Experimentally Validating Welfare Evaluation of School Vouchers*

Peter Arcidiacono† Karthik Muralidharan‡ John D. Singleton§

September 2024

Abstract

We leverage a unique two-stage experiment that randomized access to private school vouchers across markets as well as students to estimate the revealed preference value of school choice. To do this, we estimate several choice models on data only from control markets before turning to the treatment data for model validation. This exercise reveals that a model where school choice is constrained by ability-to-pay achieves better out-of-sample fit but still underpredicts experimental take-up of the voucher offer. We then present evidence from treatment markets that: a) the voucher offer also induced search; and b) private schools used program surplus to incentivize enrollment. Further, we show that a unified model incorporating these features can explain both the control and treatment data patterns. Estimates from that model imply that a targeted voucher program would have a marginal value of public funds (MVPF) of at least 3.

*We thank Michael Dinerstein, Chao Fu, Adam Kapor, Christopher Neilson, Petra Todd, and seminar participants at Cowles, Barcelona, the Federal Reserve Bank of Minneapolis, NBER Education, Princeton, Toulouse, University of Pennsylvania, and Yale for helpful comments and discussions.

†Duke University, IZA & NBER; psarcidi@econ.duke.edu

‡UC San Diego, JPAL, NBER & BREAD; kamurali@ucsd.edu

§University of Rochester, Hoover Institution & NBER; john.singleton@rochester.edu

1 Introduction

Governments around the world provide in-kind benefits to citizens, including publicly-provided schooling, health care, and food assistance (Currie and Gahvari, 2008). A central question in public economics is the relative efficiency of in-kind provision versus providing beneficiaries with a voucher to purchase the same goods or services on the open market, and there is a large empirical literature studying this question across sectors and contexts.[1] These studies typically evaluate the impact of vouchers on sector-specific outcomes, such as test scores or food consumption and nutrition. This focus may reflect the priorities of taxpayers and policymakers who care about the most cost-effective way to achieve specific outcomes of interest.

Yet, this default approach often ignores the preferences of beneficiaries themselves. For instance, vouchers may enhance their welfare by increasing choice and improving match quality. Therefore, evaluations of voucher programs should account for both impacts on outcomes that a paternalistic policymaker may care about, as well as beneficiary valuation of such programs. Estimating valuation is also critical for predicting program take up rates under different voucher values and targeting rules, which is a key input for policy. However, beneficiary valuation of publicly-provided benefits is often ignored in policy evaluation, in part because it is not easy to estimate.[2]

In this paper, we complement an existing experimental evaluation of the test score impacts of private school vouchers in rural India by also quantifying program effects on welfare based on revealed preference. We do this using a unique research design that randomized access to vouchers across both markets as well as students. Specifically, we estimate several structural econometric models of school choice on data from only control markets of the Andhra Pradesh School Choice project, whose experimental test score impacts are reported in Muralidharan and Sundararaman (2015). Our design then uses the data from treatment markets to validate the models out-of-sample, including against the choice patterns experimentally induced by the randomized voucher offers.[3]

We estimate two classes of choice models on the control markets data in which households select a primary school from among the free government and fee-charging private options in their village. The first are random coefficient logit demand models that are standard for welfare analysis in the industrial organization literature (Berry, Levinsohn and Pakes, 1995; Nevo, 2001; Petrin, 2002),

[1] Illustrative examples include Hastings and Shapiro (2018) and Banerjee et al. (2021) on food stamps or vouchers.

[2] For instance, in their work on distributional national accounts, Piketty, Saez and Zucman (2018) value public goods and publicly-provided private goods at the *cost* of providing them.

[3] To strengthen the design's credibility, the treatment data were embargoed during specification and estimation of the models and the models (and their predictions) pre-committed to in a working paper, Arcidiacono et al. (2021).

and have been applied to other contexts of school choice (e.g. Neilson 2013; Carneiro, Das and Reis 2022). Second, we develop and estimate a random utility discrete choice model that incorporates a constraint on households' ability-to-pay for private schooling in the absence a voucher. The constraint reflects the reality that liquidity and access to credit are often limited in low-income settings, such as rural India.[4] Our constrained model is in the spirit of other applications where choice sets are not observed in the data (e.g. Ben-Akiva and Boccara 1995; Barseghyan et al. 2021) and connects with prior work quantifying the salience of credit constraints.[5]

The estimates show that the ability-to-pay constrained model ascribes greater utility from private schools and characteristics associated with private schools (such as English-language instruction), than do random coefficient models that assume that all private schools are in households' choice sets. This difference is especially larger for low asset households, about a quarter of whom are estimated to be unable to choose any private school in their village, in the absence of the voucher. Arcidiacono et al. (2021) discusses the control models' estimates in full detail and presents predictions for experimental take-up of the voucher offer (and other moments) generated by simulating choices when private school tuition and fees are counterfactually set to zero.

Next, we turn to out-of-sample validation of the control model's estimates and predictions using the data from treatment markets. The validation produces three main findings. The first is that our ability-to-pay constrained model achieves relatively better out-of-sample fit, but nonetheless substantially underpredicts experimental take-up of the voucher offer. Off a base of 27% private school attendance among targeted households in the absence of the voucher program, the ability-to-pay constrained model predicts a 38 point increase. This is 10 points more than the comparable random coefficient model predicts, but the actual take-up rate among voucher winners was a 58 point increase.

The other two main findings from the validation indicate that the control market models miss key aspects of the treatment data generating process. Specifically, while both random coefficient and ability-to-pay constrained control models fit the choice patterns of ineligible and non-applicant treatment market households well, neither can rationalize that the rate of private school attendance among voucher *losers* in treated markets is much higher (15 points) than the rate among voucher applicants in control markets. This suggests that the presence of the voucher program influenced

[4]For instance, Tarozzi et al. (2014) find that micro consumer-loans substantially raised ownership and use of insecticide-treated bednets in rural India, while demand was highly elastic when households had to pay upfront. In our data, 41% of households whose child attends a government school cite "economic reasons" for their choice.

[5]Examples of papers in this set include Cameron and Heckman (2001); Keane and Wolpin (2001); Gregory (2017); Delavande and Zafar (2019).

the choices of households randomized-out from receiving an offer. The third finding is that the control models especially underpredict the rate at which voucher winners attend *low tuition* private schools. School quality as implied by the control models is positively correlated with tuition, implying voucher winners—all else equal—will prefer higher tuition private schools. This suggests the presence of school-level unobservables influencing choices that are endogenous to the program and inversely correlated with tuition.[6]

We propose two mechanisms that can potentially reconcile these findings and provide support for them from the treatment data. First, we find evidence consistent with voucher application in treatment villages inducing greater school search (see Section 5.1), which can potentially rationalize some of the control models' underprediction of voucher winners' take-up.[7] Second, private schools had strong financial incentives to enroll voucher students and, moreover, those incentives decreased with their own tuition. This is because the voucher amount, which was paid directly to private schools, was set substantially higher than the annual tuition at most private schools. This naturally raises the question of whether private schools passed through program surplus to households to incentivize enrollment. In support of this, we present evidence from a post-intervention survey suggesting school-aged siblings of voucher winners received tuition subsidies from private schools.

Since the treatment induced behavioral changes among both households and schools that were not seen in control markets, we develop a unified choice model incorporating these mechanisms (as well as an ability-to-pay constraint). We model search as a requirement that households pay a cost to reveal their match qualities at the private schools in their village. The benefit from searching for treatment market applicants is thus influenced by the voucher offer, but we assume that voucher losers face the same kind of choice environment as control market households post-search. Those that draw a sufficiently high match quality may therefore choose to attend a private school even though they have to pay tuition and fees. In addition, to model the influence of enrollment incentives on choice patterns, we include the voucher surplus—the difference between the voucher amount and a private school's tuition and fees—in voucher winners' utility (if positive). All told, these two additional mechanisms introduce three new parameters to our ability-to-pay constrained model.

We estimate the unified model on the combined control and treatment markets data. Identification of search frictions leverages the project's two-stage randomization, which in the model creates

[6]It also points away from insufficiently controlling for school-level unobservables in the control market estimation as an explanation for underpredicting take-up. We use instruments standard in the literature (e.g. Berry, Levinsohn and Pakes 1995; Hausman 1996; Nevo 2001) to address endogeneity of private school tuition and fees.

[7]This mechanism is also consistent with theory, and evidence from other settings, that information about schools is costly to acquire. See e.g. Arteaga et al. (2022); Larroucau et al. (2024); Neal and Root (2024).

exogenous variation in the return to searching for applicants across treatment and control markets. At the same time, the household-level randomization of offers within treated markets, interacted with variation in voucher surplus across private schools, identifies the effect of enrollment incentives. Note that the differences between predicted and actual take-up by treatment status (voucher winner, voucher loser in treated market, or loser in control market) are *not* first-order conditions of the maximum likelihood-based estimation routine. Thus, the fact that the results show that the unified model successfully explains choice patterns in control markets, the elevated private school attendance among voucher losers in treatment markets, and whether and where voucher winners take-up the offer increases confidence in the estimates.

We use estimates from the unified model to conduct welfare analyses under different policy counterfactuals. Our base case is a universal voucher with the same voucher value as in the experiment. This policy generates gains in consumer surplus for compliers who switch from government to private school, and also generates fiscal savings since the per-student cost of government schools is over 2.5 times the voucher value. However, a universal voucher also incurs extra fiscal costs by paying for private school attendance of those who would have attended private schools anyway. Put together, we estimate that such a program would have a marginal value of public funds (MVPF) of 1.33-3.05. We also consider an alternate policy that preserves the same voucher value, but is targeted only to asset-poor households (who were unlikely to attend private schools on their own). This policy is self-financing when the fiscal savings are large (implying an MVPF of infinity), and has an MVPF of 3 even under conservative assumptions regarding fiscal savings.

