau_j)\mathcal{U}_j^0(a_j^0, \bar{s}_{-j}(t_j); t_j)$, i.e., the objective function for fully cursed school j of type t_j .

Definition 3.1 *In the static game of incomplete information described above, a pure strategy profile $\bar{s}^* = (\bar{s}_1^* \dots, \bar{s}_j^*)$ is a **fully cursed equilibrium** if and only if for each j and for each t_j ,*

$$a_j^* \in \arg \max_{a_j} \mathcal{U}_j(a_j, \bar{s}_{-j}^*(t_j); t_j).$$

In a fully cursed equilibrium, each school correctly predicts the probability distribution over its opponents' actions, but ignores the correlation between other schools' actions and their types. As a consequence, fully cursed equilibrium does not suffer from the curse of dimensionality problem present in BNE, since only one set of equilibrium fees (and program participation decisions) needs to be computed, regardless of the number of opponents.

When schools' types are independent—meaning that for each $t_j, t'_j, t_{-j}, F(t_{-j}|t_j) = F(t_{-j}|t'_j)$ —then each type of player j and any type of player k share common beliefs about the strategy of any player $l \neq j, k$; namely, $\bar{s}_l(t_j) = \bar{s}_l(t'_j) = \bar{s}_l(t_k) = \bar{s}_l(t'_k)$, for any t_j, t'_j, t_k, t'_k .

School j 's best-response functions satisfy,

$$p_j^1(a_j^1, \bar{s}_{-j}(t_j); t_j) = \max \left\{ 0, c_j^H - v^u - \frac{\sum_i (1 - D_i) P_{ij}(a_j^1, \bar{s}_{-j}(t_j); t_j)}{\sum_i (1 - D_i) \frac{\partial P_{ij}(a_j^1, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^1}} - \psi \right\}, \quad (2)$$

$$p_j^0(a_j^0, \bar{s}_{-j}(t_j); t_j) = \max \left\{ 0, c_j^H \frac{\sum_i (1 - D_i) \frac{\partial P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^0}} + \right. \\ \left. c_j^L \frac{\sum_i D_i \frac{\partial P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^0}} - v^u - \frac{\sum_i P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\sum_i \frac{\partial P_{ij}(a_j^0, \bar{s}_{-j}(t_j); t_j)}{\partial p_j^0}} - \psi \right\}, \quad (3)$$

$$\tau_j(a_j, \bar{s}_{-j}(t_j); t_j) = \mathbb{1} \{ (\mathcal{U}_j^1(a_j^1, \bar{s}_{-j}(t_j); t_j) - \kappa_j) - \mathcal{U}_j^0(a_j^0, \bar{s}_{-j}(t_j); t_j) > 0 \}, \quad (4)$$

where equations (2) and (3) are the first-order conditions for p_j^1 and p_j^0 , respectively (in addition to the corresponding Kuhn-Tucker complementary slackness conditions). Optimal p_j^1 is a function of the marginal cost of educating a high-income student (c_j^H), the universal voucher (v^u), the markup the school charges over the price under perfect competition with subsidy, and the relative weight of enrollment in the school's objective function. Note that the markup term depends on the price elasticity of demand of high-income students, and not on that of low-income students, since it is only high-income students that face this fee. Also, the markup is lower the more price-elastic is the demand, as is standard in imperfectly competitive markets. Both the universal voucher and the weight on enrollment enter equation (2) linearly, and therefore act as marginal cost reducers. Likewise, optimal p_j^0 depends on the marginal cost of educating a student, the universal voucher, a markup term, and ψ . Here, the marginal cost is a convex combination of c_j^H and c_j^L , with weights proportional to the share of each socioeconomic group in the population. The markup term depends on the price elasticity of demand of all (low- and high-income) students, since all students face this fee. Equation (4) is the optimality condition for τ_j , and states that school j joins the targeted program if and only if the net payoffs associated with joining the program are greater than the payoffs associated with not joining.

As stated in Definition 3.1, a fully cursed equilibrium in this game is the fixed point of all schools' first-order conditions.

3.2.4 Monte Carlo Simulations

To illustrate the properties of fully cursed equilibrium (FCE) relative to other equilibrium concepts, I carry out Monte Carlo simulations, whose details are consigned to Appendix C. I am mostly interested in the comparison of FCE and BNE, and in assessing how well FCE approxi-

mates BNE in the case the observed data were generated from a Bayesian Nash equilibrium. In addition, I study an intermediate concept, which I call partial cursed equilibrium (PCE), which assumes that each school is Bayesian with respect to its (geographically) closest opponent, and is fully cursed with respect to all others. This equilibrium concept closely relates to Eyster and Rabin (2005)'s definition of (partial) cursed equilibrium.

I consider a simplified school choice and competition model, very similar to the main model described above, and compare the equilibrium strategies implied by each equilibrium concept under the same primitives.

Model Specification. I posit a single market where $I = 500$ students enroll in their most preferred school among the J schools available in the market. I vary the total number of schools across simulations to analyze how the size of the game affects equilibria. I experiment with $J = 5, 10, 15, 20, 25, 30$.

Similar to Fack et al. (2019), I model the market as a disc with a radius of 1. Both schools and students are uniformly distributed throughout the market area. Figure C.1 in Appendix C displays the spatial configuration of an example of a simulated market.

On the demand side, I adopt a parsimonious version of the indirect utility function described in Section 3.1:

$$U_{ij} = \alpha_j^\zeta + \beta_1^\zeta p_j^\zeta + \beta_2^\zeta \tau_j + \beta_3^\zeta d_{ij} + \varepsilon_{ij},$$

where, the superscript $\zeta \in \{L, H\}$ denotes the student's income group. Similarly, $p_j^\zeta = \tau_j(1 - D_i)p_j^1 + (1 - \tau_j)p_j^0$, with D_i denoting student i 's low-income status, τ_j is school j 's decision to participate in the program, and (p_j^1, p_j^0) are school j 's counterfactual fees when the school joins the program and when it opts out, respectively. The parameter α_j^ζ represents school j 's attractiveness to students of group ζ , net of fees, program participation, and proximity. Likewise, β_1^ζ , β_2^ζ , and β_3^ζ capture the preferences of students of income group ζ for school j 's fees, program participation, and distance, respectively. The idiosyncratic error term ε_{ij} follows a Type I extreme value distribution.

On the supply side, for simplicity, I abstract from the universal voucher, and I assume that schools are profit-seeking. The school's problem is the following,

$$\max_{\tau_j \in \{0,1\}, p_j^1 \geq 0, p_j^0 \geq 0} E_{t-j} [\tau_j (\Pi_j^1(a_j^1, s_{-j}(t-j); t) - \kappa_j) + (1 - \tau_j) \Pi_j^0(a_j^0, s_{-j}(t-j); t)],$$

where,

$$\begin{aligned}\Pi_j^1(a_j^1, s_{-j}(t_{-j}); t) &= (p_j^1 - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^1, s_{-j}(t_{-j}); t) \\ &\quad + (v^t - c_j^L) \sum_i D_i P_{ij}(a_j^1, s_{-j}(t_{-j}); t), \\ \Pi_j^0(a_j^0, s_{-j}(t_{-j}); t) &= (p_j^0 - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \\ &\quad + (p_j^0 - c_j^L) \sum_i D_i P_{ij}(a_j^0, s_{-j}(t_{-j}); t).\end{aligned}$$

The cost structure terms, (τ_j, c_j^L, c_j^H) , are constructed as follows,

$$\kappa_j = \omega_1^\kappa + \omega_2^\kappa u_j^\kappa, \quad (5)$$

$$c_j^L = \omega_1^{c^L} + \omega_2^{c^L} u_j^{c^L}, \quad (6)$$

$$c_j^H = \omega_1^{c^H} + \omega_2^{c^H} u_j^{c^H}, \quad (7)$$

where $(u_j^\kappa, u_j^{c^L}, u_j^{c^H})$ are independent, uniformly distributed on the interval $[0, 1]$ random draws, for each j .

Assumptions on the parameter values are described in Appendix C.1.

Data Generating Processes. The simulated data are constructed under three different data generating processes (DGPs). For each DGP, I generate $M = 500$ independent samples. For each sample $m = 1, \dots, M$, I randomly draw students' geographic coordinates, income status, and schools' idiosyncratic cost terms. The DGPs considered are the following:

1. Bayesian Nash Equilibrium. This DGP considers a setting in which schools play a BN equilibrium; that is, each school takes into account all possible types of its opponents, and the probabilities of each type's occurrence.
2. Fully Cursed Equilibrium. This DGP considers the case in which schools play an FC equilibrium; that is, each school assumes that each of its opponents plays the average strategy, according to the distribution of types, regardless of their realized type.
3. Partial Cursed Equilibrium. This DGP considers a setting in which schools play a PC equilibrium; that is, each school is Bayesian with respect to its (geographically) closest opponent, and is fully cursed with respect to all other opponents.

Appendix C.2 provides extensive detail on the DGPs.

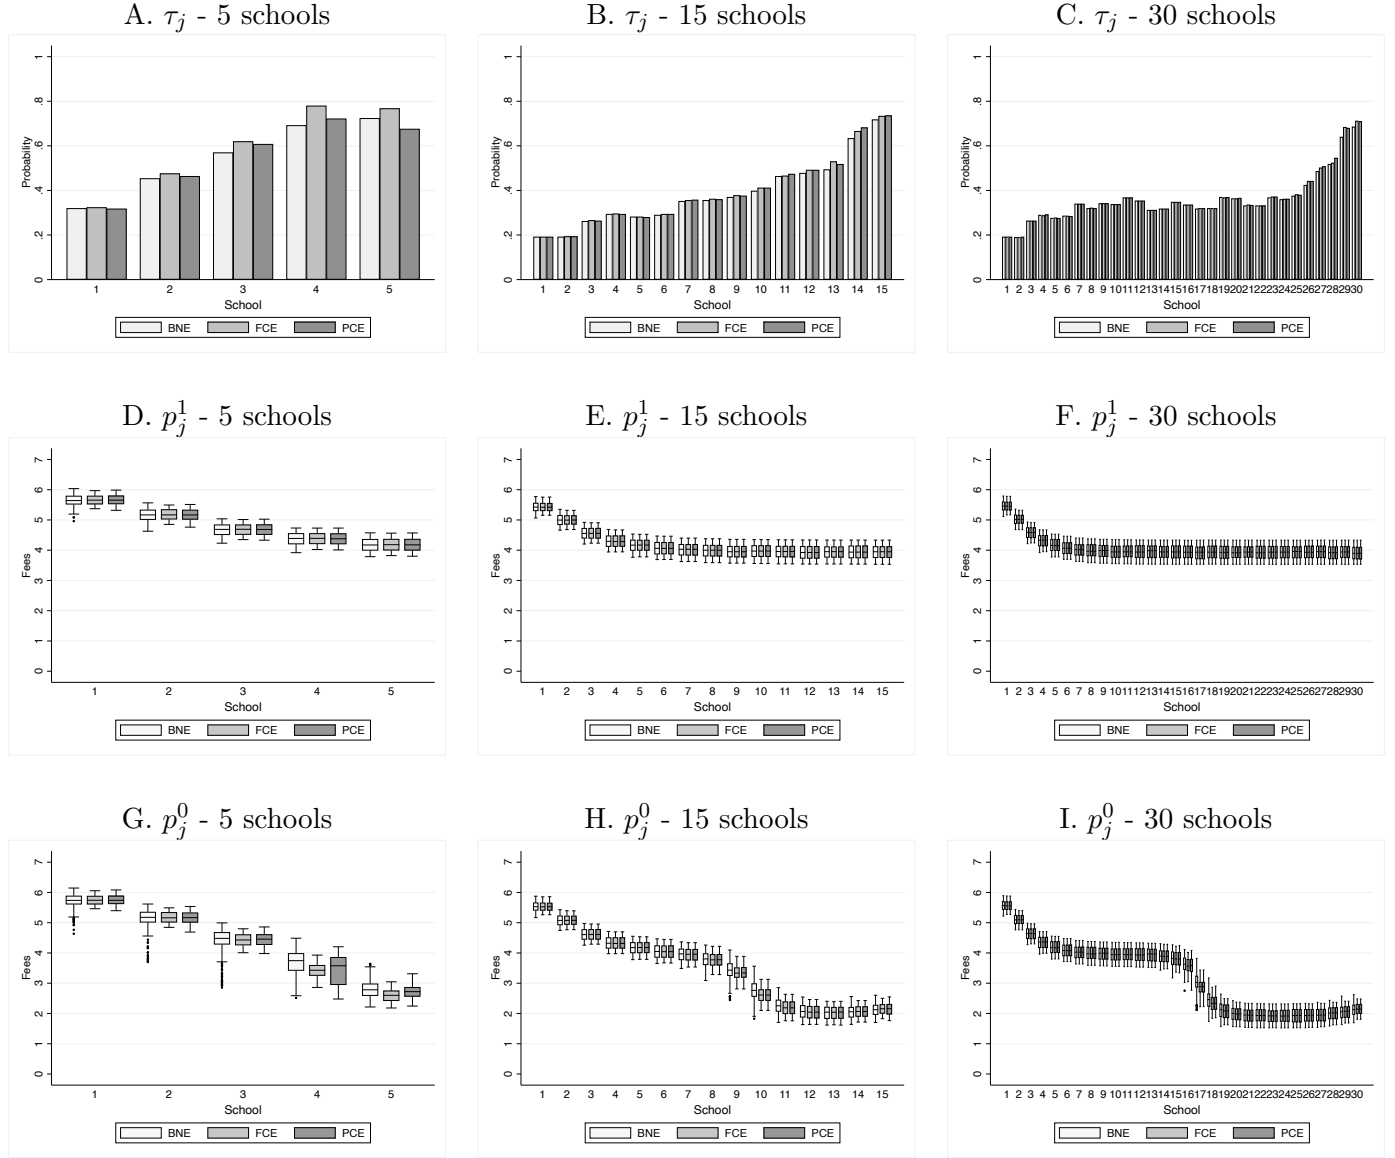
Simulation Results. Figure 2 presents a summary of the results from the Monte Carlo simulations. It displays the distributions of schools' equilibrium strategies under each equilibrium

concept, for markets with 5, 15, and 30 schools. Panels A–C show each school’s equilibrium program participation (τ_j) strategy distribution across different market sizes. Panels D–F do analogously for schools’ equilibrium counterfactual fees in the in-program regime, p_j^1 . Panels G–I do the same for schools’ equilibrium counterfactual fees in the not-in-program regime, p_j^0 . Simulations’ results for other market sizes and equilibrium effective prices are shown in Figures C.2–C.6 in Appendix C.3.

Results indicate that while differences across equilibrium concepts are nontrivial in very small markets, the predicted strategies converge quickly as market size grows. More specifically, in small markets, FCE tends to overpredict program participation relative to BNE, while PC equilibria differ from both. Notably, PCE strategies are not always a convex combination of BNE’s and FCE’s, which reflects the model’s high nonlinearity. Importantly, these differences become minimal as markets become large.

Equilibrium fee distributions display similar convergence. Relative to BNE, FCE and PCE yield more compressed counterfactual fee distributions in small markets, with fewer extreme values. By contrast, these differences largely vanish as J rises.

Figure 2: Schools' Equilibrium Strategy Distribution Under BNE, FCE, PCE



Notes: This figure displays schools' equilibrium strategies under BNE, FCE, and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. Panels A–C show each school's equilibrium program participation (τ_j) distribution across simulations for markets with 5, 15, and 30 schools, respectively. Panels D–F do analogously for schools' equilibrium counterfactual fees in the in-program regime, p_j^1 . Panels G–I do the same for schools' equilibrium counterfactual fees in the not-in-program regime, p_j^0 .

This evidence on the relationship between BNE, FCE, and PCE assumptions, together with the fact that urban education markets are large, usually consisting of tens and even hundreds of

schools, strengthens the argument for the adoption of fully cursed equilibrium in the analysis of Chile’s primary education markets. The convergence of equilibrium strategies under BNE, FCE, and PCE in sufficiently large markets suggests that FCE provides not only a computationally tractable, but also an empirically valid approximation for modeling schools’ behavior in sizable markets. This is particularly pertinent given the computational challenges associated with solving for Bayesian Nash equilibria in large-scale games where agents compete in both discrete and continuous strategies.

On the other hand, partial cursed equilibrium also presents an attractive alternative to BNE that is computationally lighter, and that preserves the Bayesian behavior for geographically close schools.¹⁴ However, in practice, PCE does not yield much different equilibrium strategies than FCE, nor are PCE strategies closer to BNE strategies. Moreover, FCE outperforms PCE in computational time required to find equilibria.

The analysis of the Monte Carlo simulations just presented through the lens of the theory of models of Bayesian games with cursed/coarse agents is instructive (Eyster and Rabin, 2005; Jehiel, 2005; Spiegel, 2011). The key element in these models is that cursed/coarse players do not assign informational value to others’ actions with respect to their private information—or, at most, they only partially connect others’ actions and private types. For instance, in a bilateral trade game with adverse selection (Eyster and Rabin, 2005; Spiegel, 2011), the cursed buyer that trades with a fully rational seller may choose a higher or a lower bid, relative to the benchmark BNE case, depending on the parameters defining the game. This ambiguity lies in two contradictory forces that emanate from the cursed buyer’s perceptual behavior. On one hand, the cursed buyer fails to take into account adverse selection, which results in an overvaluation of the traded object and in a higher bid than in the fully rational case. On the other hand, the buyer does not realize that a higher bid enhances the expected quality of the traded object, resulting in underbidding relative to the benchmark.

In the context of school competition in the Monte Carlo exercise, cursed schools overlook the connection between other schools’ actions and idiosyncratic costs. This misperception has important implications, that are observed in the simulations’ results. First, fully cursed schools have somewhat higher rates of program participation than fully rational schools, especially among schools with a high probability of joining the program in small markets. According to the theory, this result is explained by fully cursed schools overpredicting the competitiveness of others, since they cannot infer a competitor’s high participation private cost from the observation of a low participation rate.

Simulations also show a more compressed distribution of counterfactual prices for cursed

¹⁴Closeness between schools may alternatively be defined based on a school characteristic different from geographic distance, but that is also relevant for students’ demand, and therefore for supply side competition.

schools relative to the benchmark. Again, this phenomenon is a direct implication of the theory. Cursed schools underreact to opponents' high prices by choosing a price that is lower than the Bayesian Nash equilibrium price. Likewise, cursed schools do not realize that their opponents' low price is necessarily implied by a low idiosyncratic private cost, resulting in a cursed equilibrium price that is higher than the benchmark. Taken together, these two behavioral observations lead to a more compressed price distribution for cursed schools relative to fully rational schools.

Lastly, results from the simulations show a rapid convergence between bounded-rationality and full-rationality models' equilibrium strategies, suggesting that in sufficiently large games the distinction between fully and this form of bounded rationality does not considerably matter empirically. While interesting, this result is not entirely novel in the literature. In fact, the convergence result just presented is similar in spirit to the relation between Oblivious Equilibrium and Markov Perfect Equilibrium in dynamic models of imperfect competition (Weintraub et al., 2008, 2010).

Estimation Performance. I explore the performance of an estimation routine that (mistakenly) assumes data were generated from an FCE, while they were actually generated from a BNE. Specifically, I simulate data under the assumption of Bayesian Nash equilibrium, the standard rational equilibrium assumption in static Bayesian games, and examine whether an estimation routine that treats data as if they were generated from fully cursed schools can recover the true parameters. I conduct estimation analyses for varying market sizes.

For each Monte Carlo sample $m = 1, \dots, 50$, I generate 20 independent and identically distributed markets, each with the same primitives—i.e., number of students, number of schools, preference parameters, and cost structure parameters—but with independent draws of students' locations, income status, and schools' idiosyncratic cost terms. For each market, I solve for the BNE strategies, and then pool the resulting data across the 20 markets for estimation.

To focus on the supply side parameter estimation, I proceed as if the econometrician knew the true parameters governing the demand side of the model. This assumption leaves only the cost structure parameters to be estimated; that is, $\omega = (\omega_1^\kappa, \omega_2^\kappa, \omega_1^{c^L}, \omega_2^{c^L}, \omega_1^{c^H}, \omega_2^{c^H})$ in equations (5)–(7).

For each simulation sample $m = 1, \dots, 50$ and market size $J = 20, 25, 30$, I adopt a GMM-MPEC estimation approach, in which I minimize a GMM objective function subject to markets' equilibrium constraints. I work with market sizes $\{20, 25, 30\}$ because they better align with the market sizes in the actual data from Chile's primary education. The GMM objective function includes moments associated to schools' program participation, schools' effective prices, participant schools' fixed effects in students' indirect utility function (α_j^L and α_j^H), low- and high-income students' enrollment in participant schools, effective prices paid by low- and high-income stu-

dents, distance travelled by low- and high-income students, and low- and high-income students' enrollment share in participant schools.

Table 2 presents the estimation results. The first column displays the true parameters used in the data generating process. For each market size, the “mean” column reports the average of the estimated parameters across the 50 Monte Carlo samples, and the “std. dev.” column shows the standard deviation of these estimates' distribution. The results indicate that the GMM-MPEC estimator based on the fully cursed equilibrium assumption is able to recover the true parameters with little bias and reasonable precision, even when the data are generated from a Bayesian Nash equilibrium. Moreover, the estimates become more precise and closer to the true values as the market size increases, reflecting the earlier finding that equilibrium strategies under BNE and FCE become nearly indistinguishable in large markets.

Table 2: Monte Carlo Estimation Results

parameter	true value	20 schools		25 schools		30 schools	
		mean	std. dev.	mean	std. dev.	mean	std. dev.
ω_1^κ	-50	-61.56	22.38	-58.58	22.05	-58.23	22.19
ω_2^κ	150	177.42	80.26	163.29	81.44	159.42	80.69
ω_1^{cL}	0.4	0.16	0.29	0.27	0.31	0.30	0.34
ω_2^{cL}	0.8	0.68	0.15	0.74	0.15	0.75	0.17
ω_1^{cH}	0.2	0.23	0.04	0.24	0.06	0.25	0.09
ω_2^{cH}	0.8	0.81	0.02	0.82	0.03	0.82	0.05

Notes: This table reports the true parameter values, the mean of the GMM-MPEC estimates across 50 Monte Carlo samples, and the standard deviation of the estimates' distribution, for each market size (20, 25, and 30 schools). Estimation is performed under the Fully Cursed equilibrium assumption using data generated from a Bayesian Nash equilibrium. For each Monte Carlo simulation and market size, 20 independently drawn markets are considered.

The results from Table 2 provide strong empirical support for the use of Fully Cursed equilibrium as a computationally tractable and reliable approximation in the estimation of large discrete-continuous Bayesian games, such as the one studied in this paper in the context of primary education markets.

4 Data and Educational Markets

I combine various administrative data sets for Chilean schools and students for the year 2013. I use the registry of all operating schools, in which I observe schools' administrative category, fees (only for private-voucher schools), subsidies, participation in the targeted voucher program, geocoded address, and other characteristics such as religious orientation, single/double shift, and

chain membership. I adjoin to these school-level data tax records for school providers, that allow me to identify the probable for-/nonprofit status of each private-voucher school.¹⁵ I also utilize the registry of all students attending primary education in the country. I observe students’ grade and school of attendance, eligibility for the targeted program, (non-geocoded) residential address, and other demographic characteristics. I additionally use records on students’ performance in mandatory standardized tests taken by all 4th graders in the country. The standardized exams data include a questionnaire to parents that provides information on demographic characteristics such as parents’ level of education, household income, and house amenities (e.g., computer and internet availability), which I also use.¹⁶

In addition, to construct instruments, I use pre-reform administrative data from 2007, one year prior to the introduction of the program, including the censuses of schools and students. I also utilize the national household survey (CASEN) for the year 2013 to aid in the construction of instrumental variables.