This paper makes both substantive and methodological contributions. Substantively, we show that voucher programs can have economically meaningful welfare impacts, especially when private provision is more efficient or when vouchers target those with limited ability-to-pay. Our focus on welfare, as implied by revealed preference, highlights the importance of treating schools as differentiated products.[8] This contrasts with much of the prior school choice literature, which has focused mainly on impacts on student outcomes (Epple, Romano and Urquiola, 2017). Consistent with value from expanded school choice beyond impacts on test scores, we estimate that compliers under a universal voucher would pay 17% of median annual household consumption per capita for the program but would pay about 3% to increase math value-added by a standard deviation. The

[8]Our paper thus connects with the significant line of work measuring preferences over schools and school attributes using revealed preference (e.g. Hastings, Kane and Staiger 2005; Rothstein 2006; Bayer, Ferreira and McMillan 2007; Abdulkadiroğlu et al. 2020). Kamat and Norris (2020) and Sahai (2023) are recent papers estimating welfare effects of school vouchers.

most similar paper to ours is Carneiro, Das and Reis (2022), who estimate the value of private schools in Pakistan. Our validation exercise using *experimental* variation additionally highlights the importance of search costs and ability-to-pay constraints in school choice settings, connecting with recent work on heterogeneity in valuations of school quality (Bau, 2022) and the effects of information provision (Andrabi, Das and Khwaja, 2017; Allende S.C., Gallego and Neilson, 2019).

Methodologically, our paper offers lessons for combining structural econometric models with data from randomized experiments to improve the credibility of such models.[9] Contrasting approaches are helpfully illustrated by two papers on PROGRESA: Attanasio, Meghir and Santiago (2012)—which uses treatment data to *fit* the model; and Todd and Wolpin (2006)—which uses treatment data to *validate* the model.[10] Our paper set out along the path of model validation, reflecting concerns with "structural data-mining" when treatment data are used for estimation (Schorfheide and Wolpin, 2012, 2016).[11] However, we arrived at needing the treatment data to test and identify mechanisms. Crucially, we discovered that changes in household and school behavior in response to the voucher program resulted in models estimated with only control market data being meaningfully inaccurate. Thus, our experience suggests that credible model-based welfare and policy analysis is likely to require using treatment or policy variation to estimate equilibrium models. Moreover, the kind of behavioral responses our original design failed to anticipate—from households and schools (more generally, firms)—are likely to feature in many settings and applications.

2 Background and Research Design

Our data are drawn from a randomized controlled trial conducted in 180 villages in the Indian state of Andhra Pradesh (AP). Motivated by evidence of large differences in learning levels between students attending private and government schools, the AP School Choice project was designed to study the impact of private school vouchers on student learning outcomes.[12] Villages selected for the project had to have at least one private school that agreed to participate in the voucher

[9] See Todd and Wolpin (2020) and Galiani and Pantano (2021) for recent discussions.

[10] Lagakos, Mobarak and Waugh (2023) is another example of using experimental data to directly fit a structural model. In some cases, using experimental data in estimation is combined with holding out another part of the data for validation (e.g. Duflo, Hanna and Ryan 2012; Galiani, Murphy and Pantano 2015).

[11] Our decision to also pre-commit to predictions and blind estimation to the treatment data is similar to Pathak and Shi (2021), which validates structural school choice models fit prior to a policy change in Boston. However, two distinguishing elements of our setting are: 1) endogenous tuition and fees charged by private schools; and 2) unobserved heterogeneity in households' choice sets arising from ability-to-pay constraints (and search costs).

[12] This learning gap is reflected also in the project data. Table A1 shows that at baseline the average private school student scored three fifths of a standard deviation higher in math than the average government school student.

program. Across project villages at baseline (2008), more than one of every two primary school students (57%) attended a private school.[13]

The program was targeted to students likely to otherwise attend government schools. Students randomized into treatment status were offered a voucher covering the costs of tuition and required expenses (e.g. books and uniforms) at government-recognized, participating private schools in their village for the duration of primary schooling. At the average private school in the project, tuition and fees were otherwise about Rs. 1,900 per year in 2008 (Table A2)—equivalent to nearly 8% of median annual consumption per capita.[14] Expenses for transportation, however, were not covered by the voucher and, unlike government schools, private schools do not provide free mid-day meals. Beyond costs of attendance, the bundles of characteristics associated with private and government school also differ in notable ways. For instance, private schools have less qualified teachers; but they have lower rates of multi-grade teaching and teacher absence (Table A2). Private schools were also much more likely to feature English as the medium of instruction, and to allocate class time to teaching Hindi (the national language), whereas government school instruction was entirely in the local language, Telugu.

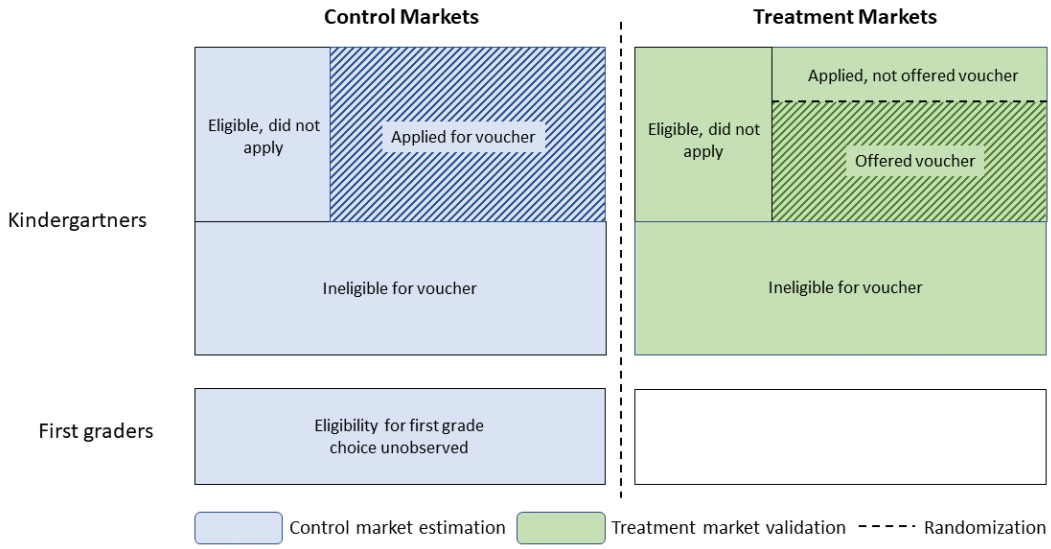
Participation in the project at the school level was voluntary, but participating private schools were not allowed to screen or selectively admit voucher students. The design stipulated that lotteries would be held to allocate places in oversubscribed schools, but in practice this proved to not be needed. The annual voucher value was set at around the 90th percentile of the tuition fee distribution among private schools. For each voucher recipient verifiably enrolled, Rs. 2,600 was paid up front and directly to schools' bank accounts.

2.1 Research Design

An important feature of the AP School Choice project is its two-stage randomization: At baseline, parents of eligible students were invited to apply for the program with the knowledge that the voucher would be allocated by lottery and that applying would not guarantee receipt. After eliciting interest from eligible households, the project first randomized villages into 90 treatment and and 90 control markets. Applicant households in treatment villages were then randomized into or out

[13]Calculated using household survey data in the AP project districts (ASER, 2018). This high private-school market share is conditional on the village having at least one private school (which was a requirement to be in the study). The unconditional private school market share among primary school-going students in AP in 2008 was around 33% (ASER, 2018). The study was conducted from 2008-12, and is set in the erstwhile undivided state of Andhra Pradesh, which was later divided into the two states of Andhra Pradesh and Telangana in 2014.

[14]Median household expenditure per capita was about Rs. 24,000 per the 2011-12 India Human Development survey in comparable rural villages (with a private school) of Andhra Pradesh.



Notes: Figure visually represents the paper's research design, which 1) estimates models of primary school choice on first grade choices of kindergartners and first graders in control markets; and 2) validates the models using data from treatment markets. Eligiblity for AP voucher (for first grade choice) determined by attending Anganwadi (government-run pre-school). Vertical dashed line represents first stage market-level randomization; horizontal dashed line dividing treatmet market applicants represents household-level randomization. Shading with upward-sloping diagonal lines represents experimental validation sample.

Figure 1: Control Market Sample and Treatment Market Validation for First Grade Choice

of the voucher treatment group in the second stage.[15]

Our research design leverages the market-level randomization to first estimate school choice models on data from control markets before validating the models against choice patterns in treatment markets. The AP project experiment was conducted in parallel on two cohorts of students: a younger cohort, who had yet to enter primary schooling at time of baseline and who we term kindergartners, and an older cohort of first graders already attending a primary school at baseline. Our empirical models focus on households' first grade primary school choice. This has two consequences for our research design: 1) in fitting the models, we use kindergartners' choices subsequent to baseline and first graders' *retrospective* choices; and 2) the experimental validation by necessity focuses on the choices of kindergartner voucher winners.

Figure 1 presents a visual representation of the design. Shown by the light blue shading, several alternative empirical models of primary school choice, detailed in the next section, are fit to only the data from control markets. Note that the control markets data are purely observational; we observe school choices made by households, characteristics of those households, and attributes of the school options (including the tuition and fees). As shown in the figure, the kindergarten control

[15]This double randomization design facilitated estimating spillover effects on non-participants in the program, and also provided a pure control group (in the control villages) that would be unaffected by such spillovers. See Muralidharan and Sundararaman (2015) for details.

markets sample contains several subgroups: those who were ineligible for the voucher; those who who were eligible and did not apply; and those who were eligible and applied but were randomized-out at the village-level. Eligibility was determined by whether the child attended a government preschool (Anganwadi) or not, which was used as a proxy means test, since students from better off households were more likely to be attending a private preschool.[16] Table A1 compares the eligible kindergartners with the private- and government-attending first grade populations, showing that students eligible for the program are more similar in background demographics and socioeconomic status to government school students. Private school students are more likely to have parents who both completed primary school; more likely to have a parent who completed secondary school; and more likely to live in a pucca (brick or stone) house, have a water facility in the home, and to have a household toilet (Table A1).

The out-of-sample validation step of our research design uses the data from treatment markets, which include parallel subgroups except that the household-level randomization further splits kindergartner applicants into voucher winners and voucher losers. For kindergartners in treatment villages that did not receive a voucher (e.g. those ineligible), we evaluate how the models fit out-of-sample under the assumption that the program did not impact their choices. For those that received a voucher offer, we evaluate the models based on their predictions for choice patterns—e.g. what share would take-up the voucher offer. We do this by counterfactually setting tuition and fees at government-recognized private schools to zero in the models. This experimental validation step is visually represented by the boxes overlaid with upward-sloping diagonal lines.