The administrative records provide geocoded addresses for schools, but only name addresses for students’ residence. Furthermore, not all students have valid name addresses. To address this data limitation, I implement a two-step procedure to recover students’ geocoded addresses, and later the corresponding distance from each student’s home to every school in the market. First, I use a combination of GIS tools to obtain geographic coordinates for students with valid name addresses. Second, for the rest of the students in the sample, I perform a random forest prediction of their residence coordinates based on a large vector of observable demographics combined with the latitude-longitude data retrieved from students with valid addresses. I end up with 66% of students having either correctly geocoded addresses or precisely predicted coordinates. Although coverage is not complete, my two-step procedure represents a significant improvement over the 50% valid geocoded addresses reported in Neilson (2025) or the imputed coordinates based on the municipality centroid in Gazmuri (2024)—both studies examine samples of primary education students very similar to this paper’s. In my empirical application, I impute all missing distance observations with a value of zero, and include a dummy indicator for non-missing in the original variable.¹⁷

I also collect data on private-non-voucher schools’ fees. This information is not included in the administrative data from the Ministry of Education. I collect fee amounts manually from

¹⁵The identification of the school’s for-/nonprofit status based on tax records is not perfect. Thus, I cautiously interpret the for-/nonprofit status from the data only as a likely indicator of the profit-seeking motive of the school.

¹⁶See Appendix A for a more detailed description of the administrative data sets I use in this paper.

¹⁷More specifically, the variable distance from home to school is transformed in the following way,

$$d' = d \times \mathbf{1} [d = \text{non-missing}].$$

For completeness, I include both d' and $\mathbf{1} [d = \text{non-missing}]$ variables in the choice model to be estimated.

schools' websites and through telephone inquiries. I successfully retrieve fee values for all private-non-voucher schools in the country.

An intermediate step in constructing the final sample is to define educational markets. This task is challenging in the context of urban centers in Chile, as is the case in other Latin American countries (Allende, 2019; Dinerstein et al., 2023; Neilson, 2025). Unrestricted choice combined with a large supply of public and private schools makes it difficult to draw clear boundaries for school markets. As an illustration, a student who lives close to the border of a city may very well choose to attend a school located in a different nearby city.

To define educational markets, I develop an algorithm inspired by Neilson (2025), which constructs markets based on student enrollment data. The algorithm proceeds as follows.

1. For a particular municipality, which I call *origination* municipality, I consider the school choices of all students residing in that municipality.
2. From those school choices, I identify all municipalities to which students travel to attend school. I call these *destination* municipalities. Note that the origination municipality can simultaneously be a destination municipality—in practice, that is always the case.
3. I rank these destination municipalities according to popularity; that is, based on the share of students who choose a school in each destination municipality.
4. I join all municipalities with a popularity rate of 5% or higher, to form the first iteration of the market associated with the origination municipality.
5. I repeat steps 1–4, but consider this first iteration of the market instead of the origination municipality in step 1. I stop when no new municipality is added to the group of destination municipalities that define a market.

I apply this algorithm for all municipalities in my sample.

I select large educational markets for my final sample. This selection criterion is necessary to bring my model, which features a supply-side game with many players, to the data. Specifically, I include markets with 10,000 or more primary education (grades 1st–8th) students. In addition, and only to work with markets that are more similar in size to each other, I leave the capital city, Santiago, out of the final sample. Including Santiago is straightforward and does not change the policy lessons drawn from the empirical analyses. The final sample consists of 27 markets, which enroll ~44% of the student population and include ~41% of all schools.¹⁸ Lastly, I keep only 4th graders in the estimation sample, since data on standardized test scores are only available for these students.¹⁹

¹⁸Figure B.1 in Appendix B presents an example of an educational market created with the geocoded data.

¹⁹See Neilson (2025) for a similar sample selection criterion.

Some market-level descriptive statistics summarized in Table 3 are worth mentioning. The average (median) market has 2,876 (2,695) 4th grade students, of which 54% (57%) are eligible for the targeted voucher. The largest market has 7,116 4th grade students. There are 77 (76) schools in the average (median) market, 37 (38) of which are public, 35 (31) are private-voucher, and 4 (3) are private-non-voucher. The largest market has 105 private-voucher schools. Among all private-voucher schools in the typical market, about two-thirds participate in the targeted voucher program. These figures underscore the wide reach of the targeted voucher among students and schools, as well as the large size of the supply-side game played by private-voucher schools, which makes my model suitable for the context.

Table 3: Educational Markets' Characterization

	mean	std. dev.	min	median	max
students	2,876	1,493	1,025	2,695	7,116
% low-income students	54	12	17	57	71
schools	77	36	27	76	151
public schools	37	20	8	38	76
private-voucher schools	35	23	10	31	105
private-non-voucher schools	4	4	0	3	14
% private-voucher schools in targeted program	66	18	28	67	89

Notes: Summary statistics for all 27 geographic educational markets included in the empirical analysis.

Student enrollment in the different administrative categories of schools follows some interesting patterns. Table 4 summarizes enrollment, top-up fees, and test scores for schools in my final sample. The vast majority (95%) of students attend a subsidized school, either public or private-voucher. Furthermore, low-income students' enrollment decisions are clearly skewed toward schools that charge them no fees: half of these students attend a public school, and about 39% attend a private-voucher school that is part of the targeted voucher program. This pattern reflects the marked financial constraints experienced by low-income families when choosing schools, which originally motivated the introduction of the targeted voucher. High-income students' choices are more dispersed among the different categories of schools, with less than a third attending a public zero-fee school and 11% attending a high-fee private-non-voucher school.

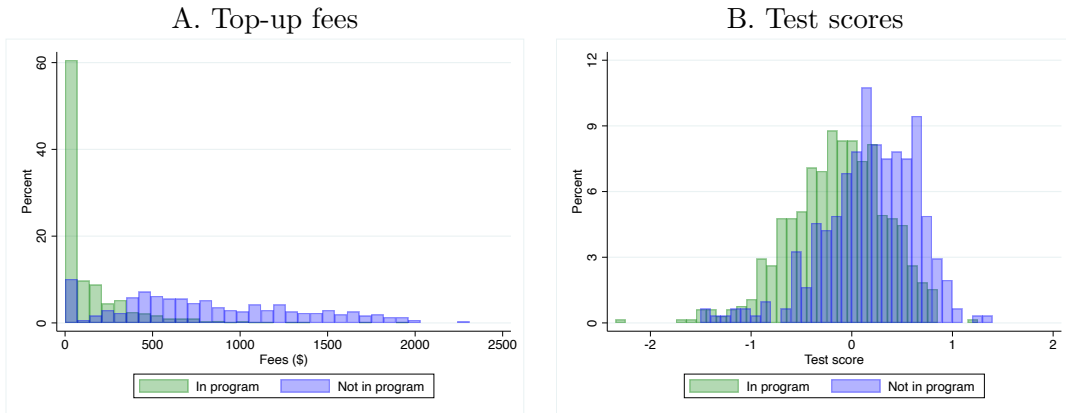
Table 4: School Characteristics, by Schools' Administrative Category

<i>category:</i> <i>in targeted voucher program:</i>	public	private-voucher		private-non-voucher
	yes	yes	no	no
enrollment (%)	41	35	19	6
enrollment - low-income (%)	50	39	11	1
enrollment - high-income (%)	30	30	28	11
top-up fees (\$)	0	150	869	6,687
test scores (s.d.)	-0.26	-0.01	0.28	0.78

Notes: Characteristics of primary schools in the analysis sample, depending on whether the school is public, private-voucher, or private-non-voucher, and on whether it participates in the targeted voucher program. Fees are in real 2013 prices, and were converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD). Test scores are the school average of math and verbal exams and are standardized to have mean zero and standard deviation one at the student level.

Private-voucher school selection into the targeted voucher program is also notable. It is mostly low-fee, low-test score schools that decide to join the program, as shown in Table 4 and Figure 3. This pattern reflects the (at least partial) profit-seeking motive of private-voucher schools, as discussed in Section 2. It also shows that many high-test score schools chose to stay out of the program, which raises doubts about the program's ability to make high-achieving schools more easily available to low-income students.²⁰

Figure 3: Private-Voucher Schools' Fees and Test Score Distributions



Notes: Fee levels are in real 2013 prices, and were converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD). Test scores are the school average of math and verbal exams among students, and are in standard deviation units.

²⁰A similar school sorting pattern was recently observed in the Louisiana Scholarship program, the Louisiana state-funded targeted voucher program studied in Abdulkadiroglu et al. (2018).

5 Estimation and Identification

I estimate the model’s parameters sequentially. First, in three steps, I obtain the school popularity and taste heterogeneity parameters in the preferences model, a measure of school quality in auxiliary test scores equations, and then link these two sets of parameters to finalize preferences estimation. Then, given demand parameters, I estimate the parameters that enter schools’ fixed participation and marginal cost functions.

5.1 Demand

I proceed in three steps. First, I estimate school popularity fixed effects and the parameters of heterogeneity in students’ preferences. Then, given that school quality is not directly observed in the data, I obtain a school quality measure by estimating test scores equations using a selection-on-observables approach, similar to Allende (2019), Allende et al. (2019), and Neilson (2025). Finally, I link school quality and other characteristics to school popularity estimates to finalize the estimation of preferences for schools.

School Popularity and Preference Heterogeneity. I use Maximum Likelihood to estimate preferences for proximity, taste heterogeneity, and mean utilities or school popularity. Heterogeneity in preferences is captured by a set of students’ observable demographic characteristics, which in practice are mother’s educational level dummies. Mean utilities vary by the student’s socioeconomic status (low/high-income), and absorb the remaining preference components from the indirect utility function,

$$\delta_j^\zeta = \beta_1^\zeta p_j^\zeta + \beta_3^{\zeta'} X_j^\beta + \beta_4^\zeta q_j + \beta_5 \tau_j^\zeta + \xi_j^\zeta.$$

The corresponding log-likelihood function is:

$$LL(\beta) = \sum_i \sum_j e_{ij} \ln \left(\frac{\exp \left((\beta_{1i}^\zeta - \beta_1^\zeta) p_j^\zeta + \beta_2^\zeta d_{ij} + (\beta_{3i}^\zeta - \beta_3^\zeta) X_j^\beta + (\beta_{5i}^\zeta - \beta_5^\zeta) \tau_j^\zeta + \delta_j^\zeta \right)}{\sum_k \exp \left((\beta_{1i}^\zeta - \beta_1^\zeta) p_k^\zeta + \beta_2^\zeta d_{ik} + (\beta_{3i}^\zeta - \beta_3^\zeta) X_k^\beta + (\beta_{5i}^\zeta - \beta_5^\zeta) \tau_k^\zeta + \delta_k^\zeta \right)} \right),$$

where e_{ij} is the choice indicator, i.e., $e_{ij} = 1$ indicates that student i attends school j .

School Quality. To estimate the measure of school quality, I follow Allende (2019), Allende et al. (2019), and Neilson (2025), and estimate the parameters of an auxiliary test score equation using an offline selection-on-observables linear model. Note that school value-added measures, as defined by the value-added literature (Rockoff, 2004; Rivkin et al., 2005; Kane and Staiger, 2008; Chetty et al., 2014a,b; Bau, 2022), cannot be estimated in this context, because prior test scores

are not observed. Instead, my school quality estimates measure a school’s accumulated quality contribution for grades 1st–4th.

In practice, I estimate by OLS the following regression,

$$Y_{ij} = W_i\gamma + \sum_j q_j e_{ij} + \epsilon_i,$$

where I include a large set of observables in W_i to capture as much variation as possible, and minimize any inconsistency in the estimates. The covariates I include are gender, low-income status, grade repetition, mother’s education, father’s education, household income, having attended pre-K, having attended kindergarten, computer availability at home, internet availability at home, indigenous status, number of books at home, and class attendance. I use the student-level average of 4th grade national standardized scores in verbal and math exams as the outcome variable. This set of observables is in line with those in Allende (2019), Allende et al. (2019), and Neilson (2025). Further, Allende et al. (2019) show that school quality estimates based on this procedure and a similar vector of observables yield comparable results to those obtained when adding prior test scores as controls. They also provide additional experimental evidence supporting the use of this strategy. I estimate the test score equations market by market. I obtain $\{\hat{q}_j\}_{j=1}^J$ as school quality estimates.

Under the selection-on-observables assumption, $E[D_{ij}\epsilon_i] = 0$, the set $\{\hat{q}_j\}_{j=1}^J$ contains consistent estimates of $\{q_j\}_{j=1}^J$; however, they are subject to estimation noise. Depending on the school size—i.e., number of students—the estimation noise may be considerable. Shrinkage estimators often help reduce the estimation noise.

Following Walters (2024), I use the following parametric normal/normal model to shrink the school quality estimates. Let $\{s_j\}_{j=1}^J$ be the collection of standard errors for $\{\hat{q}_j\}_{j=1}^J$. The first level of the hierarchical model is described by,

$$\hat{q}_j \mid q_j, s_j \sim N(q_j, s_j^2).$$

The second level of the hierarchy describes the cross-school distribution of quality,

$$q_j \mid s_j \sim N(\mu_q, \sigma_q^2),$$

with (μ_q, σ_q^2) the hyperparameters that summarize the quality distribution. I utilize common

estimators for these hyperparameters:

$$\begin{aligned}\hat{\mu}_q &= \frac{1}{J} \sum_j \hat{q}_j, \\ \hat{\sigma}_q^2 &= \frac{1}{J} \sum_j [(\hat{q}_j - \hat{\mu}_q)^2 - s_j^2].\end{aligned}$$

Subtracting s_j^2 is a bias correction that accounts for excess variance in the \hat{q}_j due to sampling error.

The posterior mean for q_j given (\hat{q}_j, s_j) is:

$$q_j^* \equiv E[q_j | \hat{q}_j, s_j] = \left(\frac{\sigma_q^2}{\sigma_q^2 + s_j^2} \right) \hat{q}_j + \left(\frac{s_j^2}{\sigma_q^2 + s_j^2} \right) \mu_q,$$

with its empirical counterpart,

$$\hat{q}_j^* = \left(\frac{\sigma_q^2}{\sigma_q^2 + s_j^2} \right) \hat{q}_j + \left(\frac{s_j^2}{\sigma_q^2 + s_j^2} \right) \hat{\mu}_q.$$

The set $\{\hat{q}_j^*\}_{j=1}^J$ contains the Empirical Bayes shrunk estimates of school quality that I use in my empirical analysis.

Linking Mean Utilities to School Quality and Other Characteristics. I use the estimated $\hat{\delta}_j^\zeta$ terms from the first step, as well as the shrinkage-estimated school quality \hat{q}_j^* from the test score regressions, to estimate the remaining mean preference parameters in a linear regression of the form:

$$\hat{\delta}_j^\zeta = \beta_1^\zeta p_j^\zeta + \beta_3^{\zeta'} X_j^\beta + \beta_4^\zeta \hat{q}_j^* + \beta_5 \tau_j^\zeta + \xi_j^\zeta. \quad (8)$$

As is standard in demand models, I assume that X_j^β is uncorrelated with ξ_j^ζ . However, I allow $p_j^\zeta = \tau_j(1 - D_i)p_j^1 + (1 - \tau_j)p_j^0$ and τ_j to be endogenous—i.e., correlated with ξ_j^ζ —and estimate equation (8) by 2SLS using the instruments described below. To account for the estimation noise in the dependent variable, $\hat{\delta}_j^\zeta$, I weight each observation by the inverse of the maximum likelihood-estimated variance of the school fixed effect.

Identification and Instruments. The term ξ_j^ζ represents group- ζ families' preferences for school j 's unobserved characteristics, such as the principal's charisma or managerial style, educational curriculum, etc., which may be correlated with the school's choice of top-up fees and program participation. A stylized example is a school with a mission to serve vulnerable families.

Such a school may be willing to develop a culture and an educational environment that are particularly welcoming to low-income students, which may intensify the school’s propensity to join the targeted program and to set low fees. Thus, consistent estimation of the linear parameters in equation (8) necessitates instrumental variables.

My choice of instruments for program participation and top-up fees follows related studies (Allende, 2019; Allende et al., 2019), and in particular Neilson (2025). To instrument program participation, I use two instrumental variables. The first is the share of eligible students in the school’s neighborhood, defined by the *zona censal* where the school is located.²¹ To avoid endogeneity issues related to reform-induced migration, I use the pre-reform share of eligible students, measured one year prior to the introduction of the policy. As discussed in Section 2, a higher share of low-income students increases the relative ex-ante profit associated with joining the program, especially for zero- and low-fee schools. The exogeneity assumption is that the pre-reform demographic composition of the school’s neighborhood is uncorrelated with families’ preferences for non-academic quality, ξ_j^ζ . In other words, families’ valuation of characteristics such as the principal’s managerial style is invariant to changes in the socioeconomic composition of the neighborhood.

A second instrument for a school’s program participation decision is the indicator for positive ex-ante net profit of joining the program. I construct this variable using the school’s pre-reform enrollment share of vulnerable students and its pre-reform fees, combined with the post-reform value of the targeted voucher, i.e., $\mathbb{1}\{\#\text{ low-income students} \times (\text{targeted voucher} - \text{top-up fees}) > 0\}$. This instrument is motivated by the discussion in Section 2, where I show that a simple ex-ante profit indicator is highly correlated with schools’ participation decisions. The validity of the instrument rests on the assumption that this ex-ante profit measure, which is based on pre-reform equilibrium strategies, is exogenous to post-reform families’ preferences for unobserved school characteristics, ξ_j^ζ .

I employ two instruments to identify families’ sensitivity to schools’ fees. First, in the spirit of classical BLP instruments (Berry et al., 1995; Nevo, 2001; Berry and Haile, 2014), I construct a distance-weighted proportion of private-voucher schools in the school’s neighborhood, with weights being the inverse of the distance to the school. This variable captures how strong the local competition faced by the school is, where the intuition is that a higher share of private schools translates into stronger competitive pressure, leading to lower fees.²² The exclusion restriction

²¹Similar to a block group in the US census, a *zona censal* in Chile’s census is the second smallest geographic unit, and is composed of a group of *manzanas*.

²²Think of a private-voucher school setting its fees facing only public schools as local competitors; this school is, in practice, a monopoly in the local private sector, with incentives to set the corresponding monopolistic fees. Now, suppose this private-voucher school is one of many such schools in the neighborhood; the school necessarily sets fees lower than the monopolistic level.

holds under the assumption that, conditional on the school characteristics included in X_j^β , the spatial distribution of competitors and their density in the local market area do not directly affect families' preferences for unobserved school attributes such as the principal's charisma or the school's culture.

In addition, I use an instrument based on wage compensating differentials for top-up fees. Following Neilson (2025), I use the 2013 national household survey to run a linear regression of log wages for highly educated workers on a quadratic polynomial in age and on industry, occupation, and municipality fixed effects. Differences in the estimated municipality fixed effects capture relative salary differences paid across local markets—i.e., wage compensating differentials. As Neilson (2025) notes, about 80% of school costs are related to teachers' salaries. Therefore, the municipality fixed effects serve as a measure of labor costs. These labor costs are unlikely to be correlated with parents' preferences for the school's unobserved characteristics.

I include an additional variable in the set of instruments: the average pre-reform top-up fees charged by competitor schools in the school's neighborhood. This instrument provides information on other schools' cost structure, and therefore on (post-reform) competitors' strategies—i.e., fees and targeted voucher take-up—which in turn determine the school's own strategies, via oligopolistic competition. The validity of this instrument relies on the assumption that other schools' pre-reform pricing decisions are not correlated with the school's post-reform unobserved characteristics, ξ_j^ζ , after controlling for the school's own observable characteristics.

An important assumption I make is to treat the school quality measure, q_j , as an exogenous characteristic. Both Allende et al. (2019) and Neilson (2025) allow schools to be strategic in this dimension, and instrument quality in demand models and contexts similar to those in this paper. Moreover, Cañedo Riedel and Sánchez (2021) and Neilson (2025) present supporting evidence showing that both participant and non-participant schools increased a measure of quality and invested in educational inputs associated with improving the quality of education in response to the introduction of the reform. A full model that incorporates endogenous quality necessitates additional cross-sectional (and possibly time-varying) variation in school quality and a model for the production function of educational quality. My supply-side model is already rich and extensive, and proves successful in capturing various market forces driving schools' responses to policies and market equilibria. Therefore, extending the set of school strategies with a quality dimension is beyond the scope of this paper. By maintaining the quality of schools fixed, I expect my counterfactual policy results to be biased toward attenuation of the real effects. As an illustration, in the extreme case of simulating a full expansion of the targeted voucher program, where all schools are mandated to join, my model would fail to account for the quality improvements associated with the extra funds injected by the program, and would likely predict lower student welfare than if I let schools adjust quality. At the other extreme, when simulating scenarios toward lower

or no targeted voucher, my model would again fail to account for schools' downward quality adjustments, predicting a smaller reduction in welfare than if quality was treated endogenously.

5.2 Supply

I parameterize schools' cost structure as follows,

$$\begin{aligned} c_j^H &= X_j^\omega \omega_{cH} + u_j^{cH} \\ c_j^L &= X_j^\omega \omega_{cL} + u_j^{cL} \\ \kappa_j &= Z_j^\omega \omega_\kappa + u_j^\kappa, \end{aligned}$$

where X_j^ω and Z_j^ω are vectors of observable (to the econometrician and all players) variables affecting marginal costs and participation fixed costs, respectively. The terms ω_{cH} , ω_{cL} , and ω_κ are parameters to be estimated. The idiosyncratic error terms, u_j^{cH} , u_j^{cL} , and u_j^κ , are independent and normally distributed with mean zero and variance terms σ_{cH}^2 , σ_{cL}^2 , and σ_κ^2 , respectively. The variance terms are also parameters to be estimated.

Conditional on observables X_j^ω and Z_j^ω , all cost elements are statistically independent, both within and across schools, i.e., $c_j^H \perp\!\!\!\perp c_{j'}^L \perp\!\!\!\perp \kappa_{j''} \mid X^\omega, Z^\omega$ for any j, j', j'' , where $\perp\!\!\!\perp$ denotes statistical independence.

I combine demand estimates and schools' optimality conditions (2)–(4) to recover the parameters governing schools' cost structure. Let $\bar{s}_j(\bar{s}_{-j}) \equiv \int_{\mathcal{S}_{t_j}} s_j(\bar{s}_{-j}; t_j) dF_t(t_j)$ be school j 's fully cursed strategy, where \mathcal{S}_{t_j} is the space of school j 's types.²³ Thus,

$$\begin{aligned} \bar{p}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} \mid X^\omega, Z^\omega) &= \Phi \left(\frac{\left(X_j^\omega \omega_{cH} - v^u - \frac{\sum_i (1-D_i) P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^1}} - \psi \right)}{\sigma_{cH}} \right) \\ &\quad \left[X_j^\omega \omega_{c^{1,H}} - v^u - \frac{\sum_i (1-D_i) P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^1}} - \psi \right. \\ &\quad \left. + \sigma_{cH} \lambda \left(\frac{-X_j^\omega \omega_{cH} + v^u + \frac{\sum_i (1-D_i) P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^1}} + \psi \right)}{\sigma_{cH}} \right), \quad (9) \end{aligned}$$

²³Recall that $\bar{s}_j(\bar{s}_{-j}) = (\bar{\tau}_j(\bar{s}_{-j}), \bar{p}_j^1(\bar{s}_{-j}), \bar{p}_j^0(\bar{s}_{-j}))$. Also, note that $\bar{s}_j^1(\bar{s}_{-j}) = (\bar{\tau}_j(\bar{s}_{-j}) = 1, \bar{p}_j^1(\bar{s}_{-j}), \bar{p}_j^0(\bar{s}_{-j}))$ and $\bar{s}_j^0(\bar{s}_{-j}) = (\bar{\tau}_j(\bar{s}_{-j}) = 0, \bar{p}_j^1(\bar{s}_{-j}), \bar{p}_j^0(\bar{s}_{-j}))$.

where $\lambda(\nu) = \frac{\phi(\nu)}{1-\Phi(\nu)}$ is the inverse Mills ratio, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal density and probability distribution, respectively. Note that,

$$P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) = \frac{e^{V_{ij}(\bar{s}_j^1)}}{e^{V_{ij}(\bar{s}_j^1)} + \sum_{k \neq j} e^{V_{ik}(\bar{s}_k)}},$$

where, $V_{ij}(\bar{s}_j^1) = \beta_{1i}(1-D_i)\bar{p}_j^1 + \beta_{2i}^\zeta d_{ij} + \beta_{3i}^\zeta X_j^\beta + \beta_4^\zeta q_j + \beta_{5i} + \xi_j^\zeta$, and $V_{ik}(\bar{s}_k) = \beta_{1i}(\bar{\tau}_k(1-D_i)\bar{p}_k^1 + (1-\bar{\tau}_k)\bar{p}_k^0) + \beta_{2i}^\zeta d_{ik} + \beta_{3i}^\zeta X_k^\beta + \beta_4^\zeta q_k + \beta_{5i}\bar{\tau}_k + \xi_k^\zeta$, for $k \neq j$. Also,

$$\frac{\partial P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^1} = \beta_{1i}(1-D_i)P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) (1 - P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})).$$

Similarly,

$$\begin{aligned} \bar{p}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} | X^\omega, Z^\omega) = & \Phi \left(\frac{X_j^\omega \omega_{cH} \frac{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} + X_j^\omega \omega_{cL} \frac{\sum_i D_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} - v^u}{\sigma_0} \right. \\ & \left. - \frac{\sum_i P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} - \psi \right) \left[X_j^\omega \omega_{cH} \frac{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} \right. \\ & \left. + X_j^\omega \omega_{cL} \frac{\sum_i D_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} - v^u - \frac{\sum_i P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} - \psi \right. \\ & \left. + \sigma_0 \lambda \left(\frac{-X_j^\omega \omega_{cH} \frac{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} - X_j^\omega \omega_{cL} \frac{\sum_i D_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}}{\sigma_0} \right. \right. \\ & \left. \left. + v^u + \frac{\sum_i P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} + \psi \right) \right], \end{aligned} \quad (10)$$

where $\sigma_0 = \left[\left(\frac{\sum_i (1-D_i) \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} \right)^2 \sigma_{cH}^2 + \left(\frac{\sum_i D_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}}{\sum_i \frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0}} \right)^2 \sigma_{cL}^2 \right]^{\frac{1}{2}}$, and

$$P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) = \frac{e^{V_{ij}(\bar{s}_j^0)}}{e^{V_{ij}(\bar{s}_j^0)} + \sum_{k \neq j} e^{V_{ik}(\bar{s}_k)}},$$

with $V_{ij}(\bar{s}_j^0) = \beta_{1i} \bar{p}_j^0 + \beta_{2i}^{\zeta} d_{ij} + \beta_{3i}^{\beta} X_j^{\beta} + \beta_{4i}^{\zeta} q_j + \xi_j^{\zeta}$. Likewise,

$$\frac{\partial P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^0} = \beta_{1i} P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) (1 - P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})).$$

Finally,

$$\bar{\tau}_j(\bar{s}_j(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) = \Phi \left(\frac{\left(\mathcal{U}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) - Z_j^{\omega} \omega_{\kappa} \right) - \mathcal{U}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega})}{\sigma_{\tau}} \right), \quad (11)$$

where,

$$\begin{aligned} \mathcal{U}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) &= (\bar{p}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) + v^u - X_j^{\omega} \omega_{cH}) \sum_i (1 - D_i) P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) \\ &\quad + (v^t + v^u - X_j^{\omega} \omega_{cL}) \sum_i D_i P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) + \psi P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}), \\ \mathcal{U}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) &= (\bar{p}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) + v^u - X_j^{\omega} \omega_{cH}) \sum_i (1 - D_i) P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) \\ &\quad + (\bar{p}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} | X^{\omega}, Z^{\omega}) + v^u - X_j^{\omega} \omega_{cL}) \sum_i D_i P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) \\ &\quad + \psi P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}), \end{aligned}$$

and,

$$\begin{aligned} \sigma_{\tau} &= \left[\left(\sum_i (1 - D_i) P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) \right)^2 \sigma_{cH}^2 + \left(\sum_i D_i P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j}) \right)^2 \sigma_{cL}^2 + \right. \\ &\quad \left. \left(\sum_i (1 - D_i) P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) \right)^2 \sigma_{cH}^2 + \left(\sum_i D_i P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j}) \right)^2 \sigma_{cL}^2 + \sigma_{\kappa}^2 \right]^{\frac{1}{2}}. \end{aligned}$$

Notice that, since X_j^{ω} and Z_j^{ω} are observed by everyone in the game, school j 's cost structure is private information only because the stochastic part of the cost terms, $(u_j^{cH}, u_j^{cL}, u_j^{\kappa})$, is unob-

served by j 's competitors. That is, school j 's type is given by the realization of its idiosyncratic error terms.

Identification proceeds as follows. The term $\bar{p}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} \mid X^\omega, Z^\omega)$ is the conditional mean of a normal censored (from below at zero) variable. Normality and variation from the observables ensure the identification of ω_{cH} and σ_{cH} . Furthermore, the model is overidentified, because theory imposes the coefficient accompanying the markup term $\frac{\sum_i(1-D_i)P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\sum_i(1-D_i)\frac{\partial P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})}{\partial p_j^1}}$ to be equal to one. A similar argument identifies ω_{cL} and σ_{cL} . From the probit assumption for $\bar{\tau}_j(\bar{s}_j(\bar{s}_{-j}), \bar{s}_{-j} \mid X^\omega, Z^\omega)$, variation in Z_j^ω identifies $\frac{\omega_\kappa}{\sigma_\tau}$. The coefficients accompanying the school's payoffs, $\mathcal{U}_j^1(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j} \mid X^\omega, Z^\omega)$ and $\mathcal{U}_j^0(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j} \mid X^\omega, Z^\omega)$, are constrained to be one by the theory, which allows the identification of ω_κ and σ_κ , separately. The relative importance of enrollment in schools' payoffs, ψ , is identified from variation in $P_{ij}(\bar{s}_j^1(\bar{s}_{-j}), \bar{s}_{-j})$ and $P_{ij}(\bar{s}_j^0(\bar{s}_{-j}), \bar{s}_{-j})$.

Exogenous variation from the use of instruments in (X^ω, Z^ω) aids in the identification of the parameters. I utilize the same instrumental variables as those used to identify the demand parameters. In particular, I include the local share of private competitors, labor costs, and pre-reform competitors' fees in the vector of covariates determining marginal costs, while the pre-reform share of low-income students, an indicator for ex-ante positive profits from joining program, and the pre-reform competitors' fees are the instruments included in Z^ω .

Both cross-market and within-market variation are exploited to separately pin down the parameters of the supply side. Intuitively, cross-market variation (differences in market size, the share of low-income students, local wage levels, pre-reform competitor fees, and the density of private competitors) generates exogenous shifts in schools' cost shifters and in the composition of demand. These shifts move schools' observed participation, fees, and enrollments in predictable ways and therefore identify the parameters that govern marginal costs, the distributional variances, and the fixed cost of participation. In addition, the vector of cost determinants, (X^ω, Z^ω) , contains market fixed effects to capture unobserved market heterogeneity. On the other hand, within-market variation (the cross-sectional dispersion of fees, participation choices, and enrollment shares across schools facing the same market-level demand) identifies markups, price elasticities, and the role of the enrollment weight (ψ) in schools' objective function: conditional on the same market demand, differences in schools' observable covariates and instruments reveal how strategic interactions translate into observed prices and take-up decisions.

I utilize a Generalized Method of Moments (GMM) routine to estimate the parameters on the supply side. I combine GMM with a Mathematical Programming with Equilibrium Constraints (MPEC; Dubé et al., 2012; Su and Judd, 2012) approach that minimizes the objective function of moment conditions subject to markets' equilibria, i.e., the fixed point of schools' best responses in

each market. The estimation combines market-level moments that exploit cross-market exogenous shifts with school-level moments that exploit within-market cross-sectional variation. I choose the following sample moments to construct the GMM objective function (where the subscript m indicates market m),

$$g_\tau = \frac{1}{\sum_{m=1}^M J_m} \sum_{m=1}^M \sum_{j=1}^{J_m} Z_j^{\omega'} (\tau_{j,m} - \bar{\tau}_{j,m}(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega)), \quad (12)$$

$$g_{p^L} = \frac{1}{\sum_{m=1}^M J_m} \sum_{m=1}^M \sum_{j=1}^{J_m} X_j^{\omega'} (p_{j,m}^L - \bar{p}_{j,m}^L(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega)), \quad (13)$$

$$g_{p^H} = \frac{1}{\sum_{m=1}^M J_m} \sum_{m=1}^M \sum_{j=1}^{J_m} X_j^{\omega'} (p_{j,m}^H - \bar{p}_{j,m}^H(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega)), \quad (14)$$

$$g_{\tau,q} = \frac{1}{\sum_{m=1}^M J_m} \sum_{m=1}^M \sum_{j=1}^{J_m} \left(\frac{J_m}{\sum_{j=1}^{J_m} \tau_{j,m}} \tau_{j,m} q_{j,m} - \frac{J_m}{\sum_{j=1}^{J_m} \bar{\tau}_{j,m}} \bar{\tau}_{j,m}(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega) q_{j,m} \right), \quad (15)$$

$$g_{p^L,q}^m = \frac{1}{M} \sum_{m=1}^M \left(\frac{1}{J_m} \sum_{j=1}^{J_m} p_{j,m}^L q_{j,m} - \bar{p}_{j,m}^L(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega) q_{j,m} \right), \quad (16)$$

$$g_{p^H,q}^m = \frac{1}{M} \sum_{m=1}^M \left(\frac{1}{J_m} \sum_{j=1}^{J_m} p_{j,m}^H q_{j,m} - \bar{p}_{j,m}^H(\bar{s}_{j,m}(\bar{s}_{-j,m}), \bar{s}_{-j,m} | X^\omega, Z^\omega) q_{j,m} \right), \quad (17)$$

where $\tau_{j,m}$ is actual program participation, and $p_{j,m}^L = \tau_{j,m}0 + (1 - \tau_{j,m})p_{j,m}$ and $p_{j,m}^H = p_{j,m}$ are actual effective fees for low- and high-income students, respectively, with $p_{j,m}$ the fees observed in the data. Variables $\bar{p}_{j,m}^L = \bar{\tau}_{j,m}0 + (1 - \bar{\tau}_{j,m})\bar{p}_{j,m}^0$ and $\bar{p}_{j,m}^H = \bar{\tau}_{j,m}\bar{p}_{j,m}^1 + (1 - \bar{\tau}_{j,m})\bar{p}_{j,m}^0$ are the corresponding model's analogs for effective prices. Scalar J_m is the total number of private-voucher schools in market m . Sample moment (12) is implied by the orthogonality condition between the error term and the observables Z^ω . Moments (13) and (14) are similarly motivated. Moment (15) compares actual quality levels of participants with those implied by the model. Market-level moments (16) and (17) minimize the difference between actual market-level averages and those implied by the model for the product of school quality and effective fees for low- and high-income students, which capture some of the cross-market correlation between quality and prices.

The MPEC algorithm simultaneously searches over parameters and endogenous variables to minimize the objective function subject to equilibrium conditions to be satisfied. As in many industrial organization applications, this problem is sparse, since the equilibrium conditions need

to hold for each market separately. The sparsity of the problem allows one to find a solution relatively quickly compared with other approaches, such as nested fixed-point algorithms, which require solving for equilibria at each guess of the parameters (Dubé et al., 2012; Su and Judd, 2012).²⁴ Appendix D describes the estimation procedure in more detail.

6 Results and Policy Analysis

6.1 Estimates

6.1.1 School Quality

I perform estimation imposing a selection-on-observables assumption (Allende, 2019; Allende et al., 2019; Neilson, 2025), where I run a linear regression of test scores on a large set of observable characteristics and school indicators, market by market. I use the student-level average of 4th grade national standardized scores in verbal and math exams. The set of covariates I include in the regressions are gender, low-income status, grade repetition, mother’s education, father’s education, household income, having attended pre-K, having attended kindergarten, computer availability at home, internet availability at home, indigenous status, number of books at home, and class attendance. For brevity, I do not present each market’s estimates, but they are available upon request.

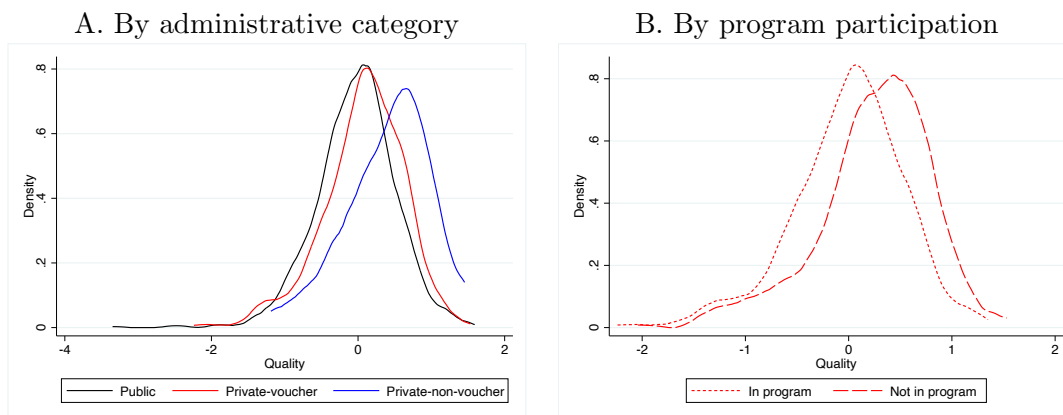
As described in Section 5, I use a parametric normal-normal procedure to shrink the school-quality estimates obtained from the test score regressions (Walters, 2024). Figure E.1 in Appendix E compares the original and shrunk school-quality estimates. Similar to the results in Neilson (2025), I find very small differences between the original and shrunk estimates. This finding is due to the reasonably large size of schools in my sample—i.e., the number of students taking the tests per school—which reduces bias due to sampling variance in the quality estimates. Nonetheless, in the remainder of my empirical analysis I use the shrinkage-adjusted quality estimates.

I summarize the shrunk school-quality estimates using density plots. Notable patterns emerge when the quality distribution is separated by schools’ administrative category. Panel A in Figure 4 shows that the quality distribution for private non-voucher schools is shifted to the right of that for private-voucher schools, which in turn is shifted to the right of the distribution for public schools. Nevertheless, there is substantial overlap between the quality distributions of public and private-voucher schools. These results are consistent with the evidence in Neilson (2025), who uses similar data and a similar procedure to estimate school quality.

²⁴In practice, I implement the GMM-MPEC routine in MATLAB/TOMLAB using the professional solver KNITRO. I provide the automatic differentiated Jacobian and (approximated) Hessian using TOMLAB’s TomSym package.

Also noteworthy is the difference in quality distributions between private-voucher schools that participate in the targeted voucher program and those that do not. Panel B in Figure 4 shows that the program attracted many low-quality schools. While there is some overlap between the two groups' distributions, most high-quality schools chose not to participate in the program.

Figure 4: Distribution of School Quality



Notes: Panel A displays nonparametric estimates of the distribution of school quality by the schools' administrative category. Panel B displays the same for private-voucher schools, by participation status in the targeted program.

6.1.2 Preferences

First Stage. As stated in Section 5, I estimate the first stage regressions for the endogenous variables (top-up fees and program participation) using instrumental variables and school characteristics. Table 5 presents the results for the low- and high-income student versions of equation (8)'s first stage. The excluded instruments are strong predictors for both endogenous variables, as indicated by the high F-statistics. The local share of eligible students and the ex-ante profit indicator are positively associated with program participation, while local competition and labor costs affect top-up fees as expected—i.e., schools best-respond to more competitors by lowering fees, and increase their price when faced with rising costs. Historically high-fee markets predict post-reform high fees and low participation rates. Quality and other school characteristics also play a role in both decisions. These first-stage estimates confirm the relevance of the chosen instruments for the subsequent 2SLS estimation of mean preference parameters.²⁵

²⁵Note that Sargan overidentification tests on the excluded instruments were performed for each set of first-stage regressions, rejecting the null hypothesis of exogeneity in the full set of instruments.

Table 5: First Stage Estimates

	top-up fees		program participation	
	coef.	std. err.	coef.	std. err.
<i>A. Low-income</i>				
pre-reform share of low-income students	0.006	0.044	0.130	0.029
local share of private competitors	-0.163	0.032	0.086	0.021
pre-reform competitors' fees	0.265	0.024	-0.078	0.016
ex-ante profit > 0	-1.351	0.041	0.572	0.028
labor costs	0.005	0.059	-0.022	0.040
quality	0.057	0.018	-0.042	0.012
secular	-0.007	0.018	0.002	0.012
public	-0.156	0.023	0.210	0.016
full day shift	0.067	0.027	-0.020	0.018
intercept	1.479	0.054	0.154	0.037
F excluded IV	278.34		112.61	
<i>B. High-income</i>				
pre-reform share of low-income students	0.344	0.116	0.141	0.039
local share of private competitors	-0.985	0.076	0.130	0.026
pre-reform competitors' fees	0.558	0.028	-0.028	0.010
ex-ante profit > 0	-2.155	0.061	0.515	0.021
labor costs	0.588	0.128	-0.086	0.043
quality	0.195	0.044	-0.101	0.015
secular	0.040	0.044	0.004	0.015
public	-0.687	0.058	0.348	0.020
full day shift	0.228	0.064	-0.042	0.022
intercept	2.799	0.104	0.072	0.035
F excluded IV	638.40		192.40	
no. of schools	2,069			

Notes: Results from the first stage of 2SLS estimation of mean preference parameters. Panel A displays estimates from schools' top-up fees and program participation on instruments and school characteristics, weighting each observation by the inverse of the maximum likelihood-estimated variance of the school fixed effect for low-income students. Panel B does the same for the estimated school mean preference for high-income students. Fee amounts are in real \$1,000 for the year 2013, and were transformed from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD).

Preference parameters. I present estimates for families' preferences in Table 6, separated by low- and high-income groups. Within each family income group, preferences for most school characteristics are allowed to vary by mother's educational level, where the omitted category

is “more than high school.” Low-income families are more price sensitive than higher-income families, which is consistent with the evidence in related studies (Neilson, 2025). There is also important heterogeneity within each family income group: families with less educated mothers are more responsive to price changes than families with more educated mothers. Interestingly, most families have a distaste for schools that participate in the targeted program, especially among high-income households, implying a “stigma” consequence of program participation.²⁶ Both low- and high-income families value higher levels of school quality, although estimates suggest that this attribute is more important for higher-income families. Preferences for secular, public, and full day shift schools also vary by income and mother’s education. Both low- and high-income families prefer schools that are closer to their homes, with the effect being stronger for higher-income families. Overall, the estimates highlight significant heterogeneity in school choice preferences across socioeconomic and mother’s educational backgrounds.

²⁶One possible explanation to this “stigma” result is that program participation changes peer composition (e.g., by attracting more low-income students) and that families are not indifferent to peers. However, participation may also proxy for other school attributes valued by higher-income families. While interesting, distinguishing whether the stigma attached to participation originates from preferences for peers versus other channels is outside the scope of this paper.

Table 6: Estimates from Demand Model

	low-income		high-income	
	coef.	std. err.	coef.	std. err.
fees	-0.591	0.233	-0.121	0.051
fees \times mother's education: less than high school	-0.980	0.060	-1.169	0.029
fees \times mother's education: high school	-0.357	0.036	-0.428	0.012
fees \times mother's education: missing	-0.029	0.044	0.046	0.014
in targeted program	-1.060	0.544	-2.218	0.273
in targeted program \times mother's education: less than high school	1.047	0.069	0.862	0.000
in targeted program \times mother's education: high school	0.619	0.056	0.678	0.037
in targeted program \times mother's education: missing	1.183	0.101	0.977	0.077
quality	0.040	0.035	0.417	0.049
secular	0.037	0.033	0.018	0.043
secular \times mother's education: less than high school	0.047	0.036	-0.050	0.037
secular \times mother's education: high school	0.049	0.035	-0.082	0.029
secular \times mother's education: missing	0.166	0.051	0.011	0.050
public	-0.906	0.085	-0.378	0.091
public \times mother's education: less than high school	1.003	0.042	0.762	0.046
public \times mother's education: high school	0.355	0.041	0.248	0.039
public \times mother's education: missing	0.983	0.059	1.055	0.068
full day shift	0.022	0.049	0.103	0.064
full day shift \times mother's education: less than high school	0.134	0.052	0.001	0.054
full day shift \times mother's education: high school	-0.011	0.050	-0.146	0.042
full day shift \times mother's education: missing	-0.002	0.074	-0.215	0.079
distance	-0.053	0.001	-0.046	0.000
distance \times mother's education: less than high school	0.046	0.001	-0.085	0.001
distance \times mother's education: high school	0.045	0.001	-0.036	0.001
distance \times mother's education: missing	-0.465	0.008	-0.019	0.001
distance: non-missing	0.070	0.061	-0.011	0.073
intercept	0.884	0.462	1.500	0.207
no. of students	77,656			

Notes: Results from maximum likelihood estimation of distance and preference heterogeneity by mother's education, and from 2SLS estimation of mean preference parameters. Omitted mother's level of education category is "more than high school". The distance-from-home-to-school variables are constructed as the interaction of a non-missing indicator and the corresponding distance variables. The non-missing dummy variable is also added in the model. Distance is measured in 100 meters. Fee amounts are in real \$1,000 for the year 2013, and were converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD).

Willingness to pay for school attributes. To complement the demand estimates, I compute families' willingness to pay (WTP) for school attributes by dividing the corresponding estimated coefficients by the negative of the price coefficient, for each income group and mother's education

level (Berry et al., 1995; Petrin, 2002). Table 7 presents these results. In general, high-income families are willing to pay substantially more than lower-income families for desirable school attributes. Also, within each socioeconomic group, more educated mothers value these attributes more than less educated mothers. One striking difference among socioeconomic groups is observed for what I call the “stigma” effect of program participation. While low-income students with mothers with no formal education are almost unresponsive to the school’s participation decision, high-income students with highly educated mothers are willing to pay up to \$18,331 to avoid enrolling in a participating school. School quality is also an attribute for which high-income families are willing to pay amounts that are orders of magnitude higher than what lower-income households would pay, and for which vulnerable students are not very responsive to improvements in—e.g., while low-income families are willing to pay up to \$68 for an additional standard deviation in school quality, many high-income families would pay more than \$3,500.²⁷ For secular, public, and full day shift attributes, WTP also varies considerably by household’s income and educational level—e.g., while low-income families with mothers with no formal education are willing to pay \$62 for their child to attend a public school, high-income families with highly educated mothers need to be paid \$3,124 to choose a public school. Every family needs to be compensated to increase the distance from their residence to the school; however, high-income families are willing to pay more than low-income households for a school closer to their home.