2.2 Treatment Data

Data collection in treatment markets mirrored the data collection in control markets. Households were surveyed at baseline, while schools were surveyed beginning the first year of the program. We process the control and treatment markets data in the exact same way (described in greater detail in Arcidiacono et al. 2021) to produce a cleaned dataset that connects each student to a village-specific set of primary schools. GPS locations were collected, facilitating the mapping of travel distances between households and schools in their village. The data contain many observable characteristics of students and household (e.g. whether belongs to a historically-disadvantaged

[16]Combining the subsamples in estimation presents three practical challenges, which are discussed fully in Arcidiacono et al. 2021): 1. the trial's sampling design; 2. attrition of kindergartners from the sample; and 3. which first graders attended an Anganwadi, and hence would have been eligible as a kindergartner, is unobserved. We construct weights to deal with the first two issues and handle the third by treating it as a latent type in estimation.

scheduled caste) as well as numerous characteristics of primary schools in the village (e.g. whether English is the medium of instruction). These variables are summarized in Tables A1 and A2.

In processing the data for the voucher winners, we restrict the sample to students for whom a specific choice of school in their village is recorded in the tracking data. This is to mirror the restrictions made for the control sample and to match the fact that the models restrict choices to schools in the student's baseline village. This yields a total sample of 629 kindergartner students who were randomly offered a voucher in treatment villages. An issue for the experimental validation to come, however, is reduced sample attrition of voucher winners, which we address in Section 4.

The randomization and symmetric data collection (and processing) imply that there should be baseline balance on average between 1) control and treatment market schools; 2) control and treatment market households who were did not (or could not) apply for the voucher; and 3) control and treatment market households who did apply for the voucher.[17] Consistent with this, Table A1 shows limited statistical differences between control and treatment market subgroups of households. Table A2 likewise shows school-level balance along most dimensions. However, treatment market private schools' are about 13% more expensive.[18]

3 Control Models and Results

In this section, we describe our empirical models of household school choice that were estimated using only data from the control villages. In our choice models, we treat households, which consist of at least one primary school aged child, as unitary decision makers. As private schools charge tuition and fees, households must weigh the expected benefits of private school attendance against foregone consumption. Such benefits potentially include a more attractive combination of school amenities as well as human capital gains.

We compare the estimates and predictions for two classes of choice models. In the first, we explicitly model the influence of an unobserved ability-to-pay constraint on choice. In relaxing this constraint, a private school voucher thereby potentially generates welfare benefits by expanding households' choice sets. We compare this model class, which places structure on how observed measures of household wealth influence choices, with random coefficient demand models that are similar to models of school choice that have been applied in other contexts.[19]

[17] There should also be balance between treatment market applicant households randomly offered and randomly not offered a voucher and there is.

[18] The treatment-control difference in average tuition is robust to controls.

[19] Arcidiacono et al. (2021) also describes and presents predictions for a model that assumes all choices are available

3.1 Ability-to-Pay Constrained Choice

In selecting a primary school, households weigh the utility of the school alternatives that belong to their village.[20] This set is denoted by $\mathcal{V}_i$ for household i. However, the tuition and fees may exceed the household's ability-to-pay. This is captured in the model through a constraint on their choice problem:

$$\max_{j \in \mathcal{V}_i} U_{ij} \geq U_{ij'} \quad \forall j' \in \mathcal{V}_i \text{ where } p_j, p_{j'} \leq \omega_i \tag{1}$$

For any school, j, the household's consumption and tuition and fees, denoted p_j, must not exceed the household's ability-to-pay, which we denote by ω_i. For government schools, p_j is zero (or nearly so). The ability-to-pay constraint represents the combination of a household's income and any liquid wealth, such as accumulated savings, with their ability to borrow against future income to finance private schooling. This "reduced-form" constraint also captures the possibility of subsistence constraints or that households may be unable to commit to the schedule of private school tuition and fees due to uncertain income streams.

Households rank the available schooling alternatives according to an indirect utility function. Letting α represent household i's marginal utility of consumption, the indirect utility to household i of school choice j can be written as:

$$U_{ij} = \alpha(y_i - p_j) + X_j'\beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \xi_j + \epsilon_{ij} \tag{2}$$

D_{ij} is the distance between school j and household i's home, while X_j represents school characteristics. $Closest_{ij}$ allows that the closest school, if a government school, is especially salient. In estimation, we include in X_j whether a school is government or private, is government recognized (if private), is English medium, offers Hindi classes, is connected to a secondary school, and three indices respectively capturing the quality of facilities, of teachers, and the characteristics of teachers. Also contained in X_j is school j's value-added in math, which we estimate from the panel of student test scores.[21] ξ_j represents an index of commonly-valued amenities of school j unobserved to the econometrician and likely correlated with tuition. ϵ_{ij} is assumed to follow a Type 1 extreme

and groups households into clusters based on observables, allowing preferences to be cluster-specific. We do not discuss this "clustered multinomial logit" demand model here given its predictions and estimates are qualitatively the same as the random coefficient model.

[20] Primary schooling is compulsory in this setting, so we do not model the choice of whether to send the child to school or not.

[21] Appendix B of Arcidiacono et al. (2021) details the value-added estimation. We include indicators for missing distance, missing value-added, and imputation of tuition and fees.

value distribution.

We subscript the parameters in equation (2) by i to denote their dependence on observed household characteristics, W_i:

$$\begin{pmatrix} \beta_i \\ \gamma_i \end{pmatrix} = \begin{pmatrix} \beta_1 \\ \gamma_1 \end{pmatrix} + \begin{pmatrix} \beta_2 \\ \gamma_2 \end{pmatrix} W_i$$

The household characteristics in W_i mediate the valuation households place on school amenities, capturing systematic heterogeneity across households in willingness-to-pay. W_i includes AP voucher program eligibility status and indicators for gender, whether belongs to a scheduled caste, is Muslim, whether an older sibling attends government school, whether both parents completed primary school, and whether one parent completed secondary school.[22] Note we do not include assets in W_i; this information enters the model via the ability-to-pay constraint.

3.1.1 Instrumenting for Private School Tuition and Fees

A first empirical challenge for estimating this model (which applies equally to the random coefficient model discussed next) on the control markets data is that ξ_j is unobserved. We implement a control function approach to address the endogeneity of private school tuition and fees (Petrin and Train, 2010). This strategy regresses tuition and fees on school characteristics and a set of instruments in a first stage:

$$p_j = X_j'\Gamma + f(Z_j) + \mu_j \tag{3}$$

where X_j are observed school characteristics (including the estimated value-added), Z_j are instruments, and $E[\xi_j \mu_j] > 0$.The utility function we then ultimately take to the data is given by:

$$U_{ij} = -\alpha p_j + X_j'\beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \kappa\hat{\mu}_j + e_j + \epsilon_{ij}$$

where $\hat{\mu}_j$ is the first stage residual for private school j and e_j is a normally-distributed random effect; both terms are zero for government schools.

Our baseline specification uses two instruments. First, we use a summary measure of each private school's location in "product space" (Berry, Levinsohn and Pakes, 1995). We do this using factor analysis applied to totals of characteristics of *other* schools in the village for each private

[22]Our specifications do not include all possible interactions of household and school characteristics. The exact interactions we do include are summarized in Table A22 of Arcidiacono et al. (2021).

school, e.g. the number of other English-medium schools. The second instrument uses the spatial environment to isolate exogenous cost differences across private schools (Hausman, 1996; Nevo, 2001). We construct the predicted tuition for each private school based on the average tuition chosen by similar private schools that are located in *other* villages.[23] Arcidiacono et al. (2021) provides additional details on construction of the instruments. First stage estimates are presented in Table A8.

3.1.2 Identifying and Estimating Ability-to-Pay

The second empirical challenge for estimating the choice problem described by equation (1) is that households' ability-to-pay, ω_i, is inherently not contained in the data. This introduces unobserved heterogeneity across households in choice sets. Mis-specifying households' choice of school as unconstrained is liable to bias estimates of willingness-to-pay and underestimate the gains of a voucher.

We specify latent ability-to-pay as a function of observed household wealth factors, given by:

$$\ln \omega_i = I_i' \lambda + \upsilon_i \tag{4}$$

In this equation, the household's log ability-to-pay at the time of choosing a primary school depends on the wealth factors, I_i, and unobservable household-specific υ_i. We assume that υ is distributed normally, with variance σ, and independent of the choice shocks. Our baseline model specification includes the household asset factor, an indicator for eligibility for the voucher program, and household size in I_i.

As this feature of the model is new, we briefly discuss estimation via maximum likelihood. Interested readers are referred to Arcidiacono et al. (2021) for more details. The basic insight is to recognize that each i can fall into one of a finite number of possible choice sets. Let j_i^* index schools in i's village in terms of ascending tuition and fees (such that J_i^* is the most expensive). Then denote by $\phi_{ij_i^*}$ the probability that household i is in choice state of being able to afford at

[23] In implementation, we match private schools within medium of instruction and focus on other schools not in nearby villages. This is to minimize the confounding influence of spatially-correlated demand shocks. As an alternative to the predicted tuition instrument, we also estimate models that include the product space IV and a cost index instrument constructed from private schools' reported costs.

most: $p_{j_i^*} \leq \omega_i \leq p_{j_i^*+1}$. We can write this as:

$$\phi_{ij_i^*} = \Phi\left(\frac{\ln p_{j_i^*+1} - I_i'\lambda}{\sigma}\right) - \Phi\left(\frac{\ln p_{j_i^*} - I_i'\lambda}{\sigma}\right)$$

where the state probability is a difference between points on the normal CDF that depend on data (tuitions and I_i) and parameters (λ and σ). $\Phi\left(\frac{\ln p_1 - I_i'\lambda}{\sigma}\right)$ is the probability of not being able to choose *any* private school in their village.[24] Combining logit expressions for choice probabilities with the state probabilities allows us to form a likelihood for each household.

3.2 Random Coefficient

We compare the latent ability-to-pay model with random coefficient models similar to classic demand estimation applications (e.g. Berry, Levinsohn and Pakes 1995; Nevo 2001; Petrin 2002) and the models of school choice in Neilson (2013) and Carneiro, Das and Reis (2022). In this class of models, the underlying choice problem is unconstrained—households are able to choose from any primary school in their village:

$$\max_{j \in \mathcal{V}_i} U_{ij} \geq U_{ij'} \tag{5}$$

where U_{ij} again represents i's indirect utility from attending school j.

The indirect utility in the random coefficient model is given by:

$$\begin{aligned} U_{ij} &= -\alpha_i p_j + X_j'\beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \xi_j + \epsilon_{ij} \\ &= -\alpha_i p_j + X_j'\beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \kappa\hat{\mu}_j + e_j + \epsilon_{ij} \end{aligned} \tag{6}$$

where the substitution reflects the control function strategy for addressing unobserved ξ_j, which is applied in the same way. While similar to the ability-to-pay constrained model, this indirect utility differs in two ways: First, note that the function allows for heterogeneity across households in their sensitivity to higher tuition and fees, reflected in the indexing by i. Specifically, we allow α_i to depend on household asset levels (e.g. whether the household own up to six possible assets) and household size. Second, the random coeffiecient demand model accommodates greater flexibility in how households value school characteristics.

Like the ability-to-pay constrained model, the random coefficient model specifies a parametric relationship between observed household characteristics, W_i, and preferences over non-tuition school

[24] The probability the household can choose from *all* private schools is given by $1 - \Phi(\frac{\ln p_{J_i^*} - I_i'\lambda}{\sigma})$.

amenities. However, the random coefficient model includes an additional stochastic component on household preferences for private schooling. Letting β_i^P indicate the marginal utility to household i of attending private school, this parameter can be expressed as:

$$\beta_i^P = \beta_1^P + \beta_2^P W_i + \nu_i \tag{7}$$

ν_i is an unobserved, continuous type that follows a mean-zero normal distribution. This additional stochastic term captures unobserved heterogeneity in preferences for private schooling across households

3.3 Control Estimation and Results

Per our research design, we estimate the empirical models above using only data from the control markets. The estimation details and results are summarized here, with full elaboration provided in Arcidiacono et al. (2021).

Estimation on the control data pools choices from several subgroups of students, shown with the light blue shading in Figure 1: kindergartners who were eligible (by virtue of attending an Anganwadi at baseline) and who applied for the voucher program; eligible kindergartners who did not apply; ineligible kindergartners; and first graders (whose retrospective choice of primary school we use in estimation). Though the model specifications allow for preferences (and ability-to-pay, in the constrained model case) to depend on AP voucher eligibility, we do not model application status.[25] However, since eligibility status is unknown for this older cohort, we model latent eligibility of these students (and use the EM algorithm in estimation). We treat the private school random effects, which adjust the variance of the private school choice shocks, as iid school- and household-specific and construct household weights to account for the project's sampling design and for attrition of kindergartners.

The full set of control model parameter estimates are reported in Tables A24 and A25 of Arcidiacono et al. (2021); selected estimates of preferred specifications focused on for validation are reported in Table A3. The ability-to-pay constrained model estimates imply that around 13% of applicant households (and 24% of low asset households) are unable to otherwise choose any private school in their village (Table 9). The ability-to-pay constrained control model accordingly ascribes greater utility from private schooling—and from school attributes associated with private

[25] As justification for this, conditional on observables, application status is not a statistically significant predictor of private school attendance in the control data.

schools—than the random coefficient model. This difference translates into important differences in welfare effects and, as we turn to in the next section, for voucher take-up.

4 Treatment Validation

We now turn to evaluating the empirical models' out-of-sample performance using the held-out treatment markets data. The treatment data allow for two kinds of validation: non-experimental and experimental. These can be understood visually from Figure 1. In the treatment data, we have several subgroups of kindergarten households who did not receive a voucher offer: those who were eligible and applied, but randomized out at the household level; those eligible who did not apply; and the ineligible. We can therefore ask how the models estimated on the control data do in explaining the choice patterns of households in treatment markets also in the "control" condition.

The primary focus of our design, however, is on validation out-of-sample against choice patterns under the voucher experiment. This experimental validation is represented by the boxes in Figure 1 filled with diagonal lines: using the empirical models, we generate predictions for the voucher take-up of kindergartner applicants. We do this by setting tuition and fees at participating private schools to zero and simulating choices. This allows us to compare model-based "treatment" moments (pre-committed to in Arcidiacono et al. 2021) with analogous moments calculated directly from the treatment group. The subsections below present the findings from these different out-of-sample validations of our empirical models in turn. Before doing so, however, we provide information about how we bring the control model to the treatment data.

4.1 Using Control Estimates on Treatment Data

To apply the control models to the treatment data, we need to construct certain variables in consistent ways. For latent factor variables, such as the asset index, we impute their values in the treatment data using the relationships between characteristics and factors in the control markets. Similarly, we use the first stage for tuition and fee endogeneity estimated on the control data to impute the residuals for treatment market private schools. Consistent with the treatment-control difference in private school tuition, shown in Table A2, this imputes a higher average unobserved quality among treatment market private schools.

But data on voucher winners yields an additional layer of complexity because of (i) how winning the voucher affected attrition and (ii) the inability of some students to use voucher even when they

intended to do so. We describe these issues next.

4.1.1 Attrition

The attrition rate, calculated as the share of households at baseline with valid tracking data, is noticeably smaller for households offered a voucher (11%) than it is for control market applicants (19%). This suggests that the voucher offer, by attracting students to private schools in their village, induced households to be more likely to stay in the final sample. We thus adjust model predictions and estimates based on the treatment data to account for selective attrition. To do this, we first solve for the number of households that would have attrited from the final offered student sample *in the absence of the voucher offer*. The calculation assumes that the attrition rates of applicants between treatment and control markets would be the same in the absence of the offer and comes to 70 of the 574 students. We then assume that the 70 students who otherwise would have attrited also belong to the subgroup of students who actually took-up the voucher offer.

Under this assumption, we use the calculation of excess attriters in two ways in the analysis. First, we assign weights to the students who actually used the voucher such that they effectively represent 70 fewer students. We also adjust the weights to account for differences in the probability of attrition between those students (as a function of observed characteristics). These weights are applied when using offered households' choice data. Second, we correct model predictions for voucher take-up for selective attrition by adding 70 students to the number of voucher users predicted by the control models.

4.1.2 Coding Voucher Take-Up

The experimental validation concerns the degree to which the models accurately predict the decisions of kindergartner students in treatment markets who were randomly offered a voucher. How "voucher use" is coded is thus a key input to the exercise, which we now discuss.

The project team collected information about voucher use as reflected in payments to private schools as well as reasons in cases on non-use. About 66% of the 629 students in the cleaned sample who were offered a voucher actually used it. Note that this number closely matches the figure stated in Muralidharan and Sundararaman (2015). However, this tabulation is not the correct one for the purposes of the experimental validation, which takes the data and sets tuition to zero in the models estimated on control markets to predict take-up. The predictions correspond to choices—as they would appear in the tracking data—with a voucher provided there were no

extenuating circumstances.

To code voucher take-up in this manner, we combine information from tracking and from the project team. Table A4 summarizes the coding. Our starting place is the majority of students (416) labeled as accepting the offer and who attend a private school in tracking data. To this group, whose use was reflected in voucher payments, we add as intended users 69 students who attend a private school in tracking data. As Table A4 shows, most of these cases are students who later "dropped out" of the voucher program or who were ex-post ineligible due to gaining admission to a private school prior to learning their voucher outcome. We further code as intended users 21 students who tried to use the voucher, but were unable by virtue of being too young (irrespective of where tracking data show them attending school).[26] For the subgroup of eventual drop outs, our analyses to come assume, consistent with the data patterns, that their tracking private school is where they initially chose to use the voucher; our analyses are agnostic about precisely where students would've used the voucher in the other cases.

Importantly, in eight treatment market villages, no students randomized-in to receive an offer actually used a voucher due to non-compliance by private schools in those villages.[27] For purposes of the experimental validation, we remove these non-complying villages from the sample entirely.[28] This leaves a sample of 574 households who were randomly offered a voucher in treatment villages. Of these, 489 (85%) intended to take-up the voucher offer. Later on, we look at choice patterns of households in the non-compliant treatment villages to test mechanisms that could explain under-prediction of take-up.

4.1.3 Data Patterns in Treatment Villages

We now summarize key data patterns in the treatment markets by examining how the choices of different groups varies across treatment and control villages, with the results shown in Table 1. First graders have similar rates of private school attendance as do ineligible kindergartners, the latter because virtually all ineligibles attend private school.

Kindergartners who are eligible for the voucher but do not apply are six percentage points

[26] There is also one student with the extenuating circumstance of "waiting list not used" that we code as intending to use.

[27] This can be clearly seen in Table A4, where excluding these "flagged" villages removes all of the offered students coded as "school rejected" from the sample. Several private schools in otherwise compliant treatment villages also reneged on participating. This was in part because of the project requirement that all students in these schools should take independent learning assessments (this concern was raised by other private schools too, which is why the assessments used in Muralidharan and Sundararaman (2015) were conducted outside school). We therefore do not set tuition and fees to 0 at these specific schools when generating model predictions.

[28] We also flag and remove one additional treatment village where take-up was not zero but was abnormally low.

Table 1: Private Schooling and Tuition and fees by Subgroup

	Attend Private				Tuition|Private			
	Control		Treat		Control		Treat	
	N	Mean	N	Mean	N	Mean	N	Mean
First graders	2766	0.57	2648	0.58	482	1.71	483	1.82
Ineligible for voucher	374	0.99	413	0.99	370	1.79	407	1.87
Eligible non-applicants	124	0.22	134	0.16	27	1.58	22	1.95
Applicants not offered voucher	987	0.32	299	0.43	316	1.85	130	2.12
Voucher winners	0	.	574	0.85	0	.	437	2.10

Notes: Table reports average private school attendance and tuition given private attendance by subgroup across control and treatment markets. Attend private for voucher winners refers to voucher use; N excludes 55 winners residing in non-complying villages who were unable to use the voucher. Note that conditioning on private attendance for Tuition|Private excludes voucher users who do not attend a private school in the tracking data.

less likely to attend in treatment villages. This result, coupled with the patterns for first graders would suggest that private schools may be slightly less attractive in treatment villages. However, voucher losers (which includes all applicants in control villages) are substantially more likely to attend private school in treatment villages. This points towards one of the main findings from the out-of-sample validation, detailed in the next section.

The second set of columns shows the average tuition among private school attendees for each of the subgroups. For each subgroup of attendees, the average tuition is higher in treatment villages, consistent with the school-level summaries presented in Table A2. But what is especially striking is the average tuition charged in schools attended by voucher winners. Namely, it is remarkably similar to that paid by applicants in treatment villages who were not offered a voucher, despite the free tuition faced by the voucher winners, suggesting that voucher winners did not on average upgrade school quality as measured by the tuition charged. This data pattern likewise figures prominently in the assessment of the control models' fitness.