²⁷A similar pattern of heterogeneous willingness to travel across socioeconomic backgrounds is found in Crema (2024) for the Charlotte-Mecklenburg Schools district in the US.

Table 7: Willingness to Pay for School Attributes

	low-income	high-income
<i>In targeted program</i>		
mother's education: less than high school	-8	-1,051
mother's education: high school	-465	-2,805
mother's education: more than high school	-1,794	-18,331
mother's education: missing	198	-16,547
<i>Quality</i>		
mother's education: less than high school	25	323
mother's education: high school	42	760
mother's education: more than high school	68	3,446
mother's education: missing	65	5,560
<i>Secular</i>		
mother's education: less than high school	53	-25
mother's education: high school	91	-117
mother's education: more than high school	63	149
mother's education: missing	327	387
<i>Public</i>		
mother's education: less than high school	62	298
mother's education: high school	-581	-237
mother's education: more than high school	-1,533	-3,124
mother's education: missing	124	9,027
<i>Full day shift</i>		
mother's education: less than high school	99	81
mother's education: high school	12	-78
mother's education: more than high school	37	851
mother's education: missing	32	-1,493
<i>Distance</i>		
mother's education: less than high school	-4	-102
mother's education: high school	-8	-149
mother's education: more than high school	-90	-380
mother's education: missing	-835	-867

Notes: Willingness to pay for school attributes is calculated by dividing the corresponding estimated coefficient from the demand model by the negative of the estimated price coefficient, for each of low- and high-income students, and for different levels of mother's education within those groups. Willingness to pay is in real USD for the year 2013, and was converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD).

6.1.3 School Costs

I estimate schools' cost structure using all schools and a 10% random sample of students per market. Using a random subsample of students substantially improves the computational speed of the estimation algorithm—which is especially useful when bootstrapping standard errors—at the cost of some precision. Estimated coefficients and bootstrapped standard errors for schools' marginal costs and participation costs are presented in Table 8. I implement a block bootstrap with 50 resamples, sampling markets with replacement (Singleton, 2019; Dinerstein and Smith, 2021). For each marginal cost and participation cost function, I include as covariates a set of school characteristics and the same instruments used for fees and program participation in the demand model. Specifically, I include indicators for religious orientation, for-profit status, and chain membership. In addition, the participation cost function includes pre-reform enrollment to capture differences in school-size affecting the participation cost across schools. The marginal cost functions include the local share of private competitors, a measure of labor costs, and pre-reform competitors' fees as cost shifters. The participation cost function also includes competitors' pre-reform fees as a cost shifter, and adds the pre-reform share of low-income students and an indicator for whether it is ex-ante profitable to join the program.

The top panel of Table 8 presents estimates for the marginal cost functions associated with educating low-income (c_j^L) and high-income (c_j^H) students. The estimated coefficients suggest that secular schools are more efficient than religious schools in the production of education for high-income students, but not so for low-income students. For-profit status considerably increases c_j^L , but decreases c_j^H . Relatedly, Sánchez (2024) and Boggiano et al. (2025) document efficient input use and test score production by for-profit schools, especially for high-ability adolescents. Chain membership is associated with reduced efficiency in the production of education. The instruments' effects on marginal costs are intuitive and consistent with demand estimates: more competition incentivizes cost efficiency, higher input (labor) costs raise marginal costs, historically high-fee neighborhoods are associated with high production costs, especially in the education of high-income students. The estimated standard deviations imply a much more dispersed distribution of c_j^H than that of c_j^L .

The middle panel reports estimates for the participation cost function, κ_j . Schools that are secular and for-profit find it more costly to join the targeted program. The opposite is observed for schools that are members of a chain. Pre-reform enrollment absorbs much of the school-size variation in the fixed cost of participation and therefore has a positive estimated coefficient. The cost shifters are strong predictors of participation cost and mirror the first-stage results of the demand model reported in Table 5. In particular, a higher share of eligible students reduces costs and thus incentivizes participation; likewise, a positive ex-ante profitability indicator increases participation. Proximity to competitors with relatively high fees raises the costs associated with

program participation.

The bottom panel presents the estimated weight on enrollment in schools' objective function. I find a reasonably large positive weight (about 1.9), which implies that, in addition to pure profit considerations, schools place important value on attracting additional students. In other words, a one-student increase in expected enrollment is valued similarly to a nontrivial increase in operating margins, consistent with objectives that combine financial performance with scale (e.g., mission, visibility, future positioning, or fixed-cost spreading). This finding is in line with evidence from other educational contexts and industries in which service providers have objectives that go beyond narrow, static profit (Hackmann, 2019; Singleton, 2019; Blanchard et al., 2025), and it empirically justifies my modeling choice.

Table 8: Estimates for Supply Model

	coef.	std. err.	coef.	std. err.
		c_j^L		c_j^H
secular	0.013	0.986	-0.030	0.203
for-profit	0.467	1.030	-0.006	0.348
belongs to a chain	0.022	1.041	0.016	0.323
local share of private competitors	-0.185	1.485	-0.053	0.356
labor costs	0.762	3.522	0.060	0.783
pre-reform competitors' fees	-0.610	1.183	0.024	0.209
intercept	-0.062	3.178	1.196	0.650
$\ln \sigma_c$	-7.504	15.605	-0.950	28.363
		κ_j		
secular	2.562	3.311		
for-profit	1.097	4.234		
belongs to a chain	-4.045	5.286		
pre-reform enrollment	5.659	7.170		
pre-reform share of low-income students	-25.315	12.030		
ex-ante profit > 0	-69.352	12.668		
pre-reform competitors' fees	14.756	15.366		
intercept	48.712	2.358		
$\ln \sigma_\kappa$	2.678	3.397		
		ψ		
enrollment weight	1.868	1.353		
market fixed effects			Y	
no. of private-voucher schools			956	

Notes: All supply parameters were estimated using a GMM-MPEC procedure. Standard errors were computed using a block bootstrap, where 50 synthetic data sets were generated resampling markets. Costs are in real \$1,000 for the year 2013, and were transformed from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD). All specifications include market fixed effects, whose corresponding estimates are shown in Table F.1 in Appendix F.

Figure 5 presents various plots describing schools' estimated cost structure. Panel A displays the predicted distribution of the fixed cost of participating in the targeted voucher program, $\hat{\kappa}_j$. The average participation cost is substantial (about \$23,649). Nevertheless, heterogeneity across schools is significant: while most schools incur a positive participation cost, one fifth are willing to pay to participate. Furthermore, 72% of the schools that would pay to join the program are low-quality, defined as schools in the bottom half of the quality distribution. On the other hand, the average estimated participation cost for high-quality schools is \$32,914, about 2.6 times the mean participation cost of low-quality schools (\$12,551), as shown in panel B. Panel C extends

this analysis to the full quality distribution by plotting the relationship between participation cost and quality; it reveals a steep positive slope, meaning that higher-quality schools face higher costs of joining the program. This result aligns with evidence in Abdulkadiroglu et al. (2018), who show a negative relationship between schools' participation in the Louisiana Scholarship Program and a measure of school effectiveness. More importantly, this evidence contradicts one of the original motivations for introducing progressive vouchers in Chile: bringing high-quality schools closer to vulnerable students by attracting their participation and thus reducing the prices these schools charge to low-income families.

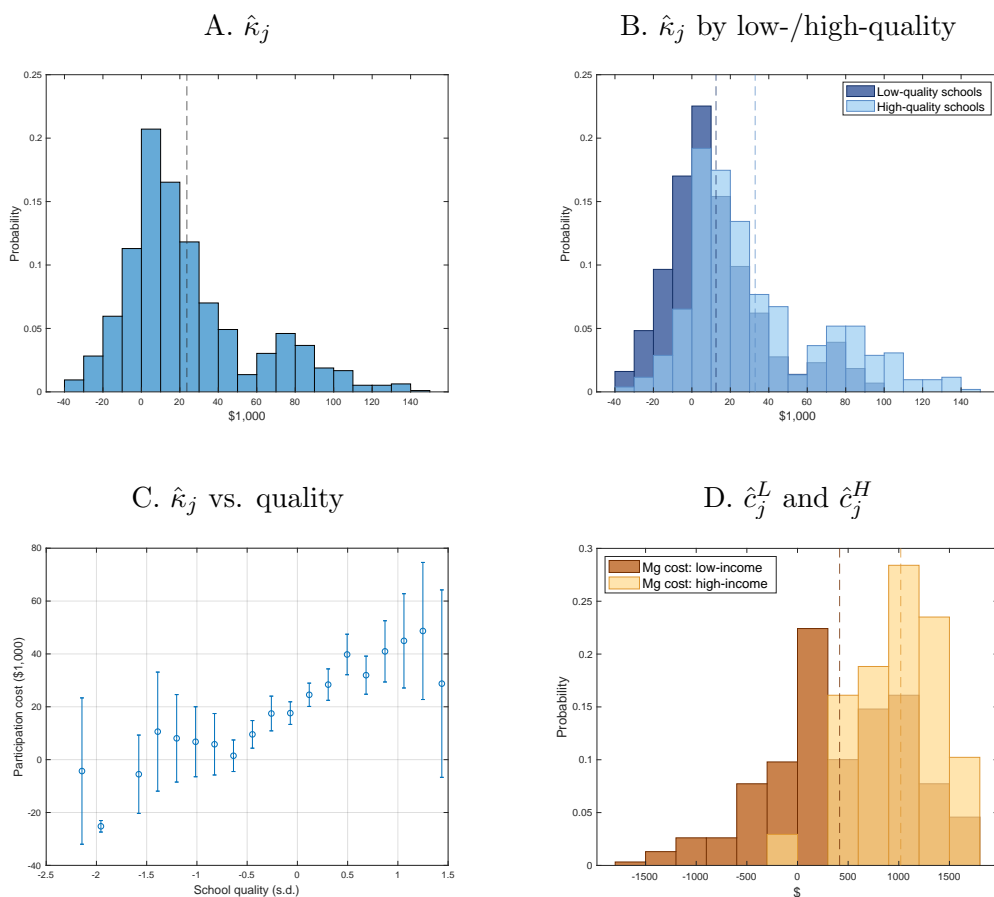
One way to interpret schools' participation cost is as a bureaucracy cost. As explained above, schools that join the program receive additional funding tied to the enrollment of socioeconomically disadvantaged students in exchange for charging zero fees to these students, but they are also required to prepare an improvement plan with clear learning goals, subject to medium-term evaluation by the Ministry of Education. In practice, many schools hire specialized education consulting organizations called ATE (*Asistencia Técnica Educativa*) to help them prepare these plans. In addition, voucher payments depend on students' attendance, which schools must report to the government on a monthly basis. As a result, the targeted voucher adds an extra layer of paperwork to schools' administrative activities. Anecdotally, as part of my interviews with school managers, one respondent mentioned that joining the targeted program forced the school to hire an additional administrative employee to handle the extra paperwork. The annual salary of such a worker in Chile is about \$20,000, which matches the average participation cost estimated by my model.

Another interpretation of the participation cost relates to how closely the targeted program's goals align with a school's own mission. For instance, a school whose mission is to serve vulnerable students may find the program an attractive tool to accomplish its objectives. In contrast, a school focused on teaching very high-ability students, or on educating student-athletes, may perceive the targeted program as poorly aligned with its goals. Vast heterogeneity in Chilean schools' missions, goals, and administrative efficiency likely explains a large share of the heterogeneity in program participation costs that my model estimates.

Panel D displays predicted marginal cost distributions for educating low-income (in dark orange) and high-income (in light orange) students. In contrast to many related studies that estimate marginal costs common to all students, my empirical approach and data allow me to separately identify the costs associated with educating low- and high-income students, thereby testing whether educational costs differ across these groups. I find that this is indeed the case. On average, the marginal cost of educating a low-income student is \$395, which is lower than both the targeted voucher (\$706) and the average marginal cost of educating a student from a higher-income family (\$1,000). Since high-income families particularly value school amenities,

educational inputs, and quality, educating students from that income group involves a greater spending effort than educating students from disadvantaged backgrounds. Nevertheless, heterogeneity is again key, and some overlap between the two distributions is observed.

Figure 5: Estimated Cost Structure



Notes: Estimated participation and marginal costs are in real \$1,000 for the year 2013, and were converted from CLP to USD using the exchange rate as of March 1, 2013 (472.96 CLP/USD). The mean values of distributions are indicated by vertical dashed lines. The marginal cost distributions in panel D are trimmed for display to the range between the 2nd and 98th percentiles.

6.1.4 Model Fit

I test the goodness of fit of my model by comparing a set of key moments in the data with those predicted by the estimated model. Table 9 presents this comparison. I focus on schools' equilibrium strategies, i.e., program participation rate, and averages of effective top-up fees for low- ($p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$) and high-income ($p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$) students. Additionally,

I include quality sorting in program participation and market-level averages of the product of school quality and effective fees. The model matches the actual data very closely, especially for schools' program participation and fee strategies. The quality of participants and even the market-level covariation between quality and fees are also well approximated by my model. These results indicate that the estimated model provides a good fit to the observed data and captures the main features of schools' strategic behavior and sorting patterns.²⁸

Table 9: Goodness of Fit

	actual	model
program participation (τ_j)	0.679	0.679
effective top-up fees for low-income (p_j^L)	262	269
effective top-up fees for high-income (p_j^H)	349	345
quality of participants ($q_j \mid \tau_j = 1$)	-0.009	0.007
$1/J_m \sum_j (q_j \times p_j^L)$	76	69
$1/J_m \sum_j (q_j \times p_j^H)$	92	84

Notes: This table displays averages for a set of key moments comparing actual data and model predictions. Program participation (τ_j) is the share of private-voucher schools participating in the targeted voucher program. Effective prices for low- and high-income students ($p_j^L = \tau_j 0 + (1 - \tau_j) p_j^0$, $p_j^H = \tau_j p_j^1 + (1 - \tau_j) p_j^0$) are average fees charged to each group. Quality of participants is the average school quality among participating schools. $1/J_m \sum_j (q_j \times p_j^L)$ and $1/J_m \sum_j (q_j \times p_j^H)$ are market-level averages of the product of school quality and effective price for low- and high-income students, respectively. Fees are in real \$1,000 for the year 2013, and were converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD). Quality is in standard deviations.

6.2 Policy Analysis and Counterfactuals

I am interested in understanding the equilibrium consequences of the actual targeted voucher policy on schools' strategies and incentives, and on students' choices and welfare. I also explore alternative, budget-neutral policy designs that are intended to tackle some of the equilibrium implications of the actual policy that are not necessarily in line with the original motivations of the program. Finally, I examine the equilibrium effects of further expanding the targeted

²⁸An additional test of how well the model with its FCE assumption fits the data compares the school choice probabilities implied by the estimated demand model using observed school strategies in the indirect utility function with those implied by the same estimated demand model but using FCE-predicted school strategies in place of actual strategies (i.e., fees and program-participation choices/probabilities). Figure E.3 in Appendix E shows the corresponding distributions (panel A) and a binscatter comparison (panel B). The distributions of school choice probabilities using actual and FCE-predicted strategies are nearly indistinguishable. The binscatter reveals small differences at low probability values, but these differences mostly reflect underprediction arising from replacing ones with values less than one in the program-participation utility term in the indirect utility function. Overall, the FCE school choice probabilities closely match those implied by the estimated demand model and the data.

voucher to enforce participation from all schools, and learn about the full potential of this policy design. I perform all these exercises by computing simulations from the model and its estimated parameters.

As school-level outcomes in the counterfactuals, I consider equilibrium program participation decisions, fee strategies (both posted and effective prices by students' income group), markups, enrollment, and profits and utility (i.e., profits plus weighted enrollment). Effective fees are the prices actually charged to each income group, and combine schools' participation decision and fee strategies with students' eligibility condition in the following way: $p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$ for low-income families, and $p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$ for high-income. I define a school's markup as $-Q/(\partial Q/\partial p)$, where Q is enrollment, and which may or may not be equivalent to the price (net of subsidies and enrollment weight) minus the marginal cost, depending on whether the optimal fee strategy is an interior or a corner solution in the school's problem—see equations (2) and (3). Nevertheless, this markup definition mirrors that in Neilson (2025) and other related studies (Armona and Cao, 2024), allowing comparison across studies and contexts.

At the student level, I examine the quality of the attended school, fees actually paid, enrollment in private-voucher schools, and welfare. I compute each of these outcomes separately for low- and high-income students. I calculate the quality of the attended school for a student by multiplying the probability of attending a particular school, as implied by the demand model, with the quality value of that school. Therefore, the average attended school quality for income group d is:

$$AQ(\mathcal{P} | d) = \frac{1}{\sum_i \mathbb{1}[D_i = d]} \sum_i \sum_j P_{ij}(\mathcal{P} | D_i = d) \times q_j,$$

where \mathcal{P} refers to counterfactual policy \mathcal{P} . Similarly, the average fees paid by a student of income group d under policy \mathcal{P} is given by,

$$F(\mathcal{P} | d) = \frac{1}{\sum_i \mathbb{1}[D_i = d]} \sum_i \sum_j P_{ij}(\mathcal{P} | D_i = d) \times p_j^d,$$

where p_j^d is the effective fees paid by a student of group d at school j .

The average student welfare is computed separately by income group using the conventional log-sum formula for consumer surplus in USD terms (Nevo, 2003):

$$CS(\mathcal{P} | d) = \frac{1}{\sum_i \mathbb{1}[D_i = d]} \sum_i \frac{\ln \left[\sum_j \exp(V_{ij}(\mathcal{P} | D_i = d)) \right]}{-\beta_{1i}}.$$

I also compute a second welfare measure, which I call “experience utility”, that subtracts the

disutility from program participation to the indirect utility in the consumer surplus formula,

$$EU(\mathcal{P} | d) = \frac{1}{\sum_i \mathbb{1}[D_i = d]} \sum_i \frac{\ln \left[\sum_j \exp \left(V_{ij}(\mathcal{P} | D_i = d) - \beta_{5i} \tau_j^\zeta \right) \right]}{-\beta_{1i}}.$$

This “experience utility” is informative to policymakers wanting to evaluate consumer gains from new policies focusing on experience as opposed to perception. In a related context, Armona and Cao (2024) call “experience utility” the net indirect utility that removes a perceptual element (in their case, utility from advertising).²⁹

In addition, I track voucher expenditure per student and total government spending in order to assess the fiscal implications of each policy scenario.

6.2.1 Actual Targeted Voucher Policy

I begin by comparing the actual scenario with one in which the targeted voucher is absent. This comparison lets me track the transition from an equilibrium with only the universal voucher to the equilibrium under the actual targeted voucher design—i.e., universal voucher plus targeted voucher at the actual value—highlighting the mechanisms that operate through both extensive (program participation) and intensive (pricing) margins.

Figure 6 summarizes the key responses at the school level. As stated above, my model permits simulating schools’ counterfactual strategies and outcomes independent of the actual choice—e.g. for a school that chooses to accept targeted vouchers, I can simulate the fees it charges when participating of the program, as well as the fees it would charge as a non-participant. I can therefore thoroughly examine the equilibrium incentives at play. Each plot in Figure 6 presents nine bars for schools’ equilibrium strategies/outcomes that are pooled into three groups of three bars each: (i) counterfactual with no targeted vouchers, (ii) counterfactual with targeted vouchers and in the regime the school participates in the program, and (iii) counterfactual with targeted vouchers and in the regime the school does not participate in the program.³⁰ Within each group, the first bar plots the unconditional average strategy/outcome, the second bar plots the average strategy/outcome for schools that join the program when this is in place, and the third bar plots the average strategy/outcome for schools that opt out. Note that strategies/outcomes in the scenario without targeted vouchers are also examined both unconditionally and conditionally on actual participation (in the scenario including targeted vouchers).

²⁹See, also, Allcott (2013).

³⁰The plot in panel B differs slightly in the presentation of the results. The first group of bars shows effective fees in the scenario without targeted vouchers, the second group shows effective fees charged to low-income students in the scenario with targeted vouchers, and the third group shows effective fees charged to high-income students with the targeted voucher policy in place.

When the targeted voucher is introduced, schools’ counterfactual fee strategies decrease on average from $E[p_j] = 653$ to $E[p_j^1] = 268$ and $E[p_j^0] = 478$ in the participation and non-participation regimes, respectively (panel A). However, selection into participation reflects an interesting market segmentation result. First, as can be seen in Figure E.2 in Appendix E, participants are in general of lower quality than non-participants—on average, participants’ quality is 0.007 s.d. and non-participants’ quality is 0.211 s.d. Second, participants set lower counterfactual prices compared with non-participants, but more importantly, participants reduce their fees conditional on participation— $E[p_j^1 | \tau_j = 1] = 112 < E[p_j | \tau_j = 1] = 647$ —while non-participants increase their fees conditional on non-participation— $E[p_j^0 | \tau_j = 0] = 839 > E[p_j | \tau_j = 0] = 666$. Consequently, some schools become cheaper (participants), while others (non-participants) turn more expensive when targeted vouchers are in place. The stigma effect identified from demand estimates (Tables 6 and 7) opens room for schools to charge high fees in the non-participation regime, given that students are willing to pay large amounts to avoid enrolling in a participant school. Consistent with the (partial) profit-seeking motive of schools, it is those that can charge the highest fees that end up opting out, all else equal. On the other hand, participant schools must lower the fees they charge to high-income students to compensate for the disutility associated with the stigma from participation.³¹

Effective fees to each income group depend on schools’ strategies and students’ eligibility status, i.e. $p_j^L = \tau_j 0 + (1 - \tau_j) p_j^0$ for low-income families, and $p_j^H = \tau_j p_j^1 + (1 - \tau_j) p_j^0$ for high-income. Equilibrium fee strategy responses to the targeted policy translate into lower effective prices on average for both low- and high-income students (panel B)—average fees to low-income students go from \$653 to \$270, whereas average fees to high-income students reduce to \$345. Nevertheless, each socioeconomic group faces a non-degenerate distribution of prices, that include cheaper and more expensive schools relative to the scenario absent the targeted policy. Participants are forced to charge zero to eligible students, but they also considerably reduce the fees they charge to children from high-income families, counteracting the stigma effect on their demand. Non-participants, on the other hand, become more expensive to everyone.

Counterfactual demand for schools in the in-program regime rises when targeted vouchers are introduced, both unconditional and conditional on actual participation (panel C). On the other hand, demand for schools substantially falls for participants and importantly rises for non-participants in the non-participation regime, relative to the no-targeted voucher case. As a result, participants slightly increase their average enrollment (from 46 to 47 students), spurred by a rise in low-income demand (panel D)—consistent with the zero-price constraint for eligible students—and non-participants see a major expansion in their demand (from 33 to 41), benefiting

³¹See Atal et al. (2024) for a similar market segmentation consequence in the retail pharmacy industry in Chile after the introduction of public options.

from the stigma effect of participation on high-income students' demand (panel E). Notably, non-participants would increase both their low-income and high-income demand were they to join the program (panels D and E), but would do so at lower fees (panel A). In contrast, the stigma effect attached to participation allows them to charge high fees without experiencing a dramatic reduction in their demand when choosing not to join the program.

I define a school's markup as $-Q/(\partial Q/\partial p)$, and explore it by income segment in panels F and G. Note that, because corner solutions in prices exist, what I call markups are not necessarily equivalent to prices (net of subsidies and enrollment weight) minus marginal cost—see equations (2) and (3). Nevertheless, this markup definition is the same as the one used in Neilson (2025) and other related studies, enabling comparison. My framework extends the analysis in Neilson (2025) by identifying counterfactual markups and examining them by students' income segment and schools' participation regime.³² Average markups for low-income students decrease only in participant schools (panel F), which is in line with the mandated zero-price for eligible students— $E[M_j^1 | \tau_j = 1] = 925 < E[M_j | \tau_j = 1] = 993$ —and the counterfactual low-fee charged by these schools were they to opt out— $E[M_j^0 | \tau_j = 1] = 961 < E[M_j | \tau_j = 1] = 993$. In contrast, non-participants enjoy an actual average markup for low-income students that is larger than absent the reform, but only if they opt out— $E[M_j^0 | \tau_j = 1] = 982$ and $E[M_j^0 | \tau_j = 0] = 1,085$ vs. $E[M_j | \tau_j = 0] = 1,006$ —which is explained by the stigma effect of participation and the market segmentation result shown in panel A. Markups for high-income students are reduced relative to the no-targeted voucher scenario (panel G). The highest reduction in average high-income students' markup is observed for participants (from \$2,633 absent the reform to \$1,890 with targeted vouchers), who also considerably reduce their price to compensate for the stigma from participation effect (panel A). Once more, the targeted reform results in a segmented market, where participants enjoy lower markups than if they opted out, and non-participants enjoy the highest markups by opting out.

Markups relative to marginal costs and vouchers are high compared with other education contexts. I estimate average actual markup/marginal cost ratios of 2.3–2.7 for low-income students and of 1.9–2.6 for high-income. In contrast, Armona and Cao (2024) estimate such a ratio at about 0.5 for the sub-baccalaureate market in the US. This finding suggests that there is still plenty of room for increasing competition in Chile's primary education industry.


The introduction of targeted vouchers in general raises both counterfactual profits and counterfactual utility (profits plus weighted-enrollment) relative to the no-targeted-voucher baseline (panels H and I)—only participants see a reduction were they not to join the program. With targeted vouchers in place, participation yields higher profits and utility than non-participation, for both actual participants and non-participants, but it is only the former who select into the

³²In contrast to Neilson (2025), I do not examine the quality markdown consequences of progressive vouchers.

program and enjoy higher profits and utility. Conversely, non-participants select out, leaving money and demand on the table. The role of the fixed participation cost is key in explaining this sorting behavior. It is implied from panel I that non-participants bear a participation cost of at least \$32,000 (on average), or about a quarter of their utility in the non-participation regime, which is substantial.³³ On top of that, as was documented in Section 6.1.3, this participation cost positively correlates with school quality, resulting in a lower-quality pool of participants relative to the overall market.

³³The value of \$32,000 is the result of subtracting the average utility from non-participation for non-participants ($E[U_j^0 | \tau_j = 0] = 121,000$) to the average utility from participation also for non-participants ($E[U_j^1 | \tau_j = 0] = 153,000$).

Figure 6: Consequences of Actual Policy - School Level



Notes: This figure displays schools' average equilibrium strategies and outcomes under two counterfactual scenarios: (i) no targeted vouchers and (ii) targeted vouchers at their actual 2013 level. In every panel, the first three bars correspond to the counterfactual with no targeted voucher, unconditional on actual participation in the targeted program, conditional on actual choice being participation, and conditional on actual choice being non-participation; the fourth to sixth bars correspond to the counterfactual with targeted vouchers and in the regime the school participates in the program (unconditional on actual choice, conditional on actual choice being participation, and conditional on actual choice being non-participation); and the seventh to ninth bars correspond to the counterfactual with targeted vouchers and in the regime the school does not participate in the program (unconditional on actual choice, conditional on the school actually being a participant, and conditional on the school actually being a non-participant). Panels A-C: counterfactual fee strategies (p_j^1 , p_j^0), effective fees to each income group ($p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$, $p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$), and total expected demand (4th grade). Panels D-E: decomposition of demand into low- and high-income enrollment. Panels F-G: markups for low- and high-income segments. Panel H: expected profits. Panel I: profits plus weighted enrollment. USD values were converted from CLP at 472.960 CLP/USD according to the exchange rate as of March 1, 2013.

At the student level, the targeted voucher policy generally benefits students in all income segments (Table 10). Both low- and high-income students on average attend schools of higher quality once the targeted vouchers are introduced. Nevertheless, the increase in quality for low-income students is relatively minor—about 0.003 s.d.—whereas high-income students experience a much larger increase, of about 0.098 s.d. Average paid fees fall for the group of low-income students (from \$91 without targeted vouchers to \$38 under the actual policy). On the other hand, high-income students’ average paid fees nearly double, reflecting both their preference to avoid participating schools and the price-based market segmentation result documented above. The private-voucher sector expands enrollment, primarily due to participants’ lower effective prices relative to the no-targeted-voucher baseline. Because families dislike attending participating schools, measured welfare falls on average when the stigma component enters utility; removing it reveals small welfare gains for both groups—\$73 and \$98 for low- and high-income students, respectively. Finally, government spending per student naturally shifts toward low-income students: average paid subsidies rise from \$1,200 to \$1,817 for low-income and fall from \$1,083 to \$936 for high-income. This last figure together with the increase in the probability of attending a private-voucher school indicates a flight from public schools to both the private-voucher and the private-non-voucher sectors for non-vulnerable students.

Table 10: Consequences of Actual Policy - Student Level

	no targeted vouchers	actual
<i>School quality (s.d.)</i>		
low-income	0.074	0.077
high-income	0.130	0.228
<i>Fees paid (\$)</i>		
low-income	91	38
high-income	799	1,519
<i>Attends private-voucher (probability)</i>		
low-income	0.607	0.632
high-income	0.406	0.450
<i>Welfare (\$)</i>		
consumer surplus - low-income	5,092	4,858
consumer surplus - high-income	29,432	24,371
experience utility - low-income	5,092	5,165
experience utility - high-income	29,432	29,530
<i>Government expenditure (\$)</i>		
low-income	1,200	1,817
high-income	1,083	936

Notes: This table shows student-level outcomes under two counterfactual scenarios: (i) no targeted voucher program (only the universal voucher is in place); and (ii) the targeted voucher set at its actual 2013 level. School quality is the average (in student-level s.d. units) of the quality of the attended school. Fees paid are average annual top-up fees actually paid by students. “Attends private-voucher” is the probability that a student is enrolled in a private-voucher school. Consumer surplus is the demand model-implied expected utility. Experience utility recomputes consumer surplus by subtracting the direct (dis)utility component associated to attending a participating school. Government expenditure is the student-average voucher spending (universal plus targeted). Each outcome is shown by student income group. USD values were converted from CLP at 472.96 CLP/USD according to the exchange rate as of March 1, 2013.

In sum, the introduction of targeted vouchers produces notable changes in the market, via school responses to incentives and competition, and students’ school choices. It attracts participation from a meaningful chunk of the private-voucher sector, but produces a segmented market: lower-quality schools self-select into participation and cut prices, while higher-quality schools opt out and raise prices to screen families averse to the program. Students benefit from the increase in funds. Low-income students face much lower effective prices and enjoy slightly higher average school quality, whereas high-income students pay higher prices and shift toward non-participant schools with higher quality; measured welfare falls because of the stigma term, but experience

utility (net of stigma) modestly increases for both groups.

An open question is whether alternative policy parameters would produce equilibrium outcomes that are better aligned with the original policy goals. The challenge of attracting high-quality schools to join the program remains for the actual policy. The results from this section’s analysis motivate designs that incorporate differentiated subsidies that are larger for schools of higher quality. A fixed financial incentive to counteract the large participation cost faced by high-quality schools is also appealing. Armona and Cao (2024) demonstrate that quality-varying policies are useful in producing desired supply responses in the context of higher education in the US. Similarly, Singleton (2019) evaluates a location-dependent subsidy design for charter schools in Florida, that is successful in inducing charter schools to serve poor communities.

6.2.2 Alternative Targeted Voucher Designs

Motivated by the results from the preceding section, I evaluate three budget-neutral redesigns of the targeted voucher intended to (i) raise the average quality of participating schools, and (ii) preserve broad access for low-income students. The first design is a two-tier schedule on quality: schools above the median of the quality distribution receive 150% of the value of the actual targeted voucher; schools at or below the median receive half the value of the actual voucher. The second is a smooth, quality-indexed schedule, $v_j^{t,l} = \gamma \tilde{q}_j$, with slope chosen so that aggregate expenditure matches the actual program—I calibrate $\gamma = 0.57$, and truncate at zero so the lowest-quality school receives zero and all others receive a strictly positive amount. The third adds a fixed participation incentive paid only conditional on joining, which I set to \$20,000, consistent with the average estimated participation cost (Figure 5) and with the annual salary cost of hiring an extra (highly-educated) employee to manage the administrative work related to the targeted program. This design also scales down the per-student targeted voucher to about half the actual value, to keep total government spending approximately constant.

Table 11 shows that all three alternatives slightly adjust overall participation (remaining near 0.66–0.68) but shift the composition: the share of high-quality participants rises (from 0.62 to 0.64–0.68) and the mean quality of participants increases substantially (from 0.007 s.d. to between 0.032 and 0.057 s.d.). The fixed incentive design produces the largest improvement in participant quality, followed by the two-tier and then the linear designs. Since there are schools that charge zero even when out of the program, the share of tuition-free schools is higher than the share of participants. The fixed incentive design is the only one that increases the share of tuition-free schools, relative to the baseline (0.80 vs. 0.78). However, the average quality of tuition-free schools is higher in all three redesigns, with the highest increase produced by the fixed incentive (from 0.022 to 0.034–0.055 s.d.). Average effective fees fall for both income groups, by about 2–16% for low-income and 3–12% for high-income students—once more, the fixed incentive

design has the lowest fees, followed by the two-tier design. This is consistent with high-quality and high-fee schools joining the program and therefore lowering their fees, together with some lower-quality and zero-fee schools opting out.

Student-level average school quality changes little for the low-income group, with the fixed incentive design being the only redesign that slightly increases the quality of the school attended (from 0.077 to 0.078 s.d.), despite the considerable shift in the pool of participants at the school level. The low willingness to pay of low-income students for the various school attributes (Table 7) explains this small change in the quality of the attended school: only large enough changes in school characteristics will induce low-income students to switch schools, and consequently to experience different school attributes. This is not the case for higher-income students, who are willing to pay significant amounts to improve the school attributes they experience. This more socioeconomically advantaged group of students sees a decrease in attended school quality, from 0.228 to 0.211–0.216 s.d. on average, which is explained by their strong aversion to participant schools combined with the reduction in the quality of non-participants. Low-income students pay roughly the same across designs—consistent with their relative immobility already evidenced—whereas higher-income students pay slightly less, with the highest effect found under the fixed incentive design (from \$1,519 to \$1,491–1,514), due to the observed reduction in effective fees. Average consumer surplus for the low-income group improves slightly in all three redesigns, with the largest improvement under the fixed incentive scheme (of about \$27). The higher mobility of high-income students translates into larger differences in average consumer surplus across designs. While the linear scheme reduces this group of students’ consumer surplus by \$8, both the linear and the fixed incentive schemes produce gains of \$21 and \$75, respectively.

Table 11: Counterfactual Results from Alternative Designs

	actual $v^{t,0}$	$1.5v^{t,0} \mid q_j > \text{p50}$ $0.5v^{t,0} \mid q_j \leq \text{p50}$	linear $v^{t,l} = 0.57 \cdot \tilde{q}_j$	fixed incentive
<i>School-level</i>				
program participation (share)	0.679	0.663	0.669	0.679
program participation - high-quality (share)	0.616	0.647	0.637	0.675
quality - participants (s.d.)	0.007	0.043	0.032	0.057
free to low-income students (share)	0.776	0.774	0.773	0.797
quality of tuition-free schools (s.d.)	0.022	0.040	0.034	0.055
effective fees to low-income students (\$)	270	256	264	228
effective fees to high-income students (\$)	345	335	341	305
<i>Student-level</i>				
school quality - low-income (s.d.)	0.077	0.074	0.076	0.078
school quality - high-income (s.d.)	0.228	0.211	0.216	0.214
fees paid - low-income (\$)	38	38	38	37
fees paid - high-income (\$)	1,519	1,504	1,514	1,491
consumer surplus - low-income (\$)	4,858	4,867	4,861	4,885
consumer surplus - high-income (\$)	24,371	24,392	24,363	24,446

Notes: This table shows results from counterfactual exercises that test alternative targeted voucher designs. The first column presents outcomes for the actual targeted voucher program. The second column presents results for a voucher scheme that pays 150% of the value of the actual targeted voucher to schools in the upper half of the quality distribution and 50% of the value of the actual voucher to schools in the bottom half of the quality distribution. Column three shows equilibrium results for a scheme that pays a higher amount to higher-quality schools in a linear fashion, with a slope of 0.61, and with the constraint that the lowest-quality school receives zero and all other schools receive strictly positive amounts. In the last column, private-voucher schools that decide to join the program receive a fixed amount of USD 20,000, and the targeted voucher is set to 49% of the actual value. Each value shows the average over the schools' or the students' distributions. The USD values were converted from CLP at 472.96 CLP/USD according to the exchange rate as of March 1, 2013.

In sum, modest quality differentiation in voucher amounts or a calibrated fixed participation incentive improves the quality mix of participants and zero-fee options without materially increasing fiscal cost. The fixed incentive approach yields the largest gains and stronger price reductions, suggesting that directly offsetting participation costs for higher-quality schools is a relatively efficient lever. On the student side, low-income students prove quite unresponsive to modest changes in the supply, resulting in little variation in their outcomes and welfare. In other words, budget-neutral redesigns of the actual policy are practically insufficient in making meaningful improvements on vulnerable students, despite these policies' sizable effects on the supply side.

6.2.3 Full Expansion Reform

The implications from the introduction of the actual targeted policy, as well as those from the simulated alternative budget-neutral redesigns, motivate the analysis of a full expansion of the targeted voucher. The actual policy represented a major budgetary increase that delivered sizable gains for low-income students. In contrast, despite significant responses on the supply side, budget-neutral redesigns cannot notably improve outcomes or welfare for low-income students relative to the actual targeted policy. The natural next step is to examine the consequences of a full expansion of the current targeted program, in which all private-voucher schools must participate and charge zero top-up fees to low-income students, while keeping both the universal and targeted vouchers at their actual levels.

I compare the implications of this full expansion to the ones under the actual design and the baseline without targeted vouchers. Table 12 summarizes the results. At the school level, participation reaches one by construction, and all schools are free to low-income students. The average quality of these tuition-free schools for low-income students is simply the average of all schools' quality, which is higher than the market-segmented actual policy (0.073 vs. 0.022 s.d.). Average effective fees for low-income students falls to zero, whereas that for high-income students rises relative to the actual policy (\$556 vs. \$345). This increase in fees for non-vulnerable students is implied by the reduced differentiation among schools: since all schools are mandated to participate, there are no schools that are free from the participation stigma anymore—i.e., non-participants—and participants do not have to lower their prices as much to become attractive to high-income students relative to non-participants as is the case with the current policy.

On the student side, the average quality of the attended school for low-income students increases by 6% relative to the actual policy (0.082 vs. 0.077 s.d.), which is larger than the 4% increase from the baseline to the actual program. Low-income families' paid fees mechanically drop to almost zero with the mandated participation, and these students are slightly more likely to enroll in a private-voucher school relative to the actual program, benefitting from the mandated zero fees at public and subsidized private schools. High-income students see a sizable decline in both school quality and paid fees relative to the actual policy, to levels on average on par with those in the scenario without targeted vouchers—0.122 s.d. and \$810 under mandated participation vs. 0.228 s.d. and \$1,519 under the actual policy. Correspondingly, these students' likelihood of attending a private-voucher school falls from 45% under the actual policy to 38%, also approaching the baseline without targeted vouchers (41%). Once more, by eliminating the participation/non-participation dimension of differentiation introduced by the actual policy, the full expansion restores the baseline competition of schools for high-income students, who choose schools of similar quality and price to those in the scenario without targeted vouchers. Average experience utility for low-income students rises under the full expansion by about \$44, or 60% of

the corresponding increase from the baseline to the actual policy. In contrast, high-income students' average experience welfare falls by \$104, on par with the level under the scenario without targeted vouchers.

Fiscal costs increase meaningfully, by about 8% relative to the actual policy (\$117,633 vs. \$108,506), which is fully explained by the additional funds required to enroll the one third of schools that initially chose not to join the current program.

Table 12: Counterfactual Results from Major Expansion Reforms

	no targeted vouchers	actual	mandated participation
<i>School-level</i>			
program participation (share)	0.000	0.679	1.000
free to low-income students (share)	0.467	0.776	1.000
quality of tuition-free schools (s.d.)	0.087	0.022	0.073
effective fees to low-income students (\$)	653	270	0
effective fees to high-income students (\$)	653	345	556
<i>Student-level</i>			
school quality - low-income (s.d.)	0.074	0.077	0.082
school quality - high-income (s.d.)	0.130	0.228	0.122
fees paid - low-income (\$)	91	38	5
fees paid - high-income (\$)	799	1,519	810
attends private-voucher - low-income (probability)	0.607	0.632	0.637
attends private-voucher - high-income (probability)	0.406	0.450	0.381
experience utility - low-income (\$)	5,092	5,165	5,209
experience utility - high-income (\$)	29,432	29,530	29,426
<i>Government-level</i>			
total expenditure (\$1,000)	88,966	108,506	117,633

Notes: This table reports equilibrium outcomes under three scenarios: (i) no targeted voucher (only the universal voucher at its 2013 level), (ii) the actual 2013 policy, and (iii) enforced participation (i.e., all private-voucher schools required to join the targeted program and charge zero top-up fees to low-income students). School-level rows are averages across schools. “Quality of tuition-free schools” refers to schools that are free to low-income students. “Effective fees” are the prices actually faced by each income group, $p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$ and $p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$. “Free to low-income students” is the share of schools with $p^L = 0$. Student-level rows are averages across students. “Experience utility” is the inclusive value (log-sum) in dollar units, removing the direct (dis)utility of attending a participating school. “Government expenditure” is total universal plus targeted payments. USD values use the exchange rate of 472.96 CLP/USD, as of March 1, 2013.

Contrary to budget-neutral redesigns, fully expanding the targeted voucher program to enforce participation from all private-voucher schools delivers substantially larger improvements in low-

income students' outcomes/welfare. Mandating participation shifts the entire private-voucher sector onto the "participant" regime, eliminating the participation margin but preserving schools' pricing responses. This policy forces all private-voucher schools to charge zero fees to low-income students and raises the average quality of the schools that are free to these students. Relative to both the no-targeted-voucher baseline and the actual policy, low-income families face strictly lower effective prices and attend higher-quality schools.

High-income students, in contrast, lose the option to avoid participating schools and face higher effective fees (moving back toward the no-targeted-voucher benchmark) and lower average school quality, as the quality gap between participant and non-participant schools disappears.

Overall, forcing all private-voucher schools into the targeted program is an effective way to improve access and moderately raise school quality for vulnerable students, but it requires a sizeable increase in public spending and redistributes surplus away from higher-income families who lose access to non-participating, high-fee, high-quality schools.

7 Conclusions

I have examined the equilibrium consequences of Chile's targeted voucher policy using a structural model of school demand and competition with endogenous participation and pricing decisions, combined with rich administrative microdata. The policy, which layers a means-tested voucher on top of a universal per-student subsidy, was intended to expand access for low-income students to higher-quality private-voucher schools by compensating schools financially for enrolling disadvantaged children and capping (at zero) the fees they may charge to eligible students.

The analysis shows that the targeted voucher meaningfully changes both schools' and families' behavior relative to a system with only a universal voucher. On the supply side, a sizable share of private-voucher schools select to participate, but participation is strongly negatively selected on quality: higher-quality schools face substantially higher fixed participation costs and therefore are less likely to join. Participating schools lower top-up fees, particularly for high-income students, while non-participants raise them, leading to a segmented market in which lower-quality schools tend to be low-price participants and higher-quality schools tend to be high-price non-participants.

On the demand side, both low- and high-income families value quality and proximity, but they differ markedly in price sensitivity and in their valuation of program participation. Low-income households are very price-sensitive and value free or heavily subsidized options; however, their willingness to pay for improvements in quality is modest. High-income households are less price-sensitive but place substantial weight on avoiding schools that participate in the targeted voucher program, generating a sizable "stigma" effect attached to participation. As a result, when the program is introduced, low-income students benefit from sharply lower effective fees and modestly

higher average quality, while high-income students sort toward non-participating, higher-quality schools and face higher effective prices.

The estimated cost structure clarifies why high-quality schools rarely join the program. The fixed cost of participation is large on average and increasing steeply in school quality. Marginal costs also differ by student socioeconomic group: it is more costly to educate high-income students than low-income students in this setting. Schools place substantial weight on enrollment in their objective function, in addition to profits, which helps explain their strong responses on both the intensive (pricing) and extensive (participation) margins when subsidies change.

Policy counterfactuals indicate that modest, budget-neutral redesigns of the targeted voucher—such as quality-differentiated per-student subsidies or a fixed participation incentive plus a lower per-student voucher—can raise the average quality of participating and tuition-free schools and reduce effective prices. However, these designs have only limited impact on the quality of schools actually attended by low-income students and on their welfare, because low-income households have low willingness to pay for observed school attributes and thus move only in response to very large changes in school characteristics or prices. In contrast, a full expansion that mandates participation by all private-voucher schools, while preserving existing voucher amounts, delivers appreciable gains for low-income students: all private-voucher schools become free for eligible students, and the average quality of tuition-free options rises. These gains come at a nontrivial fiscal cost and at the expense of high-income families, who lose access to non-participating, high-fee, high-quality schools and reallocate toward schools resembling those under a system with only the universal voucher.

Overall, the evidence suggests that Chile’s targeted voucher program, as implemented, succeeds in lowering effective prices and slightly improving school quality for low-income students, but it does so through a market segmentation mechanism that attracts a relatively low-quality set of private-voucher schools into the program while leaving many high-quality schools out. Moreover, equilibrium sorting and the stigma attached to participation dampen the welfare gains that the policy could have delivered. Carefully designed quality-dependent subsidies or fixed incentives that directly offset participation costs can improve the quality mix of participating schools without large increases in total spending, but sizable improvements in outcomes for disadvantaged students require either stronger financial commitments or additional policy instruments targeting school quality and families’ perceptions.

The framework developed in this paper—combining a rich demand system with a supply model that endogenizes both participation and pricing—provides a flexible toolkit for evaluating not only voucher-like policies in education, but also an expanding set of policies with voluntary participation on the supply side, where the ultimate impact of these policies depends on how many suppliers participate, and on the characteristics of these participants (e.g., low- vs. high-quality)—

see Abdulkadiroglu et al. (2018) and Singleton (2019) for voucher and charter school programs in the US, Dinerstein and Smith (2021) for NYC's public schools funding reform, Neilson et al. (2022), Barahona et al. (2025b) and Pal (2025) for examples in higher education, Chittaro and Sánchez (2025) and de Elejalde and Sánchez (2025) for a government-guaranteed credit program, and Atal et al. (2025) for a drug quality-regulation policy, among others.