4.2 Non-Experimental Validation

We are now in a position examine how the control models fit the treatment data. We begin by examining how well the empirical models fit the choice patterns of treatment market households who do not receive a voucher. To do so, we take the cleaned data from treatment markets for ineligible households, eligible non-applicants, and applicants who did not win a voucher and directly apply the control model estimates, which allows us to compute predictions for private school attendance.

Table 2 shows the results of validating the models out-of-sample. Both the random coefficient model and the ability-to-pay constrained model match well the private school attendance rates for ineligibles and eligible households who didn't apply for the voucher program. However, both models

Table 2: Out-of-Sample Validation: Treatment Market Predictions

	Attend Private Model			Tuition|Private Model		
	Data	RC	CC	Data	RC	CC
Ineligible for voucher	0.99	0.99	0.98	1.87	1.96	2.02
Eligible non-applicants	0.16	0.19	0.17	1.95	2.00	2.04
Voucher losers	0.43	0.29	0.28	2.12	1.98	2.00
Voucher winners	0.83	0.58	0.67	2.11	2.46	2.48

Notes: Table reports private school attendance and average tuition given private attendance among treatment market kindergartner subgroups in the treatment market data (Data) and as predicted by the control random coefficient model (RC) and control ability-to-pay constrained model (RC). Numbers for voucher winners are weighted to adjust for the group's lowered attrition.

significantly underpredict private school attendance of voucher losers by nearly 15 percentage points.

To further examine the fit of our models to these groups, we formulate the question of misspecification as hypotheses tests. To do so, we begin by fixing the indirect utility for each option j in treatment models to that predicted from the control model estimation (plus a T1EV choice shock). For empirical model m:

$$\hat{u}_{ij}^m = -\hat{\alpha}_i^m p_j + X_j' \hat{\beta}_i^m + \hat{\gamma}_i^m \ln D_{ij} + \hat{\delta} Closest_{ij} + \hat{\xi}_j^m$$

We then estimate an auxiliary model for each control model on kindergartners in treatment villages who do not win a voucher where their indirect utility at j (less an idiosyncratic choice shock) is specified as:

$$U_{ij}^m = \hat{u}_{ij}^m + \pi_T^m Private_j + \pi_L^m \mathbf{1}[VoucherLoser_i] \times Private_j + \tau^m p_j + \epsilon_{ij}$$

This specification allows us to see whether the overall private utility from the control model (which is embedded in $\hat{u}_{ij}^m$) is different for treatment village ineligibles and eligible non-applicants (given by π_C^m) and for voucher losers (given by $\pi_C^m + \pi_L^m$) as well as whether the price coefficient is different in treatment villages. Under the assumption the model is true, $\hat{u}_{ij}^m$ controls for all observed and unobserved qualities of school j. These auxiliary models thus ask whether the control models do a good job predicting which private school these students attend (as a function of tuition) as well as whether they attend private school.

The results are presented in Table 3. Columns (1) and (3) report goodness-of-fit summaries in the form of AIC stats for the random coefficient and ability-to-pay constrained control models, respectively. The fit of the random coefficient model to the choices of treatment market households

Table 3: Non-Experimental Validation: Hypotheses Tests

		RC		CC
	(1)	(2)	(3)	(4)
Private school		-0.12		0.07
		(0.50)		(0.09)
Private school × Voucher loser		2.21		2.05
		(0.56)		(0.30)
Tuition and fees (1000s of Rs.)		0.00		-0.11
		(0.10)		(0.09)
AIC	2,399	2,260	2,411	2,265

Notes: Table reports hypothesis tests of model mis-specification that examine predictive power of private voucher school constant and tuition and fees for choices of treatment market kindergartners not offered a voucher conditional on the indirect utility of the alternative implied by the control random coefficient model estimates (RC) and control ability-to-pay constrained model estimates (CC). $N = 846$ kindergartner treatment market households not offered a voucher. Excluded group is ineligible and eligible kindergartners who did not apply for AP voucher. Standard errors reported in parentheses.

who do not win a voucher is marginally better. Columns (2) and (3) show that, consistent with Table 2, the estimates of the private dummy are not different from zero for ineligibles or for non-applicants, but it is significantly positive for voucher losers irrespective of the control model. At the same time, the coefficient on price is small and insignificant, suggesting that the control estimates are providing good estimates of the tuition gradient for this sample.

Overall, the non-experimental validation suggests that both control models provide a good fit for the treatment market data for households who did not apply for a voucher or were ineligible for one, but do not fit as well for applicants who were randomized out from receiving a voucher offer at the household-level. That both models miss for this group raises the question, which we turn to later, of whether and how the voucher intervention nonetheless may have influenced this group's choices.

4.3 Experimental Validation

This subsection presents the findings from the experimental validation of the control models. We first focus on predictions for voucher take-up before again exploring sources of mis-specification in a hypothesis testing framework.

4.3.1 Predicted versus Actual Voucher Take-up

We first examine how the control models do at predicting private school attendance of voucher winners, without and with accounting for the effect of the offer on attrition. Table 4 presents actual

and model-predicted take-up. The first column reports attendance at private schools by applicant kindergartners in control markets. Overall, 27% of applicants in control markets choose to attend a government-recognized private school. Columns (3) and (4) report model predictions for voucher take-up according to the random coefficient and ability-to-pay constrained model, respectively.[29]

Table 4: Experimental Validation: Take-up of Voucher Offer

	Data		$\hat{U}se$		$\hat{U}se$ adj.	
	Control	Treat	RC	CC	RC	CC
	(1)	(2)	(3)	(4)	(5)	(6)
Overall	0.27	0.85	0.50	0.60	0.56	0.65
Female	0.24	0.86	0.50	0.59	0.55	0.64
Muslim	0.47	0.98	0.70	0.79	0.80	0.86
Lower caste	0.18	0.77	0.42	0.53	0.47	0.57
Older sibling in gov't school	0.14	0.79	0.33	0.43	0.40	0.49
Both parents completed primary school	0.41	0.88	0.64	0.70	0.69	0.74
≥ 1 parent completed secondary	0.46	0.76	0.67	0.74	0.71	0.77
Both parents laborers	0.21	0.77	0.44	0.54	0.49	0.59
Asset level < 3	0.21	0.85	0.47	0.57	0.54	0.63
Asset level $= 3$	0.29	0.85	0.51	0.61	0.56	0.65
Asset level $= 4$	0.25	0.85	0.50	0.60	0.56	0.65
Asset level > 4	0.38	0.89	0.59	0.66	0.64	0.70

Notes: Table presents average private school attendance by applicants in control markets (Control), average voucher take-up by treatment market applicants (Treat), and average voucher take-up by treatment market applicants as predicted by random coefficient (RC) and ability-to-pay constrained control models (CC) by subgroup. Columns (5) and (6) adjust the predictions upward for the reduction in winners' attrition due to the voucher offer. Predictions correspond to baseline specification described in the text and detailed in Arcidiacono et al. (2021).

The random coefficient model predicts that private school attendance under the voucher will increase by 23 points to 50%. The ability-to-pay constrained model predicts that private school attendance will more than double, increasing another 10 points to 60% of those offered. Across households, this gap is pretty uniform, though it is only six points among those where both parents completed primary school. Among households where a parent completed secondary school, the ability-to-pay constrained model underpredicts by only 2 points. Columns (5) and (6) of Table 4 then adjust the model predictions for selective attrition—that the voucher offer induced fewer students to attrit. This correction raises the predictions to 56% and 65% take-up of the voucher offer, respectively.

Column (2) of Table 4 reports take-up of the voucher in the treatment markets—what actually

[29]Note these predictions do not exactly match those reported in Arcidiacono et al. (2021). This is because Table 4 re-computes the predictions on the students offered the voucher in the treatment markets, whereas the original predictions were computed for applicants in control markets. These predictions are thus adjusted for minor treatment-control differences in observables. They also account for non-participation in the program by some schools. There is a second reason the predictions are different (and generally a little lower), which is that Arcidiacono et al. (2021) simulated take-up allowing households to use a voucher at unrecognized private schools. In practice, the voucher could only be used at government-recognized private schools.

happened. As reported earlier, 85% of applicants randomly offered the voucher used it (or intended to use it) to attend a participating private school. Compared with the random coefficient model prediction, this represents a gap of 29 points. The ability-to-pay constrained model's prediction was also too low, but by 9 fewer points. The subgroups comparisons show that the models performed especially badly at predicting take-up of students with an older sibling in government school. The ability-to-pay constrained model was off by 30 points for this group. This reflects that the control market estimates of both models assign a significant disutility to attending a private school for this group of students. The data-prediction gap in the case of the constrained model is 15 points among households without an older sibling in government schools. While these comparisons discussed pertain only to the baseline specifications of the control models, the gaps highlighted are robust across the alternative control model specifications estimated (e.g. using the alternative IVs).

Table A5 compares model predictions for elasticities of private schooling with respect to the voucher offer with those computed using the experimental variation. These comparisons likewise reveal that the models generally underpredict—albeit the ability-to-pay constrained less so—but also reveal a data pattern that the constrained model captures better. The table shows that the voucher elasticity is highest for the low asset households and lowest for the high asset households. Both models match this, but the difference in the elasticity between the low and high asset households is matched more closely by the ability-to-pay constrained model. Finally, Table A6 compares effects of the voucher offer on characteristics of households' chosen schools in terms of treatment-control differences. As expected, the offer raised tuition at chosen schools (by about Rs. 1000 on average) and increased attendance at an English medium school by 13 points. It also increased attendance at a school offering Hindi by 33 points. Both models underpredict the effect on Hindi, but produce similar ITT effects as the experiment on English and tuition. This is despite underpredicting private school attendance significantly and thus suggests the models overvalue English and overpredict use of the voucher at high tuition schools.

4.3.2 Hypothesis Tests

This section examines model fit and mis-specification by estimating auxiliary models on the choices of the treatment market applicants randomly offered a voucher while controlling for the indirect utility predicted by the control models. Specifically, we estimate models of the following form:

$$U_{ij}^{m} = \hat{u}_{ij}^{m} + \hat{\alpha}_{i}^{m} p_j + \pi_V^m PrivateVoucher_j + \epsilon_{ij} \tag{8}$$

for each empirical model m estimated on the control sample. $\hat{u}_{ij}^m + \hat{\alpha}_i^m p_j$ is treated household i's predicted indirect utility from choice j, according to the estimates from control model m. Like before, if control model m accurately captures treated students' take-up of the voucher offer (i.e. their preferences over voucher-eligibile private schools), we expect that $\pi_V^m = 0$. For households we code as intending to use the voucher but who were not able to actually use it, estimation matches their intended use with their probability of attendance at any government-recognized private school in their village with the voucher.