References

- ABDULKADIROGLU, A., P. A. PATHAK, J. SCHELLENBERG, AND C. R. WALTERS (2020): “Do Parents Value School Effectiveness?” *American Economic Review*, 110, 1502–39.
- ABDULKADIROGLU, A., P. A. PATHAK, AND C. R. WALTERS (2018): “Free to Choose: Can School Choice Reduce Student Achievement?” *American Economic Journal: Applied Economics*, 10, 175–206.
- ALLCOTT, H. (2013): “The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market,” *American Economic Journal: Economic Policy*, 5, 30–66.
- ALLENDE, C. (2019): “Competition Under Social Interactions and the Design of Education Policies,” Working paper.
- ALLENDE, C., F. GALLEGO, AND C. NEILSON (2019): “Approximating the Equilibrium Effects of Informed School Choice,” Working paper.
- ARMONA, L. AND S. CAO (2024): “Redesigning Federal Student Aid in Sub-baccalaureate Higher Education,” Working paper.
- ATAL, J. P., J. CUESTA, AND M. SAETHRE (2025): “Quality Regulation and Competition: Evidence from Pharmaceutical Markets,” Manuscript.
- ATAL, J. P., J. I. CUESTA, F. GONZÁLEZ, AND C. OTERO (2024): “The Economics of the Public Option: Evidence from Local Pharmaceutical Markets,” *American Economic Review*, 114, 615–44.
- BARAHONA, N., C. DOBBIN, J. FUENZALIDA, AND S. OTERO (2025a): “The Effects of Widespread Online Education on Market Structure and Enrollment,” Manuscript.
- BARAHONA, N., C. DOBBIN, AND S. OTERO (2025b): “Equilibrium Price Responses to Targeted Student Financial Aid,” Manuscript.
- BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2015): “Self-confirming equilibrium and model uncertainty,” *American Economic Review*, 105, 646–677.
- BAU, N. (2022): “Estimating an Equilibrium Model of Horizontal Competition in Education,” *Journal of Political Economy*, 130.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.

- BERRY, S. T. AND P. A. HAILE (2014): “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82, 1749–1797.
- BJORKEGREN, D., A. DUHAUT, G. NAGPAL, AND N. TSIVANIDIS (2025): “Public and Private Transit: Evidence from Lagos,” NBER Working Paper 33899.
- BLANCHARD, P., S. FLEITAS, AND R. GONZÁLEZ VALDENEGRO (2025): “Regulation by Public Options: Evidence from Pension Funds,” Manuscript.
- BOGGIANO, B., C. SÁNCHEZ, AND R. VALDEBENITO (2025): “To Profit or not to Profit? Evidence from Chile’s For-Profit Voucher Schools,” Manuscript.
- BUCAREY, A. (2018): “Who Pays for Free College? Crowding Out on Campus,” Manuscript.
- CAÑEDO RIEDEL, G. AND C. SÁNCHEZ (2021): “Supply-side Responses to Targeted Vouchers,” Manuscript.
- CHARNESS, G. AND D. LEVIN (2009): “The origin of the winner’s curse: A laboratory study,” *American Economic Journal: Microeconomics*, 1, 207–236.
- CHETTY, R., J. N. FRIEDMAN, AND J. E. ROCKOFF (2014a): “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” *American Economic Review*, 104, 2593–2632.
- (2014b): “Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” *American Economic Review*, 104, 2633–2679.
- CHITTARO, L. AND C. SÁNCHEZ (2025): “Selection in Crisis Lending: Evidence from Chile’s Government-Guaranteed Loans,” Manuscript.
- CONVERSE, P. E. (2000): “Assessing the Capacity of Mass Electorates,” *Annual Review of Political Science*, 3, 331–353.
- CREMA, A. (2024): “School Competition, Classroom Formation, and Academic Quality,” Manuscript.
- DE ELEJALDE, R. AND C. SÁNCHEZ (2025): “Risk-Shifting Incentives Under Government Credit Guarantees,” Manuscript.
- DEKEL, E., D. FUDENBERG, AND D. K. LEVINE (2004): “Learning to play Bayesian games,” *Games and Economic Behavior*, 46, 282–303.

- DINERSTEIN, M., D. MORALES, C. NEILSON, AND S. OTERO (2023): “Equilibrium Effects of Public Provision in Education Markets,” Working paper.
- DINERSTEIN, M. AND T. D. SMITH (2021): “Quantifying the Supply Response of Private Schools to Public Policies,” *American Economic Review*, 111, 3376–3417.
- DUBÉ, J.-P., J. T. FOX, AND C.-L. SU (2012): “Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation,” *Econometrica*, 80, 2231–2267.
- EPPLE, D., R. E. ROMANO, AND M. URQUIOLA (2017): “School Vouchers: A Survey of the Economics Literature,” *Journal of Economic Literature*, 55, 441–92.
- ESPONDA, I. (2008): “Behavioral equilibrium in economies with adverse selection,” *American Economic Review*, 98, 1269–1291.
- ESPONDA, I. AND D. POUZO (2016): “Berk-Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models,” *Econometrica*, 84, 1093–1130.
- ESPONDA, I. AND E. VESPA (2014): “Hypothetical Thinking and Information Extraction in the Laboratory,” *American Economic Journal: Microeconomics*, 6, 180–202.
- EYSTER, E. AND M. RABIN (2005): “Cursed equilibrium,” *Econometrica*, 73, 1623–1672.
- FACK, G., J. GRENET, AND Y. HE (2019): “Beyond Truth-Telling: Preference Estimation with Centralized School Choice and College Admissions,” *American Economic Review*, 109, 1486–1529.
- FERREYRA, M. M. AND G. KOSENOK (2018): “Charter school entry and school choice: The case of Washington, D.C.” *Journal of Public Economics*, 159, 160–182.
- FUDENBERG, D. (2006): “Advancing beyond Advances in Behavioral Economics,” *Journal of Economic Literature*, 44, 694–711.
- FUDENBERG, D. AND D. K. LEVINE (1993): “Self-Confirming Equilibrium,” *Econometrica*, 61, 523–545.
- GAYNOR, M. AND R. J. TOWN (2011): “Chapter Nine - Competition in Health Care Markets,” in *Handbook of Health Economics*, ed. by M. V. Pauly, T. G. McGuire, and P. P. Barros, Elsevier, vol. 2 of *Handbook of Health Economics*, 499–637.

- GAZMURI, A. M. (2024): “School segregation in the presence of student sorting and cream-skimming: Evidence from a school voucher reform,” *Journal of Public Economics*, 238, 105176.
- HACKMANN, M. B. (2019): “Incentivizing Better Quality of Care: The Role of Medicaid and Competition in the Nursing Home Industry,” *American Economic Review*, 109, 1684–1716.
- HOLT, C. A. AND R. SHERMAN (1994): “The Loser’s Curse,” *American Economic Review*, 84, 642–652.
- HUNEEUS, F., J. P. KABOSKI, M. LARRAIN, S. L. SCHMUKLER, AND M. VERA (2025): “Crisis Credit, Employment Protection, Indebtedness, and Risk,” Manuscript.
- JEHIEL, P. (2005): “Analogy-based expectation equilibrium,” *Journal of Economic Theory*, 123, 81–104.
- JEHIEL, P. AND F. KOESSLER (2008): “Revisiting games of incomplete information with analogy-based expectations,” *Games and Economic Behavior*, 62, 533–557.
- JEHIEL, P. AND D. SAMET (2007): “Valuation equilibrium,” *Theoretical Economics*, 2, 163–185.
- JOHNSON, E. (2023): “Essays on Economics of Higher Education,” University of Chicago PhD Dissertation.
- KAGEL, J. H. AND D. LEVIN (1986): “The Winner’s Curse and Public Information in Common Value Auctions,” *American Economic Review*, 76, 894–920.
- (2002): *Common Value Auctions and The Winner’s Curse*, Princeton: Princeton University Press.
- KANE, T. J. AND D. O. STAIGER (2008): “Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation,” NBER Working paper 14607.
- LAKDAWALA, D. AND T. PHILLIPSON (1998): “Nonprofit Production and Competition,” NBER Working paper 6377.
- LI, F., D. MARTÍNEZ, M. A. MARTÍNEZ, AND C. SÁNCHEZ (2025): “Equilibrium Effects of Immigrants on School Competition,” Manuscript.
- NEILSON, C. (2025): “Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students,” Working paper.

- NEILSON, C., S. GALLEGOS, F. CALLE, AND M. KARNANI (2022): “Screening and Recruiting Talent at Teacher Colleges Using Pre-College Academic Achievement,” Princeton University, Industrial Relations Section Working Paper no. 636.
- NEVO, A. (2001): “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- (2003): “New Products , Quality Changes , and Welfare Measures,” *The Review of Economics and Statistics*, 85, 266–275.
- PAL, J. M. (2025): “Education Policy and the Quality Public Servants,” Manuscript.
- PETRIN, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110, 705–729.
- RIVKIN, S. G., E. A. HANUSHEK, AND J. F. KAIN (2005): “Teachers, Schools, and Academic Achievement,” *Econometrica*, 73, 417–458.
- ROCKOFF, J. E. (2004): “The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data,” *American Economic Review*, 94, 247–252.
- SÁNCHEZ, C. (2024): “The Effects of For-profit and Nonprofit Subsidized Schools on Academic Achievement,” Manuscript.
- SINGLETON, J. D. (2019): “Incentives and the Supply of Effective Charter Schools,” *American Economic Review*, 109, 2568–2612.
- SPIEGLER, R. (2011): *Bounded Rationality and Industrial Organization*, New York: Oxford University Press.
- SU, C.-L. AND K. L. JUDD (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, 80, 2213–2230.
- THALER, R. H. (1988): “Anomalies: The Winner’s Curse,” *The Journal of Economic Perspectives*, 2, 191–202.
- WALTERS, C. (2024): “Empirical Bayes Methods in Labor Economics,” NBER Working Paper 33091.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): “Markov Perfect Industry Dynamics With Many Firms,” *Econometrica*, 76, 1375–1411.

WEINTRAUB, G. Y., B. V. ROY, AND C. L. BENKARD (2010): “Computational Methods for Oblivious Equilibrium,” *Operations Research*, 58, 1247–1265.

WOLLMANN, T. G. (2018): “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles,” *American Economic Review*, 108, 1364–1406.

A Data

Below, I present a detailed description of the main data sets used in this paper. These data sets were kindly provided by the Chilean Ministry of Education and Agencia de Calidad de la Educación.

- *Registry of students, 2007, 2013.*

These data provide information on students' gender, date of birth, age, residential address (only in 2013), educational track, grade, class, grade repetition status, special education status, and various characteristics of the school attended, such as municipality, administrative category (public, private-voucher, private-non-voucher), single/double shift schedule, and urban status.

- *Registry of schools, 2007, 2013.*

These data provide information on schools' municipality, administrative category, urban status, address (only in 2013), tuition, religious orientation, and educational track offered.

- *Registry of students that are eligible to participate in the targeted voucher program, 2008–2013.*

These data provide information on the characteristics of students who are eligible to participate in the targeted voucher program, and on whether each eligible student is an actual beneficiary of the program.

- *Registry of schools that participate in the targeted voucher program, 2013.*

These data provide information on schools that participate in the targeted voucher program, including municipality, administrative category, urban status, number of disadvantaged students eligible for the targeted voucher subsidy, and number of students who are beneficiaries of the program.

- *For-/nonprofit status for private-voucher schools, 2013.*

These data contain a school-level indicator of probable for-/nonprofit status for each private-voucher school. The data come from tax records, which allow imperfect identification of the profit motive of each school.

- *National standardized exams (SIMCE) for 4th graders, student-level, 2013.*

These data provide information on students' test scores for three different subjects: verbal, mathematics, and natural sciences.

- *4th grade SIMCE's questionnaire to parents and tutors, 2013.*

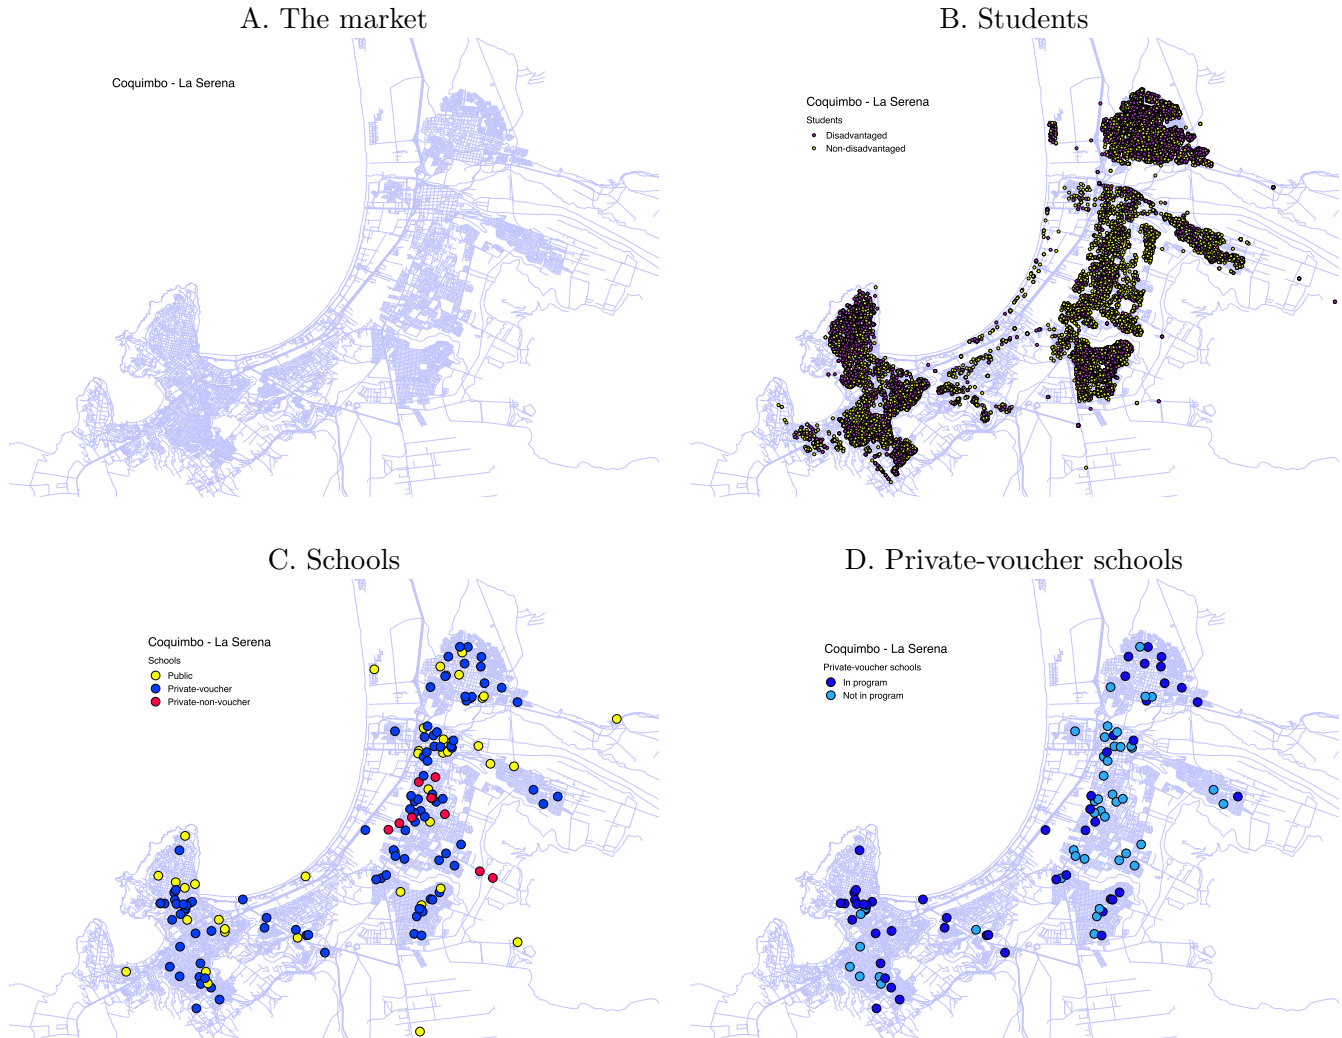
These data consist of the responses to a survey that parents and tutors answer during the days when the national standardized tests are taken. The survey is voluntary, though

more than 90% of parents respond to it every year. It provides information on students' household size, house amenities, total number of books available in the house, household total monthly income, parents and tutors' time use, education, indigenous identification, occupation, health insurance, participation in social programs, reasons for the choice of the school, beliefs about the student's future educational attainment, satisfaction with the school, knowledge of the school's average performance in standardized tests, total monthly expenses related to the student's education other than tuition, and school's admission criteria, tuition and fees.

B A Market Example

Figure B.1 presents an example of an educational market created with the geocoded data. The market is formed by the municipalities of Coquimbo and La Serena in Northern Chile. Panel A displays the streets and roads layout for the market. Panel B displays the spatial distribution of students' residences within the market. It distinguishes between economically disadvantaged (in purple) and non-disadvantaged (in yellow) students. Notice that it is possible to identify neighborhoods with high and low concentrations of disadvantaged students. Panel C displays the spatial distribution of schools within the market, distinguishing between public (in yellow), private-voucher (in blue), and private-non-voucher (in red) schools. Here, we can also identify areas with different concentrations of privately managed schools. Finally, Panel D displays the spatial distribution of private-voucher schools, distinguishing between schools that participate in the targeted voucher program (in blue) and those that do not (in light blue). Not surprisingly, neighborhoods with high concentrations of disadvantaged students (in Panel B) also have high concentrations of schools that opted to participate in the targeted voucher program. Nonetheless, both groups of schools are found in all neighborhoods.

Figure B.1: Educational Market: Coquimbo-La Serena



Notes: This figure presents the educational market formed by the municipalities of Coquimbo and La Serena, in Northern Chile.

C Monte Carlo Simulations

This appendix provides details on the Monte Carlo exercise I perform to study the relation between Bayesian Nash equilibrium (BNE), fully cursed equilibrium (FCE), and partial cursed equilibrium (PCE). I compare the equilibrium strategies resulting from each equilibrium concept under the same primitives of the model. I consider a simplified school choice and competition model, very similar to the model I study in the main text.

Section C.1 specifies the model. Section C.2 describes the data generating processes. Section C.3 compares the equilibrium strategies resulting from the BNE, FCE, and PCE equilibrium assumptions.

C.1 Model Specification


Market Size. I consider a single market where $I = 500$ students enroll in their most preferred school among the J schools that are available in the market. I vary the total number of schools across simulations to analyze how the size of the game affects equilibria. I experiment with $J = 5, 10, 15, 20, 25, 30$.

For simplicity, I consider only private-voucher schools. As in the main model, schools play a static game of incomplete information, with strategies being program participation and fees, for each school type or realization of the idiosyncratic component of the cost structure.

Spatial Configuration. Similar to Fack et al. (2019), the market is modeled as a disc with a radius of 1. Both schools and students are uniformly distributed throughout the market area. The Cartesian distance between student i and school j is denoted by d_{ij} .

Figure C.1 displays the spatial configuration of one of the simulated markets analyzed in the Monte Carlo exercise. The market example consists of 5 schools and 500 students.

Figure C.1: Spatial Configuration of a Simulated Market



Notes: This figure illustrates the spatial distribution of students and schools in one of the simulated markets analyzed in the Monte Carlo simulations. It includes 500 students (in blue) and 5 schools (in red).

Student Preferences. I adopt a parsimonious version of the indirect utility function described in Section 3.1:

$$U_{ij} = \alpha_j^\zeta + \beta_1^\zeta p_j^\zeta + \beta_2^\zeta \tau_j + \beta_3^\zeta d_{ij} + \varepsilon_{ij},$$

where, as in the main text, the superscript $\zeta \in \{L, H\}$ denotes the student's income group. Similarly, $p_j^\zeta = \tau_j(1 - D_i)p_j^1 + (1 - \tau_j)p_j^0$, with D_i denoting student i 's low-income status, τ_j is school j 's decision to participate in the program, and (p_j^1, p_j^0) are school j 's counterfactual fees when the school joins the program and when it opts out, respectively. The parameter α_j^ζ represents school j 's attractiveness to students of group ζ , net of fees, program participation, and proximity. Likewise, β_1^ζ , β_2^ζ , and β_3^ζ capture the preferences of students of income group ζ for school j 's fees, program participation, and distance, respectively. The idiosyncratic error term ε_{ij} follows a Type I extreme value distribution. Consequently, the probability that student i chooses school j is

$$P_{ij} = \frac{\exp\left(\alpha_j^\zeta + \beta_1^\zeta p_j^\zeta + \beta_2^\zeta \tau_j + \beta_3^\zeta d_{ij}\right)}{\sum_k \exp\left(\alpha_k^\zeta + \beta_1^\zeta p_k^\zeta + \beta_2^\zeta \tau_k + \beta_3^\zeta d_{ik}\right)}.$$

I set the share of low-income students to 50%.

School fixed effects are specified as follows:

$$\begin{aligned} \{\alpha_j^L\}_{j=1}^{30} &= \{10, 10.5, \dots, 24.5\}, \\ \{\alpha_j^H\}_{j=1}^{30} &= \{24.5, 24, \dots, 10\}. \end{aligned}$$

The rest of the parameters are set to: $(\beta_1^L, \beta_1^H, \beta_2^L, \beta_2^H, \beta_3^L, \beta_3^H) = (-0.9, -0.3, 0.1, -0.1, -1, -0.5)$.

Schools' Problem. For simplicity, I abstract from the universal voucher. The school's problem is the following:

$$\max_{\tau_j \in \{0,1\}, p_j^1 \geq 0, p_j^0 \geq 0} E_{t-j} \left[\tau_j \left(\Pi_j^1(a_j^1, s_{-j}(t_{-j}); t) - \kappa_j \right) + (1 - \tau_j) \Pi_j^0(a_j^0, s_{-j}(t_{-j}); t) \right],$$

where,

$$\begin{aligned}
\Pi_j^1(a_j^1, s_{-j}(t_{-j}); t) &= (p_j^1 - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^1, s_{-j}(t_{-j}); t) \\
&\quad + (v^t - c_j^L) \sum_i D_i P_{ij}(a_j^1, s_{-j}(t_{-j}); t), \\
\Pi_j^0(a_j^0, s_{-j}(t_{-j}); t) &= (p_j^0 - c_j^H) \sum_i (1 - D_i) P_{ij}(a_j^0, s_{-j}(t_{-j}); t) \\
&\quad + (p_j^0 - c_j^L) \sum_i D_i P_{ij}(a_j^0, s_{-j}(t_{-j}); t).
\end{aligned}$$

I set the targeted voucher parameter to be $v^t = 1.6$. The cost structure terms, (κ_j, c_j^L, c_j^H) , are constructed as follows:

$$\kappa_j = \omega_1^\kappa + \omega_2^\kappa u_j^\kappa, \quad (\text{C.1})$$

$$c_j^L = \omega_1^{c^L} + \omega_2^{c^L} u_j^{c^L}, \quad (\text{C.2})$$

$$c_j^H = \omega_1^{c^H} + \omega_2^{c^H} u_j^{c^H}, \quad (\text{C.3})$$

where $(\omega_1^\kappa, \omega_2^\kappa, \omega_1^{c^L}, \omega_2^{c^L}, \omega_1^{c^H}, \omega_2^{c^H}) = (-50, 150, 0.4, 0.8, 0.2, 0.8)$, and $(u_j^\kappa, u_j^{c^L}, u_j^{c^H})$ are independent, uniformly distributed on the interval $[0, 1]$ random draws, for each j .

C.2 Data Generating Processes

The simulated data are constructed under three different data generating processes (DGPs). For each DGP, I generate $M = 500$ independent samples. For each sample $m = 1, \dots, M$, I randomly draw students' geographic coordinates, $d_{ij}^{(m)}$, income status, $D_i^{(m)}$, and schools' idiosyncratic cost terms, $(u_j^{\kappa, (m)}, u_j^{c^L, (m)}, u_j^{c^H, (m)})$.

DGP 1: Bayesian Nash Equilibrium. This DGP considers a situation where schools play a BN equilibrium; that is, each school takes into account all possible types of its opponents, and the probabilities of each type's occurrence. Moreover, I assume that the distribution of types is known to all schools, and that all schools share the same beliefs over those distributions.

In a BNE, school j 's best-response functions are the following:³⁴

$$\begin{aligned} \hat{p}_j^{1,(m)} \left(a_j^{1,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right) &= \int_{\mathcal{S}_{t_{-j}}} \left[c_j^{H,(m)} - \frac{\sum_i (1 - D_i) P_{ij} \left(a_j^{1,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\sum_i (1 - D_i) \frac{\partial P_{ij} \left(a_j^{1,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^1}} \right] dF(t_{-j}), \\ \hat{p}_j^{0,(m)} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right) &= \int_{\mathcal{S}_{t_{-j}}} \left[c_j^{L,(m)} \frac{\sum_i D_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} + \right. \\ &\quad \left. c_j^{H,(m)} \frac{\sum_i (1 - D_i) \frac{\partial P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} - \frac{\sum_i P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} \right] dF(t_{-j}), \\ \hat{\tau}_j^{(m)} \left(a_j^{(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right) &= \mathbb{1} \left\{ \int_{\mathcal{S}_{t_{-j}}} \left[\left(\Pi_j^1 \left(a_j^{1,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right) - \kappa_j^{(m)} \right) - \Pi_j^0 \left(a_j^{0,(m)}, \hat{s}_{-j}^{(m)}; t_j^{(m)} \right) \right] dF(t_{-j}) > 0 \right\}, \end{aligned}$$

where $F(\cdot)$ is the joint distribution of school types. School j best-responds to its opponents' strategies, which in turn depend on all school types in a highly nonlinear fashion that necessitates equilibrium solving. I utilize an iterative procedure to solve for a BNE, similar to the one in Fack et al. (2019):

1. I define an initial distribution of school types, $F^0(\cdot)$, thus of schools' strategies, for each $m = 1, \dots, 500$ simulated samples.
2. I calculate schools' BN best-responses using initial beliefs $F^0(\cdot)$. I denote these schools' best-responses by $\hat{s}^{(m)}(F^0) = (\hat{\tau}^{(m)}(F^0), \hat{p}^{1,(m)}(F^0), \hat{p}^{0,(m)}(F^0))$.
3. The distribution of schools' best-responses across the M samples jointly determines the "posterior" empirical distribution of school types, $F^1(\cdot)$.
4. Schools best-respond using $F^1(\cdot)$ as their beliefs, and steps 2. and 3. are repeated until a fixed point is found, which occurs when the posterior distribution $F^t(\cdot)$ is consistent with schools' beliefs $F^{t-1}(\cdot)$. The resulting BN equilibrium beliefs are denoted by $F^*(\cdot)$.

The BNE simulated data consist of schools' BN strategies $(\hat{\tau}^{(m)}(F^*), \hat{p}^{1,(m)}(F^*), \hat{p}^{0,(m)}(F^*))$, for each m .

DGP 2: Fully Cursed Equilibrium. This DGP considers the case where schools play an FC equilibrium; that is, each school assumes that each of its opponents plays the average strategy, according to the distribution of types, regardless of their realized type. I, again, assume that the

³⁴For convenience, I set the parameters in the artificial data such that all equilibrium fees are strictly positive.

distribution of types is known to all schools, and that all schools share the same beliefs over those distributions.

I calculate schools' FCE strategies $(\check{\tau}_j^{(m)}, \check{p}^{1,(m)}, \check{p}^{0,(m)})$ for all j, m , in two steps. First, I solve for a fixed point in the following system of equations describing FC schools' average best-responses:

$$\begin{aligned} \bar{p}_j^{1,(m)}(\bar{a}_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= \bar{c}_j^H - \frac{\sum_i (1 - D_i) P_{ij}(\bar{a}_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\sum_i (1 - D_i) \frac{\partial P_{ij}(\bar{a}_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^1}}, \\ \bar{p}_j^{0,(m)}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= \bar{c}_j^L \frac{\sum_i D_i \frac{\partial P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}} + \\ &\quad \bar{c}_j^H \frac{\sum_i (1 - D_i) \frac{\partial P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}} - \frac{\sum_i P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\sum_i \frac{\partial P_{ij}(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}, \\ \bar{\tau}_j^{(m)}(\bar{a}_j^{(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= \Pr \left\{ \left(\Pi_j^1(\bar{a}_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) - \kappa_j^{(m)} \right) - \Pi_j^0(\bar{a}_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) > 0 \right\}, \end{aligned}$$

where $\bar{c}_j^L = \omega_1^{c^L} + \frac{\omega_2^L}{2}$ and $\bar{c}_j^H = \omega_1^{c^H} + \frac{\omega_2^H}{2}$.³⁵ In this first step, other schools evaluate Π_j^1 and Π_j^0 at these average cost values. The only remaining private-information term is u_j^κ , so $\Pr \left\{ (\Pi_j^1 - \kappa_j) - \Pi_j^0 > 0 \right\} = \max \left\{ 0, \min \left\{ 1, \frac{1}{\omega_2^\kappa} (\Pi_j^1 - \Pi_j^0 - \omega_1^\kappa) \right\} \right\}$, due to the uniform $[0, 1]$ distribution assumption.

In the second step, I solve for a fixed point in the following system of equations that defines schools' best-responses to other schools' FC (average) strategies, $(\bar{\tau}_j^{(m)}, \bar{p}^{1,(m)}, \bar{p}^{0,(m)})$, including

³⁵Also, note that $\bar{a}_j^{(m)} = (\bar{\tau}_j^{(m)}, \bar{p}_j^{1,(m)}, \bar{p}_j^{0,(m)})$, $\bar{a}_j^{1,(m)} = (1, \bar{p}_j^{1,(m)}, \bar{p}_j^{0,(m)})$, and $\bar{a}_j^{0,(m)} = (0, \bar{p}_j^{1,(m)}, \bar{p}_j^{0,(m)})$.

schools' own realized idiosyncratic cost terms, $(u_j^{\kappa,(m)}, u_j^{c^L,(m)}, u_j^{c^H,(m)})$.

$$\begin{aligned}
\check{p}_j^{1,(m)}(a_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= c_j^{H,(m)} - \frac{\sum_i (1 - D_i) P_{ij}(a_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\sum_i (1 - D_i) \frac{\partial P_{ij}(a_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^1}}, \\
\check{p}_j^{0,(m)}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= c_j^{L,(m)} \frac{\sum_i D_i \frac{\partial P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}} + \\
&\quad c_j^{H,(m)} \frac{\sum_i (1 - D_i) \frac{\partial P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}} - \frac{\sum_i P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\sum_i \frac{\partial P_{ij}(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)})}{\partial p_j^{0,(m)}}}, \\
\check{\tau}_j^{(m)}(a_j^{(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) &= \mathbb{1} \left\{ \left(\Pi_j^1(a_j^{1,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) - \kappa_j^{(m)} \right) - \Pi_j^0(a_j^{0,(m)}, \bar{s}_{-j}^{(m)}; t_j^{(m)}) > 0 \right\}.
\end{aligned}$$

The resulting vector of schools' FCE strategies, $(\check{\tau}^{(m)}, \check{p}^{1,(m)}, \check{p}^{0,(m)})$, for each m , constitutes the simulated data.

DGP 3: Partial Cursed Equilibrium. This DGP considers a situation where schools play a PC equilibrium; that is, each school is Bayesian with respect to its (geographically) closest opponent, and is fully cursed with respect to all other opponents. The assumption that the distribution of types is known to all schools, and that all schools share the same beliefs over those distributions is maintained.

In a PCE, school j 's best-response functions are the following:

$$\begin{aligned}
\tilde{p}_j^{1,(m)} \left(a_j^{1,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right) &= \int_{\mathcal{S}_{t_{l_j}}} \left[c_j^{H,(m)} - \frac{\sum_i (1 - D_i) P_{ij} \left(a_j^{1,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\sum_i (1 - D_i) \frac{\partial P_{ij} \left(a_j^{1,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^1}} \right] dF(t_{l_j}), \\
\tilde{p}_j^{0,(m)} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right) &= \int_{\mathcal{S}_{t_{l_j}}} \left[c_j^{L,(m)} \frac{\sum_i D_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} + \right. \\
&\quad c_j^{H,(m)} \frac{\sum_i (1 - D_i) \frac{\partial P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^{0,(m)}}}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} - \\
&\quad \left. \frac{\sum_i P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\sum_i \frac{\partial P_{ij} \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right)}{\partial p_j^{0,(m)}}} \right] dF(t_{l_j}), \\
\tilde{\tau}_j^{(m)} \left(a_j^{(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right) &= \mathbb{1} \left\{ \int_{\mathcal{S}_{t_{l_j}}} \left[\left(\Pi_j^1 \left(a_j^{1,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right) - \kappa_j^{(m)} \right) - \right. \right. \\
&\quad \left. \left. \Pi_j^0 \left(a_j^{0,(m)}, (s_{l_j}^{(m)}, \bar{s}_{-j \setminus \{l_j\}}^{(m)}); t_j^{(m)} \right) \right] dF(t_{l_j}) > 0 \right\},
\end{aligned}$$

where l_j is the index for school j 's closest competitor, and $-j \setminus \{l_j\}$ denotes the set of j 's competitors excluding l_j . Similar to my solving for a BNE, I use an iterative procedure to solve for a PCE:

1. I define an initial distribution of closest school types, $F^0(\cdot)$, thus of these schools' strategies, for each $m = 1, \dots, 500$ simulated samples.
2. I calculate schools' PC best-responses using initial beliefs $F^0(\cdot)$. I denote schools' best-responses by $\tilde{s}^{(m)}(F^0) = (\tilde{\tau}^{(m)}(F^0), \tilde{p}^{1,(m)}(F^0), \tilde{p}^{0,(m)}(F^0))$.
3. The distribution of schools' best-responses across the M samples jointly determines the "posterior" empirical distribution of school types, $F^1(\cdot)$.
4. Schools best-respond using $F^1(\cdot)$ as their beliefs, and steps 2. and 3. are repeated until a fixed point is found, which occurs when the posterior distribution $F^t(\cdot)$ is consistent with schools' beliefs $F^{t-1}(\cdot)$. The resulting PC equilibrium beliefs are denoted by $F^*(\cdot)$.

The PCE simulated data consist of schools' PC strategies $(\tilde{\tau}^{(m)}(F^*), \tilde{p}^{1,(m)}(F^*), \tilde{p}^{0,(m)}(F^*))$, for each m .

C.3 Equilibria Comparison

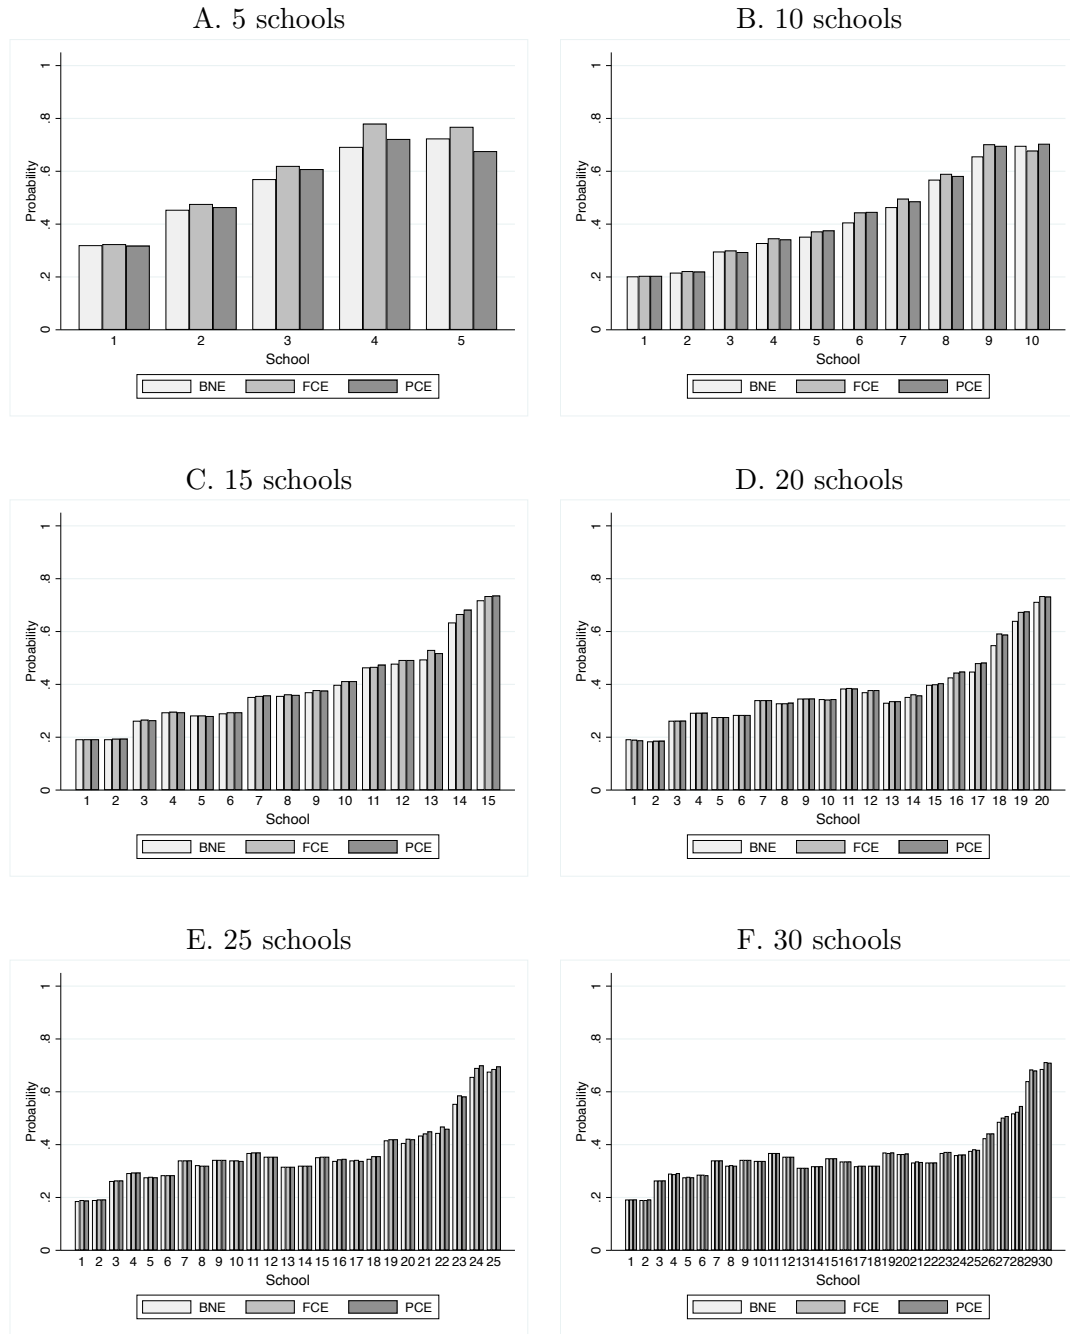
I utilize the data from the Monte Carlo simulations from all three DGPs to examine the relationship between Bayesian Nash equilibrium, fully cursed equilibrium, and partial cursed equilibrium. I compare the equilibrium strategies resulting from each equilibrium concept under the same primitives of the model. Figures C.2–C.4 display the distributions of schools’ equilibrium strategies for program participation (τ_j), fees under participation (p_j^1), and fees under non-participation (p_j^0), respectively, across different market sizes ($J = 5, 10, 15, 20, 25, 30$). Figures C.5 and C.6 present analogous results for equilibrium effective prices for low-income ($p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$) and high-income ($p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$) students.

Results indicate that, although non-negligible differences between equilibrium strategies across BNE, FCE, and PCE exist in small markets (e.g., 5 schools), these differences diminish considerably as the number of schools increases. In particular, Figure C.2 shows that FCE tends to overpredict program participation relative to BNE, especially in very small markets (i.e., 5 schools). PCE program participation strategies also differ from BNE strategies, with these differences being smaller the larger the market. Interestingly, PCE program participation strategies are not necessarily a convex combination of BNE strategies and FCE strategies, which is an implication the highly nonlinear aspect of the game.

In sufficiently large markets, equilibrium program participation strategies from all three equilibrium concepts resemble each other substantially, especially for schools with low participation rates.

Note, too, that schools that are more preferred by low-income students and less preferred by high-income students—i.e., schools with high values of α_j^L and low values of α_j^H in students’ indirect utility function—participate in the program with a relatively higher frequency, across all equilibrium concept assumptions and market sizes.

Figure C.2: Schools' Equilibrium Strategy Distribution Under BNE, FCE, PCE - τ_j

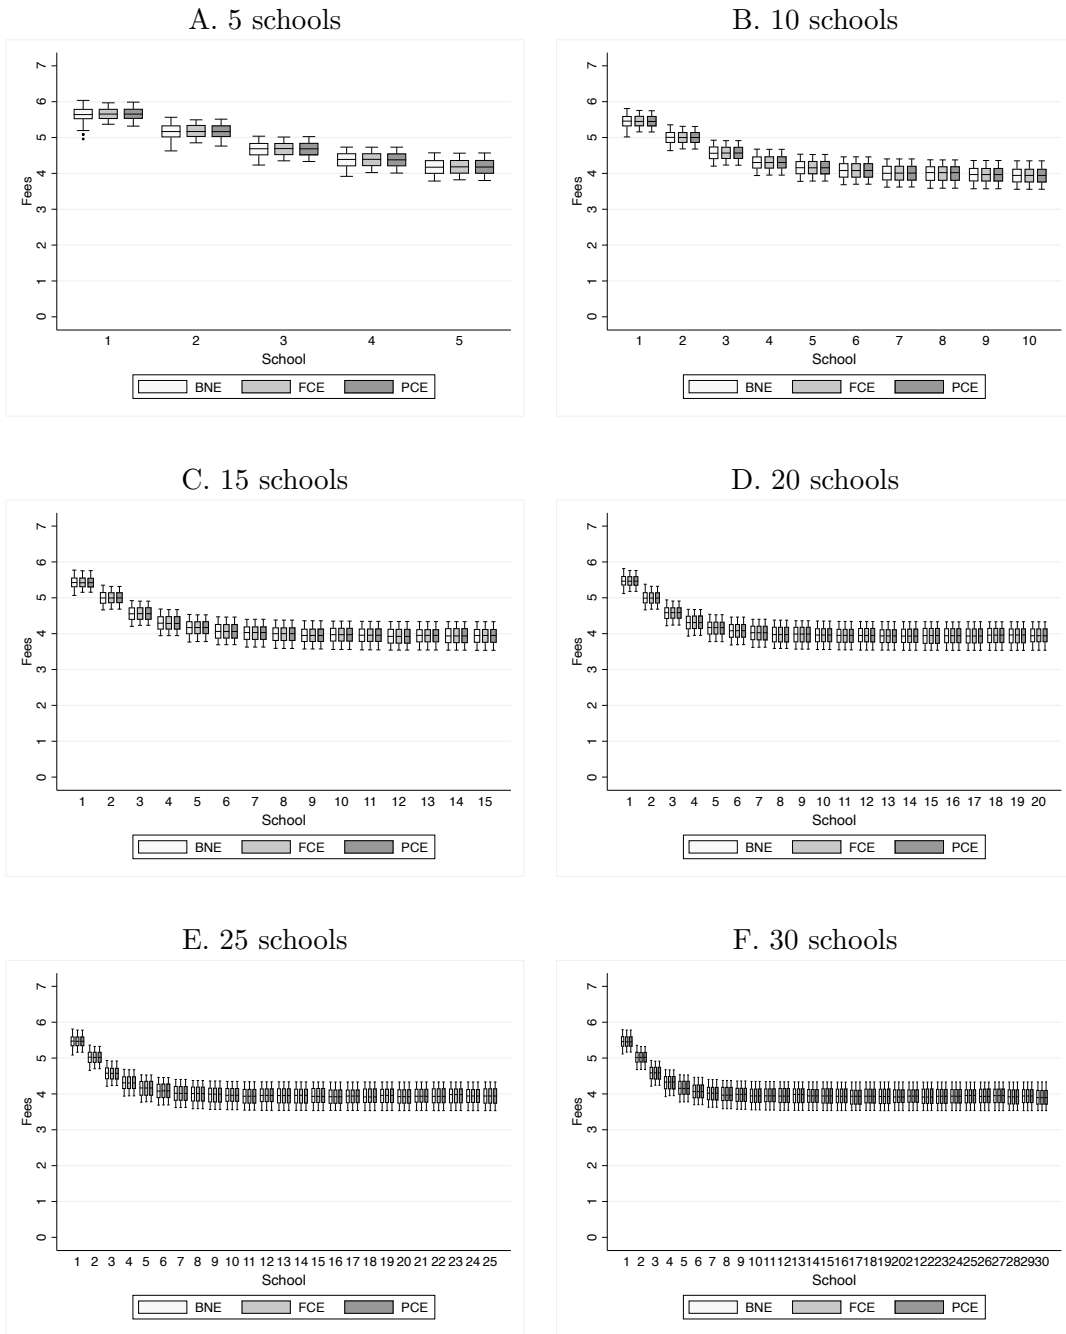


Notes: This figure displays schools' equilibrium program participation, τ_j , under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Panel A shows each school's strategy distribution across simulations for a market with 5 schools. Panels B–F do analogously for markets with 10, 15, 20, 25, and 30 schools.

Schools' counterfactual price distributions are presented in Figures C.3 and C.4. Again, we observe that equilibrium price distributions under the three equilibrium concepts converge rapidly as the number of schools in the market increases. In general, and especially in small markets, compared with BNE, FCE prices present a more compressed distribution, with a lower incidence of extreme values. PCE counterfactual price distributions are also more compressed than BNE distributions, and are somewhat closer to FCE distributions. Nevertheless, all three sets of price distributions are strikingly similar in sufficiently large markets.