Table 5: Experimental Validation: Hypothesis Tests Comparing Random Coefficient and Ability-to-pay Constrained Models

		RC			CC	
	(1)	(2)	(3)	(4)	(5)	(6)
Private voucher school		4.72	7.49		2.60	5.28
		(0.30)	(0.46)		(0.22)	(0.40)
Tuition and fees (@ voucher school)			-1.32			-1.32
			(0.17)			(0.16)
$\hat{\text{Use}}$	0.56	0.84	0.84	0.65	0.84	0.84
AIC	1,496	1,198	1,135	1,400	1,235	1,164

Notes: Table reports hypothesis tests of model mis-specification that examine predictive power of private voucher school constant and tuition and fees for voucher winners' choices conditional on the indirect utility of the alternative implied by the control random coefficient model estimates (RC) and control ability-to-pay constrained model estimates (CC). $N = 574$ kindergartner treatment market households offered a voucher (not in non-complying treatment villages). Standard errors reported in parentheses.

Columns (1) and (6) of Table 5 report measures of goodness-of-fit to offered students' choices under the voucher; the constrained model achieves a lower AIC. Columns (2) and (7) then insert an intercept for private (voucher-eligible) schools, as in the hypothesis testing framework outlined above. This added provides an alternative way to quantify underprediction of take-up between control models: the coefficient on the intercept is 40% larger in the random coefficient model.

Columns (5) and (9) of Table 5 simultaneously estimate an intercept for voucher-eligible private schools and a "slope" on tuition at those schools:

$$U_{ij}^m = \hat{u}_{ij}^m + \hat{\alpha}_i^m p_j + \pi_V^m PrivateVoucher_j + \tau_V^m p_j + \epsilon_{ij}$$

The result is surprising: while both models underpredict voucher use, they *over*-predict usage at higher tuition private schools. In other words, offered students use the voucher at lower tuition schools than expected. Further, the coefficient on tuition is remarkably similar between models and, though not shown in the table, this pattern holds across levels of household wealth. This

finding is key for understanding the sources of mis-specification in the control market estimates. In particular, it suggests that conventional unobserved school characteristics (i.e. insufficiently addressing tuition endogeneity in the control markets) are not the issue. This is because offered students do not have to pay the tuition, so the presence of school unobservables unaccounted for by the control models would instead predict a positive slope on tuition. Rather, if there is an unobservable school "quality" that voucher winners are sorting on, it is negatively correlated with tuition.

We use the hypothesis testing framework to examine several other kinds of mis-specification of the ability-to-pay constrained model that we pre-committed to (see Table 17 in Arcidiacono et al. 2021). We focus on the constrained model henceforth because it achieves better out-of-simple fit to the choices of voucher winners. Column (7) of Table A7 adds interactions between students' baseline math scores and school characteristics—a private school intercept, whether English medium, estimated math valued-added, and whether offers Hindi instruction—to the model to test for ability sorting. The control models did not include ability heterogeneity. The results suggest that some mis-specification may come from greater take-up among higher ability students, but higher ability students actually "prefer" lower value-added schools (and vice versa). Column (8) allows for the possibility that offered students value voucher-eligible private schools' attributes differently than implied by the control models. These results indicate higher disutility of travel to voucher schools, much weaker preferences for English medium instruction and for value-added, and greater preferences for Hindi classes. While interesting, the inclusion of these covariates does little to explain the overall underprediction of voucher take-up nor does their inclusion meaningfully modify the negative coefficient on voucher school tuition.

5 Unified Model

The results of the validation using treatment market data point to several important findings. First, the ability-to-pay constrained model achieves relatively better fit to the experimental patterns. Nonetheless, a large gap between predicted take-up and experimental take-up persists. A key question for the welfare analysis is thus what this gap represents. In this section, we follow-up on two clues revealed during the course of control model validation: First, the intervention appears to have caused voucher losers to attend private schools more than they otherwise would have. Second, conditional on taking-up the offer, voucher winners appear to prefer lower tuition private schools

all else equal.

In this section, we first advance explanations for these findings and provide supporting evidence from the treatment data. We propose that the voucher impacted choice through search, including of voucher losers who anecdotally anticipated they would also get vouchers, and that private schools used program surplus to incentivize voucher recipients to enroll. We then detail a unified empirical model that incorporates these new mechanisms (along with an ability-to-pay constraint) which we estimate on the combined control and treatment markets data. We show the unified model successfully rationalizes the data patterns and finally use it to estimate welfare effects.

5.1 What Was Missing?

5.1.1 Search

That voucher losers in treatment villages—in contrast with those eligible who did not apply and ineligibles—enrolled in private schools at much higher rates than what the control models predicted suggests the intervention impacted their choice in some way. While numerous possibilities exist, our proposed explanation is that voucher losers expected they may get a voucher too and, hence, searched for private school options.[30] When it was later revealed they had to pay tuition, they nevertheless chose a private school on account of the information gained from searching.

Our principal evidence for this explanation comes from the handful of treatment market villages mentioned earlier where, in the end, no household was able to actually use a voucher because the private schools in these villages chose to not participate in the voucher program. Call these villages "flagged." We estimate a linear probability model where the dependent variable is private school attendance. We control for whether the student applied for the voucher interacted with treatment village and flagged villages, whether they were offered a voucher interacted with flagged villages, where they were eligible for the voucher.

The results are presented in Table 6. The interaction between applying for a voucher and treatment village is large and positive (and matches the evidence shown earlier of a 15 point discrepancy), as is the effect of winning a voucher. But what is interesting is what happens in flagged treatment villages: namely, we no longer see an effect of winning a voucher nor do we see that flagged villages have private school attendance that is any different from other villages for non-applicants. Yet, both voucher winners and voucher losers attend private schools at similar

[30] Note there are only minor imbalances between the groups on observables and that the survey evidence we have is consistent with voucher losers paying tuition at private schools.

Table 6: Private School Attendance in Non-compliant Treatment Villages

	Attend voucher private	
Offered AP voucher	0.377***	0.412***
	(0.030)	(0.031)
Offered × Flagged village		-0.399***
		(0.115)
Applied for AP voucher	0.068**	0.068**
	(0.033)	(0.033)
Applied × Treatment village	0.155***	0.154***
	(0.039)	(0.039)
Applied × Flagged village		-0.001
		(0.115)
Treatment village	-0.029	-0.025
	(0.026)	(0.027)
Flagged village		-0.043
		(0.061)
Ineligible for AP voucher	0.622***	0.623***
	(0.031)	(0.030)
Constant	0.197***	0.197***
	(0.030)	(0.030)
Observations	2,960	2,960

Notes: Table reports estimates of linear probability models of private school attendance among kindergartners to examine differences in attendance pattern in "flagged" non-complying treatment villages where the program was not successfully implemented. Excluded group is AP voucher-eligible households who did not apply. Standard errors reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

rates to voucher losers in other treatment villages and correspondingly attend at higher rates than applicants in control villages.

Overall, the patterns in Table 6 are consistent with voucher losers and voucher winners in flagged villages equally anticipating a voucher offer, searching private schools under that pretense, and—for some—drawing sufficiently high match qualities to rationalize elevated private school attendance even after it was later revealed they would have to pay tuition. At the sane time, they are inconsistent with alternative explanations, chiefly peer effects, since the peer effect on private attendance in flagged villages should be sharply attenuated, but voucher losers are still equally likely to attend private schools in those villages as in any other treatment village.

5.1.2 Enrollment Incentives

Why do voucher winners appear to prefer low tuition private schools? We propose pass through of the voucher surplus an the explanation. Such a supply-side response make rational sense given the program's design: the voucher's yearly value was set to the 90th percentile tuition level, i.e. about 44% more than the annual tuition and fees charged by the average private school. A profit-

maximizing private school with tuition below the voucher amount would thus try to attract voucher students by sharing the surplus generated and, importantly, this incentive will be stronger for lower tuition private schools.

Table 7: Voucher Pass-through: Survey Responses for Focal Child and their Siblings

	(1)	(2)	(3)
	Private	Tuition and fees (Rs.)	
Offered voucher	0.542***	-2,742***	-580.5***
	(0.0277)	(199.5)	(113.1)
Constant	0.220***	3,153***	760.5***
	(0.0424)	(263.1)	(127.1)
Observations	948	395	941
Sample	All	Private=1	All
	Siblings (ages 5-9)		
Offered voucher	0.152***	-860.9**	289.2
	(0.0470)	(392.4)	(179.0)
Constant	0.265***	1,396***	313.6*
	(0.0851)	(444.9)	(181.5)
Observations	452	183	441
Sample	All	Private=1	All

Notes: Table reports ITT estimates of voucher offer impact on private school attendance (column 1) and spending on tuition and fees (column 3) according to post-intervention survey data on focal study child (upper panel) and their primary school-aged siblings (lower panel). Column (2) examines differences in spending on tuition and fees conditional on the focal study child attending a private school. Each upper panel observation is a study kindergartner; each lower panel observation is a school-aged sibling of a study kindergartner. Standard errors reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

While rational, a challenge for this explanation is how the surplus could feasibly be shared with voucher students. We present evidence one way this is achieved is by offering scholarships to voucher students' siblings. Specifically, we examine post-intervention survey responses of households in control and treatment markets regarding private school attendance and their expenditure on tuition and fees. Importantly, the survey includes responses pertaining to the focal child, who did (treatment) or would have received a voucher (control), as well as for their siblings in the household. The top panel of Table 7 shows, as expected, that randomly offered households report 54 point greater private school attendance for the main child (column 1) and report spending about Rs. 600. less on the main child's tuition and fees (column 3). The bottom panel of the table reports analogous intent-to-treat estimates for school-aged siblings of the main child. The key finding is: the offer raises the probability their sibling attends private school by 15 points (column 1) without

changing the household's spending on tuition and fees for the sibling child (column 3).[31]

5.2 Unified Model

In this subsection, we detail an empirical model that we then take to the entire dataset that has two new features: 1) search—households must pay a cost to reveal their match qualities at private schools and all voucher applicants in treatment villages anticipate receiving a voucher; and 2) enrollment incentives—participating private schools in treatment villages share a fraction of the program's surplus with voucher recipient.

The ex-post utility from a participating private school that voucher applicants in treatment villages expect (minus the preference shock) in our unified model is given by:

$$u_{ij}^V = u_{ij} + \alpha p_j + \theta(V - p_j) \times \mathbf{1}[V > p_j] \tag{9}$$

where u_{ij} is the "control" utility previously given by equation (2). αp_j is added to this because these households anticipate receiving a voucher (recall that the coefficient on tuition and fees in u_{ij} was α). The strength of the enrollment incentive is governed by θ, the parameter on the difference between the voucher amount ($V = 2.6$ since the program paid private schools Rs. 2,600 for each voucher enrollee) and private school j's tuition and fees; no incentive is applied if the school's fees exceed the vouhcer amount. It is intuitive to see from this equation that the enrollment incentive will be larger at low-tuition private schools, potentially reconciling why more voucher winners than expected by the control models attend such private schools.

The second extension is to introduce search. In particular, households have full information about government schools, but must pay a cost to reveal their match (represented by the preference shocks, the ϵs) with private schools. Letting c_i represent the cost and G_i denote the set of government schools in i's village, applicants in treatment villages search when:

$$\begin{aligned} c_i &< \ln\left(\sum_{j\in\mathcal{V}_i} \exp u_{ij}^V\right) - \ln\left(\sum_{j\in G_i} \exp u_{ij}\right) \\ &< -\ln(P_{iG|S}^V) \end{aligned} \tag{10}$$

where $P_{iG|S}^V$ is the probability i chooses a government school conditional on searching and receiving

[31] The middle column of Table 7 (column 2) shows that, conditional on the focal child attending a private school, offered households report spending essentially zero on the main child's tuition and fees and report spending about 60% less than control households on their siblings' tuition and fees.

a voucher. In contrast, "control" households will search for private schools when

$$\begin{aligned} c_i &< \ln\left(\sum_{j\in\mathcal{V}_i} \mathbf{1}[p_j \leq \omega_i] \exp u_{ij}\right) - \ln\left(\sum_{j\in G_i} \exp u_{ij}\right) \\ &< -\ln(P_{iG|S}) \end{aligned} \tag{11}$$

where, recall, ω_i represents unobserved ability-to-pay. Thus, absent a voucher, constrained households will be less likely to pay the search cost because many of the private schools will be outside of their price range regardless, limiting the benefits. With this added mechanism (characterized by two new parameters, the location and scale of c_i which we assume is exponentially distributed), the voucher thus can affect private school attendance both through searching (by increasing the expected gains from search) and by making private schools more attractive conditional on searching.

It is this search channel that also provides a mechanism to explain higher private school attendance by applicants in treatment villages who did not receive a voucher. Specifically, we treat these households, consistent with the patterns presented earlier, as *expecting* to get the voucher, as in equation (10). Then, at the stage where they must make a decision as to which school to attend, they receive no enrollment incentive and must pay full price at participating private schools (i.e. their ex-post utility is given by u_{ij}, not u_{ij}^V). In the case of ineligible students, we assume they paid the search cost earlier; the reason they were ineligible for the voucher program is that they were attending a private school pre-kindergarten.

We estimate the unified model on the combined the control and treatment markets data, pooling households across all subgroups visually represented in Figure 1. Details of the estimation, which like the control model estimation uses the EM algorithm, are included in the Appendix.

5.3 Unified Model Results and Fit

Table 8 presents selected parameter estimates of the unified model alongside those obtained from the control ability-to-pay model; the full set of indirect utility parameter estimates are presented in Table A9.

The unified model estimates are generally similar to those obtained on the control markets data alone by the ability-to-pay constrained model, e.g. the utility of attending an English medium school. The estimates imply that the average voucher program applicant would pay over Rs. 500 for a one standard deviation (on the student distribution) increase in the math value-added of

Table 8: Estimates: Selected Parameters—Control Ability-to-Pay and Unified Models

	Control	Unified
Tuition and fees (1000s of Rs.)	-1.28	-1.54
	(0.58)	(0.07)
First stage residual	1.77	1.60
	(0.63)	(0.07)
Enrollment incentive		2.26
		(0.19)
Private school	11.35	9.74
	(2.35)	(0.39)
× Eligible for AP voucher	-10.13	-5.51
	(1.76)	(0.45)
Private random effect σ	2.66	1.77
	(0.27)	(0.09)
Search		
Location		-0.24
		(0.09)
Scale		0.36
		(0.03)
Ability-to-pay constraint		
Intercept	2.96	3.39
	(0.55)	(0.68)
Eligible for AP voucher	-1.29	-0.77
	(0.41)	(0.36)
Asset factor	1.09	1.20
	(0.23)	(0.28)
σ	1.34	1.48
	(0.28)	(0.32)
N households	4,251	8,374
N observations	35,796	69,413

Notes: Table reports selected parameter estimates (and standard errors in parentheses) of control ability-to-pay constrained model (Control) and unified model estimated on entire dataset (Unified), including ability-to-pay constraint and search cost parameters. Parameter on total siblings in the constraint not reported; indirect utility parameter estimates for both models are reported in Table A9.

their primary school. This translates into about 3% of median consumption per capita. Table 8 shows, however, that eligible students' utility from attending a private school (all else equal) is much larger according to the unified model. This coefficient increases substantially (about $4x$) due to the incorporation of search costs. The coefficient on the enrollment incentive term of the unified model reported in Table 8 is large and positive. Recall that this coefficient is identified from the types of schools voucher winners attend relative to what we would expect based on the behavior of those in control villages.

Table 9: Estimates: Ability-to-pay Constraint and Search Probability

	Share unable to pay for...				Search privates	
	any private		*priciest* private			
	Control	Unified	Control	Unified	Control	Unified
First graders						
Overall	0.09	0.05	0.18	0.10	1.00	0.53
Lower caste	0.13	0.07	0.25	0.114	1.00	0.43
Both parents completed primary	0.11	0.05	0.23	0.10	1.00	0.35
Asset level < 3	0.24	0.14	0.44	0.26	1.00	0.43
Asset level $= 3$	0.09	0.04	0.20	0.09	1.00	0.51
Asset level $= 4$	0.03	0.01	0.08	0.03	1.00	0.56
Asset level > 4	0.01	0.01	0.03	0.01	1.00	0.60
Voucher program applicants						
Control markets	0.13	0.06	0.25	0.13	1.00	0.47
Voucher losers	0.12	0.07	0.27	0.14	1.00	0.81
Voucher winners	0.16	0.08	0.33	0.17	1.00	0.82

Notes: Table reports estimates for shares of households constrained by ability-to-pay absent the voucher and who search private school options by subgroup per the estimates of the control ability-to-pay constrained model (Control) and unified model (Unified). Any and priciest private schools refer to among those in the household's village.

The share of each group that paid the search costs as well as how binding the ability-to-pay constraint is are shown in Table 9. Voucher winners are substantially more likely to search for private than applicants in control villages, who are in turn more likely to search than eligible non-applicants in either control or treatment villages. Voucher losers in treatment villages are also more likely to search as we treat them as expecting to receive a voucher at the search stage in order to reconcile their high rates of private school attendance.

The estimates in Table 8 suggest that ability-to-pay is less constraining when search is accounted for and this is also reflected in the numbers in Table 9. The control model forces any search effects to instead operate through the ability-to-pay constraint. While less binding in general according the unified model, the constraint is still meaningful: 6-8% of voucher applicants cannot afford any private school and more than 13 percent cannot afford the most expensive private school in their

village per the unified model estimates.

Table 10: Unified Model Goodness-of-Fit

	Attend Private		Tuition|Private	
	Data	Unified	Data	Unified
First graders				
Overall	0.57	0.58	1.71	1.70
Lower caste	0.34	0.36	1.65	1.62
Both parents completed primary	0.27	0.28	1.48	1.60
Asset level < 3	0.28	0.33	1.45	1.57
Asset level $= 3$	0.52	0.54	1.72	1.70
Asset level $= 4$	0.68	0.66	1.84	1.76
Asset level > 4	0.78	0.78	1.67	1.69
Voucher program applicants				
Control markets	0.34	0.32	1.88	1.65
Voucher losers	0.48	0.45	2.13	1.91
Voucher winners	0.81	0.79	2.09	2.13

Notes: Table presents private school attendance and tuition given private school attendance by subgroup in the data with numbers implied by the unified model estimates to assess goodness-of-fit. Note that differences in the Data column with Table 1 are due to the numbers in this table accounting for attrition weights, which are used in the unified model estimation. Also, winners unable to use the voucher are treated as voucher losers in the unified model estimation and accordingly included in that row in the table.

Table 10 shows that these two additional features—search costs and enrollment incentives—significantly improve the fit of the model, both with regard to the rate at which different groups attend private school but also by providing a better match with the posted tuition of the schools that voucher winners attend. The first set of columns shows actual private school attendance for different groups of students, the predicted rates using the control model, and the predicted rates using the unified models. The predicted rates of private school attendance, both for voucher winners and voucher losers in treatment villages, now are within three percentage points of what is observed in the data. The second set of columns repeats the exercise but focuses on tuition conditional on attendance. Enrollment incentives are important here, with expected tuition now in line with what is observed for voucher winners who attend private school.

5.4 Implications for Welfare from the Unified Model

In this subsection, we use the model estimates to study welfare impacts of counterfactual voucher programs. To do so, we simulate the school choices of households in control villages given a voucher amount and eligibility criteria that we specify. We implement a voucher as a coupon provided to households that pays up to its full value towards tuition at any government-recognized private schools. This exercise allows us to calculate effects on social welfare given the choice environments—

household locations, school location, and school amenities—that exist in the data. This exercise thus abstracts away from strategic and general equilibrium adjustments that would be expected at scale, such as entry and exit of private schools.[32]

We begin by considering a program that makes a voucher worth up to the 90th percentile of private school tuition universally available to all households. As Table 11 shows, our estimates predict that this program would raise the private school share from 58% to 75% of all households (17 points). We estimate that the average household who would otherwise attend a government school (i.e. the average complier) would be willing-to-pay (WTP) Rs. 1,460 for the program, which in present value terms translates to about 17% of median annual consumption per capita.[33] This value is given by the added inverse of the estimated compensating variation—the amount of income that each household would need to be compensated to keep their utility level with the program the same as without it. Though the voucher funds complier households' choice to switch to a private school, Rs. 1,690 on average are spent on every household who would have attended a private school anyway.