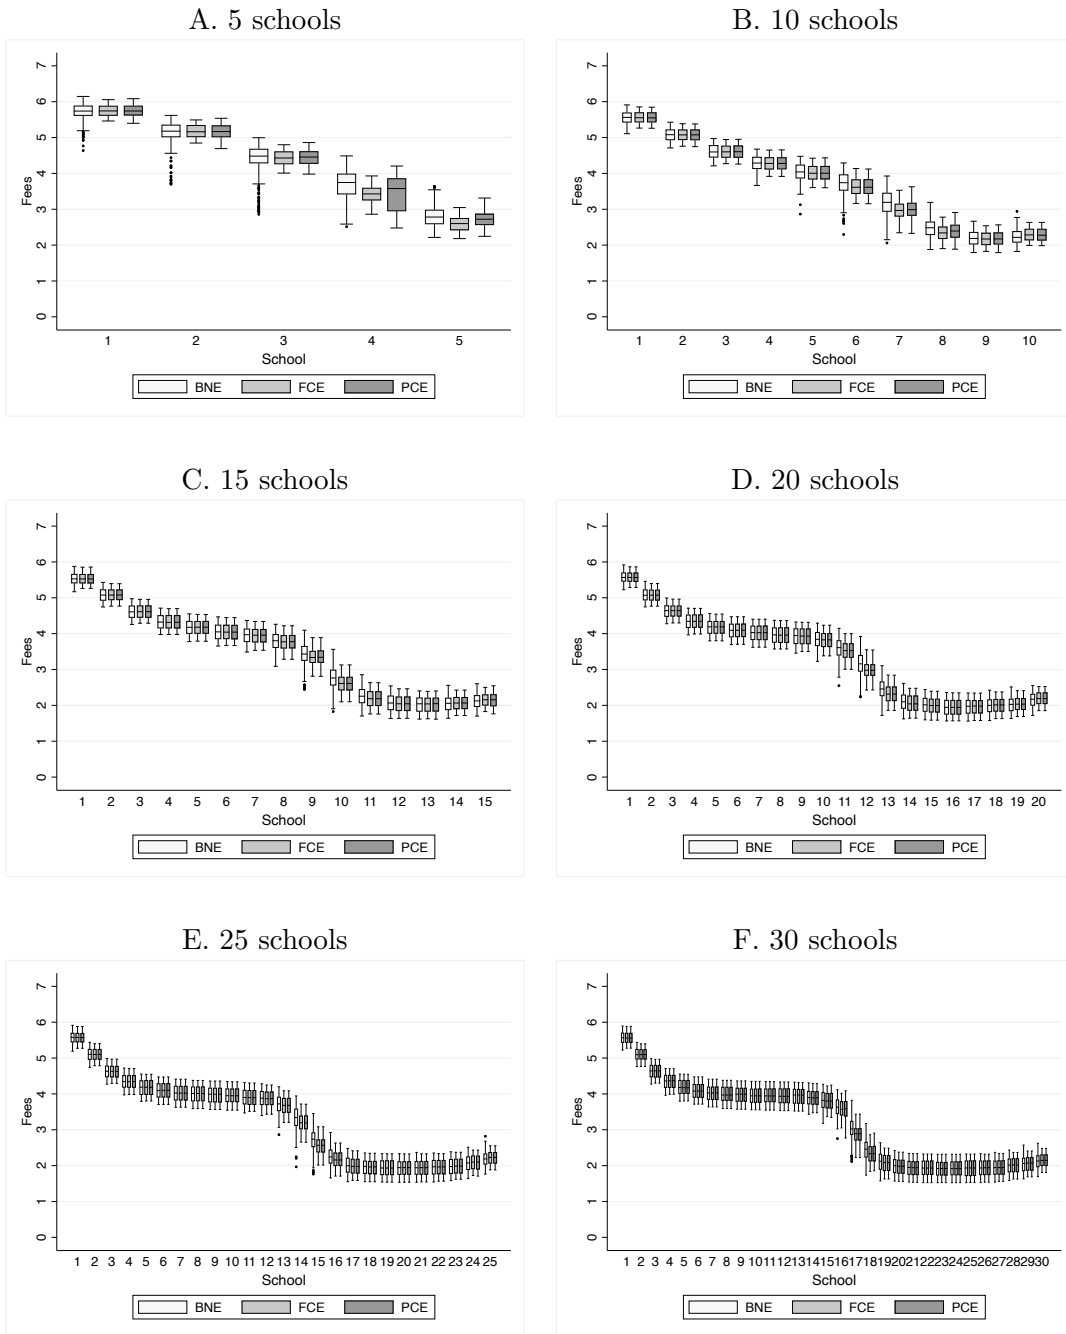
Another important observation is the relatively higher equilibrium prices set by schools that are more preferred by high-income students and less preferred by low-income students—these are schools with a high α_j^H and a low α_j^L in students' indirect utility function. This pattern is present across all three equilibrium concepts, as well as for every market size.

Figure C.3: Schools' Equilibrium Strategy Distribution Under BNE, FCE, PCE - p_j^1



Notes: This figure displays the distribution of schools' equilibrium counterfactual fees conditional on joining the program, p_j^1 , under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Panel A shows each school's strategy distribution across simulations for a market with 5 schools. Panels B-F do analogously for markets with 10, 15, 20, 25, and 30 schools.

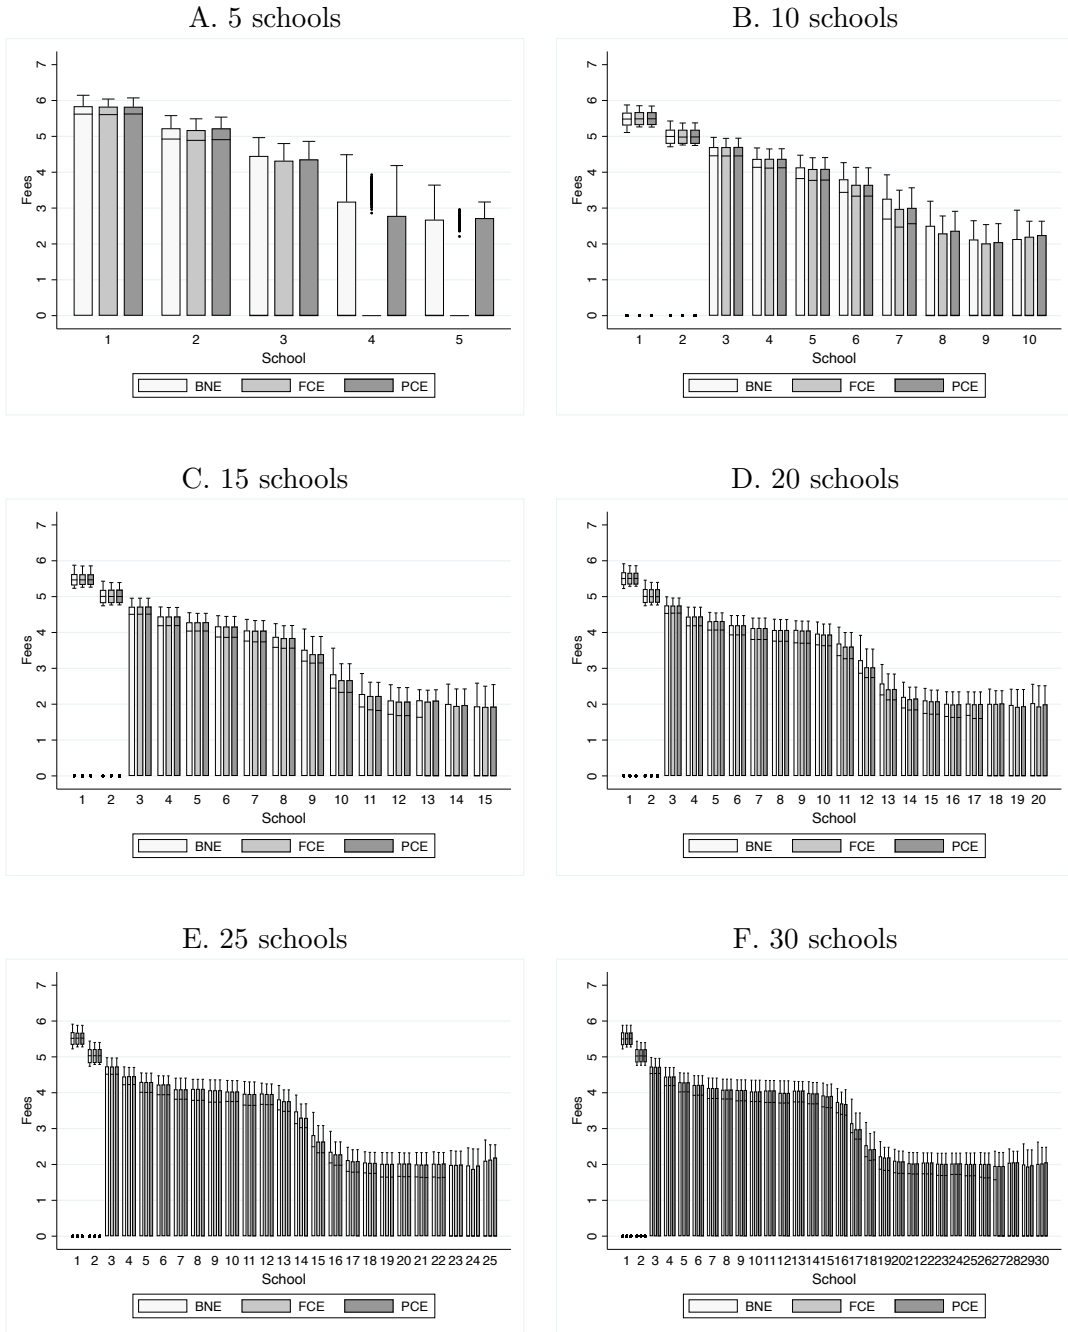
Figure C.4: Schools' Equilibrium Strategy Distribution Under BNE, FCE, PCE - p_j^0



Notes: This figure displays the distribution of schools' equilibrium counterfactual fees conditional on opting out of the program, p_j^0 , under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Panel A shows each school's strategy distribution across simulations for a market with 5 schools. Panels B-F do analogously for markets with 10, 15, 20, 25, and 30 schools.

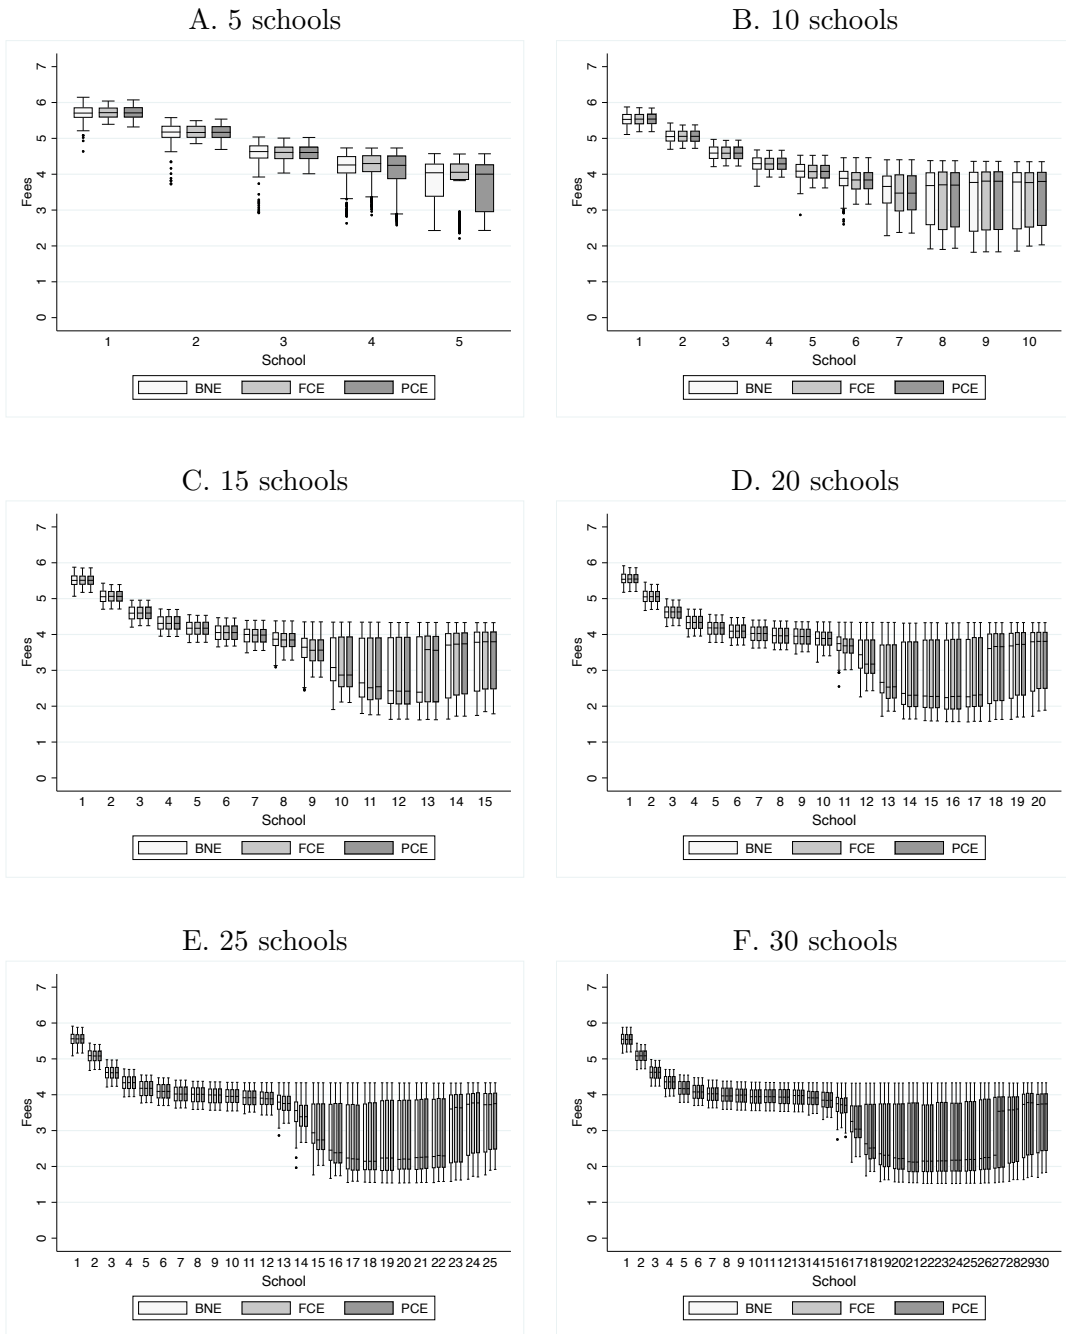
Distributions for equilibrium effective prices for low-income students, $p_j^L = \tau_j 0 + (1 - \tau_j)p_j^0$, and for high-income students, $p_j^H = \tau_j p_j^1 + (1 - \tau_j)p_j^0$, are presented in Figures C.5 and C.6, respectively. Effective prices are empirically relevant, as they are what students observe when shopping for schools, and what students are charged once enrolled in their school of choice. The results for equilibrium effective prices mechanically follow those shown in Figures C.2–C.4, since their definitions are direct functions of schools’ equilibrium strategies. More precisely, equilibrium effective prices across BNE, FCE, and PCE assumptions converge rapidly as the number of schools in the market increases. Some appreciable differences in equilibrium effective prices are observed in the smallest of markets, but they shrink considerably in larger markets. By the same token, FCE and PCE equilibrium effective price distributions are somewhat more compressed than BNE distributions—much less so in sufficiently large markets—and schools more preferred by high-income students and less preferred by low-income students set relatively higher equilibrium effective prices.

Figure C.5: Effective Price Distribution Under BNE, FCE, PCE - p_j^L



Notes: This figure displays the distribution of schools' equilibrium effective prices charged to low-income students, $p_j^L = \tau_j 0 + (1 - \tau_j) p_j^0$, under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Panel A shows each school's effective price distribution for low-income students across simulations for a market with 5 schools. Panels B-F do analogously for markets with 10, 15, 20, 25, and 30 schools.

Figure C.6: Effective Price Distribution Under BNE, FCE, PCE - p_j^H



Notes: This figure displays the distribution of schools' equilibrium effective prices charged to high-income students, $p_j^H = \tau_j p_j^1 + (1 - \tau_j) p_j^0$, under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Panel A shows each school's effective price distribution for high-income students across simulations for a market with 5 schools. Panels B-F do analogously for markets with 10, 15, 20, 25, and 30 schools.

Table C.1 summarizes the distributional evidence for schools' equilibrium strategies presented above. Specifically, Table C.1 presents average equilibrium school strategies at the market level, under each of BNE, FCE, and PCE assumptions, and for $J = 5, 10, 15, 20, 25, 30$ market sizes. The similarity of equilibrium strategies across equilibrium assumptions is more evident when focusing on market averages. See, for instance, the evidence for the average of p_j^1 (second column). For most market sizes, average p_j^1 across equilibrium concepts are indistinguishable from each other.

Also in line with the distributional evidence presented in the preceding figures, PCE average strategies are not necessarily closer to BNE than FCE strategies are.

Notice, too, that average equilibrium counterfactual prices $E[p_j^1]$ and $E[p_j^0]$ (second and third columns) decrease as the number of schools in the market increases, suggesting the presence of a competition effect.

Table C.1: Average Equilibrium Strategies Under BNE, FCE, PCE

	$E[\tau_j]$	$E[p_j^L]$	$E[p_j^0]$	$E[p_j^L]$	$E[p_j^H]$
<i>5 schools</i>					
BNE	0.552	4.812	4.354	2.124	4.699
FCE	0.594	4.818	4.272	1.934	4.701
PCE	0.558	4.815	4.305	2.053	4.660
<i>10 schools</i>					
BNE	0.418	4.349	3.734	2.367	4.121
FCE	0.435	4.349	3.695	2.283	4.109
PCE	0.435	4.349	3.703	2.294	4.117
<i>15 schools</i>					
BNE	0.385	4.210	3.485	2.300	3.879
FCE	0.395	4.211	3.462	2.260	3.878
PCE	0.395	4.211	3.463	2.262	3.881
<i>20 schools</i>					
BNE	0.373	4.148	3.354	2.224	3.739
FCE	0.382	4.148	3.333	2.192	3.741
PCE	0.382	4.148	3.334	2.193	3.742
<i>25 schools</i>					
BNE	0.366	4.107	3.266	2.167	3.648
FCE	0.373	4.107	3.250	2.141	3.647
PCE	0.373	4.107	3.250	2.142	3.649
<i>30 schools</i>					
BNE	0.359	4.075	3.211	2.138	3.583
FCE	0.363	4.075	3.199	2.120	3.582
PCE	0.364	4.075	3.200	2.120	3.584

Notes: This table displays the market's average strategies in equilibrium under BNE, FCE and PCE assumptions. Data are generated from 500 Monte Carlo simulations using the same primitives for students' preferences and schools' cost structure. I use a separate set of 500 simulations for different market sizes, where I vary the number of schools in $J = 5, 10, 15, 20, 25, 30$. Columns report schools' average equilibrium strategies, $E[\tau_j]$, $E[p_j^L]$, and $E[p_j^0]$, as well as the market's average effective prices, $E[p_j^L] = E[\tau_j 0 + (1 - \tau_j)p_j^0]$ for low-income students, and $E[p_j^H] = E[\tau_j p_j^1 + (1 - \tau_j)p_j^0]$ for high-income students.

D Estimation Procedure for Supply

The GMM-MPEC problem for the supply side of the model is

$$\min_{\omega, \bar{s}} g(\omega, \bar{s})' W g(\omega, \bar{s}),$$

subject to

$$h(\omega, \bar{s}),$$

where $g(\omega, \bar{s})$ is the vector of sample moments, and W is a GMM weighting matrix of the form


$$W = \text{diag} \left(\left(\frac{1}{\sum_{m=1}^M J_m} Z^{\omega'} Z^{\omega} \right)^{-1}, \left(\frac{1}{\sum_{m=1}^M J_m} X^{\omega'} X^{\omega} \right)^{-1}, \left(\frac{1}{\sum_{m=1}^M J_m} X^{\omega'} X^{\omega} \right)^{-1}, I_3 \right),$$

where $\text{diag}(v)$ is a matrix with v on its diagonal and zeros elsewhere, and I_3 is the identity matrix of dimension 3. Finally, let $h(\omega, \bar{s})$ be the $3 \sum_{m=1}^M J_m \times 1$ vector of equilibrium conditions (9)–(11).

The MPEC algorithm simultaneously searches over parameters ω and endogenous variables \bar{s} to minimize the objective function subject to the equilibrium conditions. In practice, I implement the GMM-MPEC routine in MATLAB/TOMLAB using the professional solver KNITRO. I provide the automatically differentiated Jacobian and (approximate) Hessian using TOMLAB's TomSym package.

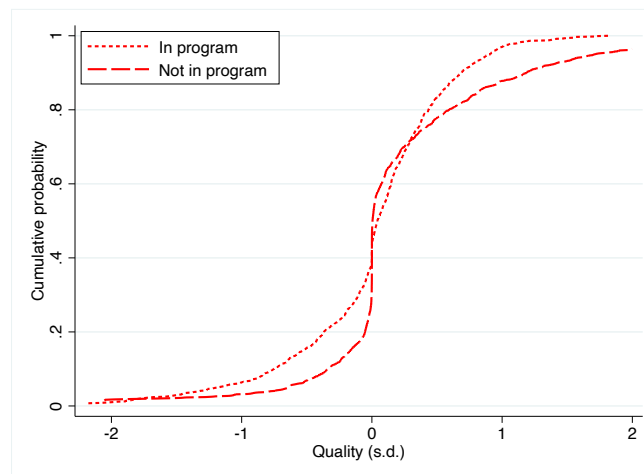
E Additional Figures

Figure E.1: Original and Shrunk School Quality Estimates



Notes: Original school quality estimates from test score regressions are plotted against their shrunk versions, after applying a parametric normal/normal shrinkage procedure.


Figure E.2: Predicted Quality Distributions of Participant/Non-participant Schools




Notes: Predicted quality distributions of participant and non-participant private-voucher schools.

Figure E.3: Actual and FC-Predicted School Choice Probabilities

A. Densities



B. Binscatter



Notes: Comparison of school choice probabilities implied by the estimated demand model using observed strategies with those implied by the same model using FCE strategies in place of actual strategies (i.e., fees and program-participation choices/probabilities). Panel A shows kernel density estimates; panel B shows a binscatter comparison.

F Additional Table

Table F.1: Estimates for Supply Model: Market Fixed Effects

	coef.	std. err.	coef.	std. err.	coef.	std. err.
	c_j^L		c_j^H		κ_j	
market 2	-1.423	3.178	-0.005	0.650	20.864	2.358
market 3	-0.315	15.605	0.492	28.363	26.856	3.397
market 4	-0.039	3.469	0.228	0.893	8.775	7.554
market 5	-0.016	3.018	-0.456	0.809	35.006	7.250
market 6	-0.339	4.665	-0.170	0.572	3.221	5.000
market 7	-0.256	4.481	-0.169	0.516	9.703	17.294
market 8	-0.069	5.494	-0.100	0.857	9.328	10.701
market 9	0.927	6.487	-0.153	14.040	18.926	5.780
market 10	0.424	5.801	0.100	13.974	23.861	11.228
market 11	0.884	7.965	0.374	14.465	18.960	181.173
market 12	0.981	12.185	-0.184	8.488	15.368	174.115
market 13	0.855	27.930	0.290	0.712	14.876	14.451
market 14	1.038	49.598	-0.545	0.702	22.723	12.016
market 15	1.346	51.305	-0.184	55.585	29.779	11.434
market 16	0.862	46.029	-0.043	1.466	1.495	7.817
market 17	0.702	35.412	-0.566	55.528	11.882	75.608
market 18	0.672	24.603	-0.839	1.812	-7.374	86.265
market 19	0.786	3.398	-0.333	1.615	7.194	78.219
market 20	1.239	4.637	0.196	1.242	5.804	14.574
market 21	1.252	4.561	-0.588	1.285	12.145	17.709
market 22	0.059	11.584	-0.692	1.089	-5.463	10.531
market 23	-0.399	18.830	0.185	12.875	8.086	14.632
market 24	3.960	97.184	-2.610	68.618	0.342	11.476
market 25	0.599	9.977	0.352	1.914	25.231	14.892
market 26	-3.745	23.707	0.006	8.915	-23.954	8.523
market 27	1.751	9.170	-1.783	27.491	15.605	2.764
no. of private-voucher schools			956			

Notes: All supply parameters were estimated using a GMM-MPEC procedure. Standard errors were computed using a block bootstrap, where 50 synthetic data sets were generated by resampling markets (Singleton, 2019; Dinerstein and Smith, 2021). Costs are in real \$1,000 for the year 2013, and were converted from CLP to USD according to the exchange rate as of March 1, 2013 (472.96 CLP/USD).