Table 11: Welfare Impacts of Universal Voucher

Outcomes			
Increase in private schooling			0.17
Average complier HH's WTP (1000s of Rs.)			1.46
Average inframarginal HH's cost (1000s of Rs.)			1.69
			Fiscal ext.
Welfare metrics			1/3 2/3
MVPF			1.33 3.05
Benefit/cost ratio	MEB	0.5	0.93 1.29
		1.5	0.70 1.06

Notes: Table reports welfare impacts of voucher program universally available to all households that covers tuition up to Rs. 2,600 if implemented in control markets, as estimated by the unified model. 1000 simulations. Willingness-to-pay calculated by compensating variation. Two fiscal externality scenarios are considered: one where, for every household induced to switch to a private school (i.e. complier), 1/3rd of government spending per pupil (Rs. 8,400) is cut; another where 2/3rds is cut. Benefit/cost ratios calculated assuming marginal excess burdens (MEB) of 0.5 and 1.5.

In evaluating the overall welfare impact of the voucher program, the inefficiency of paying the tuition of inframarginal households must be weighed against the potential fiscal externality: how much of per pupil spending in government schools (about Rs. 8,390 in Andhra Pradesh per

[32]Note that by providing the voucher to households as a coupon, neither private schools nor households can keep any surplus of the voucher amount above the school's price. For this reason, and because we would expect their effect to differ from the AP project in counterfactuals where siblings are also voucher-eligible, our simulations do not allow private schools to use voucher surplus to offer enrollment incentives.

[33]To convert to present value terms, we multiply the gain by $1 + \delta + \delta^2 + \delta^3 + \delta^4$ where δ is the product of 0.90 (a 10% annual discount rate) and 0.79 (the annual probability that a voucher recipient remains in private school.

Dongre 2012) could be cut for every student who uses the voucher to exit government schooling? We consider two scenarios: a small impact scenario, where just 1/3rd could be recovered, and a large impact scenario, where 2/3rds—approximately equal to the share of spending allocated to teachers—could be recovered. As Table 11 reports, we find marginal values of public funds (MVPFs) of 1.33 and 3.05, respectively. We also compute benefit/cost ratios of the program assuming different values of marginal excess burden of taxation (0.5 and 1.5).[34] If the fiscal externality is large, we find that the benefit/cost ratio of the universal voucher program exceeds one in either case, indicating that the program improves social welfare.

How do these welfare estimates compare with those generated by the control models? Appendix Table A10 presents side-by-side comparisons, showing that only the unified model estimated on control and treatment markets data finds that a universal voucher increases social welfare (given a large fiscal externality and large marginal excess burden). How it differs from the respective control model estimates is revealing. The control ability-to-pay constrained model arrives at meaningfully large welfare metrics (e.g. MVPF of 2.14) but for the wrong reasons: it ascribes even greater WTP to complier households, who it estimates are much more constrained than the unified model, while predicting only somewhat less take-up (1-2 points less). In contrast, the control random coefficient model predicts significantly less take-up (7 points lower) and estimates the average complier's WTP is over 25% lower.

While the unified model results highlight scenarios in which a universal voucher program raises welfare, it stands to reason that targeted programs could be more efficient by minimizing voucher use by inframarginal households. With this motivation, we consider a counterfactual program offering a voucher of the same value (Rs. 2,600), but which is only available to the bottom quarter of households in terms of socioeconomic status (i.e. own two or fewer assets). Table 12 reports welfare impacts from this program.[35] Though its impact on private schooling is smaller (7 points), the average complier's willingness-to-pay is comparable with the universal case. However, far less of each inframarginal household's tuition gets paid for by the targeted program. The result is greater efficiency. The MVPF is nearly three in the small fiscal impact scenario and equals infinity—the targeted voucher program is self-financing—when the fiscal externality is large. In this latter scenario, the targeted program would be expected to generate up to $2 in social welfare for every

[34] Note that MVPF $= \frac{\overline{WTP}}{\overline{cost} - \text{Take-up}^* \psi^* 8.4}$, while benefit/cost $= \frac{\overline{WTP} + \text{Take-up}^* \psi^* 8.4^* (1+MEB)}{\overline{cost}^* (1+MEB)}$ where $\psi \in [1/3, 2/3]$. MVPF is defined as ∞ (the program pays for itself) when the denominator is negative. $\overline{WTP}$ and $\overline{cost}$ are the average household's willingness-to-pay and tuition paid for by the program, respectively.

[35] Note that our simulations do not consider the effect targeted programs may have on the tuition charged by private schools to non-eligible students.

$1 in tuition the program pays for.

Table 12: Welfare Impacts of Voucher Targeted to Asset-Poor Households

Outcomes				
Increase in private schooling			0.07	
Average complier HH's WTP (1000s of Rs.)			1.52	
Average inframarginal HH's cost (1000s of Rs.)			0.22	
Welfare metrics			Fiscal ext. 1/3	2/3
MVPF			2.93	∞
Benefit/cost ratio	MEB	0.5	1.27	1.99
		1.5	1.05	1.76

Notes: Table reports welfare impacts of voucher program targeted only to households with two or fewer assets that covers tuition up to Rs. 2,600 if implemented in control markets, as estimated by the unified model. 1000 simulations. Willingness-to-pay calculated by compensating variation. Two fiscal externality scenarios are considered: one where, for every household induced to switch to a private school (i.e. complier), 1/3rd of government spending per pupil (Rs. 8,400) is cut; another where 2/3rds is cut. Benefit/cost ratios calculated assuming marginal excess burdens (MEB) of 0.5 and 1.5.

6 Conclusion

Our paper makes two sets of contributions. The first are empirical. Here we show that a model of school choice with ability-to-pay constraints, search costs, and supply-side responses matches the high voucher take-up rates observed in the Andhra Pradesh School Choice project. We estimate substantial welfare gains from the voucher program (as well as counterfactual voucher programs) in part due to the costs of government schools being significantly higher than their private counterparts. Further, our results show that the gain in consumer surplus is economically meaningful for many students induced into private schools by vouchers because of the presence of ability-to-pay constraints that otherwise constrain some households from consuming school quality up to the value of the numeraire good.

The second set of contributions are methodological. The control models successfully fit the out-of-sample choice patterns of "control" households in treatment markets. In addition, while we anticipated that underpredicting experimental take-up could likely stem from inadequate instruments (or a mis-specified control function)—concerns that occupy attention in the demand estimation literature—our experimental validation identifies other issues as first-order. Rather, in initially holding out the entirety of the treatment markets data, we missed the intervention's apparent effects on how households choose schools and on private schools' behavior. Moreover, our later empirical quantification of these mechanisms relies on the pairing of control with treatment data

for identification. Our experience thus points towards the necessity of developing and estimating equilibrium models (e.g. that incorporate supply-side responses) using both treatment and control variation for credible policy analysis, as in Attanasio, Meghir and Santiago (2012) and as with our unified model.

References

Abdulkadiroğlu, Atila, Parag A Pathak, Jonathan Schellenberg, and Christopher R Walters. 2020. "Do parents value school effectiveness?" *American Economic Review*, 110(5): 1502–1539.

Allende S.C., Claudia, Francisco Gallego, and Christopher Neilson. 2019. "Approximating the Equilibrium Effects of Informed School Choice."

Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja. 2017. "Report Cards: The Impact of Providing School and Child Test Scores on Educational Markets." *The American Economic Review*, 107(6): 1535–1563. Publisher: American Economic Association.

Arcidiacono, Peter, Karthik Muralidharan, Eun-young Shim, and John D Singleton. 2021. "Experimentally validating welfare evaluation of school vouchers: Part i." National Bureau of Economic Research.

Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman. 2022. "Smart matching platforms and heterogeneous beliefs in centralized school choice." *The Quarterly Journal of Economics*, 137(3): 1791–1848.

ASER. 2018. *Annual Status of Education Report.* Aser Centre.

Attanasio, Orazio P., Costas Meghir, and Ana Santiago. 2012. "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA." *The Review of Economic Studies*, 79(1): 37–66.

Banerjee, Abhijit, Rema Hanna, Benjamin A. Olken, Elan Satriawan, and Sudarno Sumarto. 2021. "Food vs. Food Stamps: Evidence from an At-Scale Experiment in Indonesia." National Bureau of Economic Research Working Paper 28641. Series: Working Paper Series.

Barseghyan, Levon, Maura Coughlin, Francesca Molinari, and Joshua C. Teitelbaum. 2021. "Heterogenous Choice Sets and Preferences."

Bau, Natalie. 2022. "Estimating an equilibrium model of horizontal competition in education." *Journal of Political Economy*, 130(7): 1717–1764.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan. 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy*, 115(4): 588–638.

Ben-Akiva, Moshe, and Bruno Boccara. 1995. "Discrete choice models with latent choice sets." *International Journal of Research in Marketing*, 12(1): 9–24.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Cameron, Stephen V., and James J. Heckman. 2001. "The Dynamics of Educational Attainment for Black, Hispanic, and White Males." *Journal of Political Economy*, 109(3): 455–499. Publisher: The University of Chicago Press.

Carneiro, Pedro, Jishnu Das, and Hugo Reis. 2022. "The value of private schools: Evidence from Pakistan." *Review of Economics and Statistics*, 1–45.

Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature*, 46(2): 333–383.

Delavande, Adeline, and Basit Zafar. 2019. "University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints." *Journal of Political Economy*, 127(5): 2343–2393.

Dongre, A. 2012. "What is the Per Child Expenditure in Government Schools?" In *Accountability Initiative*. New Delhi.

Duflo, Esther, Rema Hanna, and Stephen P. Ryan. 2012. "Incentives Work: Getting Teachers to Come to School." *The American Economic Review*, 102(4): 1241–1278.

Epple, Dennis, Richard E. Romano, and Miguel Urquiola. 2017. "School Vouchers: A Survey of the Economics Literature." *Journal of Economic Literature*, 55(2): 441–492.

Galiani, Sebastian, Alvin Murphy, and Juan Pantano. 2015. "Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment." *The American Economic Review*, 105(11): 3385–3415.

Galiani, Sebastian, and Juan Pantano. 2021. "Structural Models: Inception and Frontier." National Bureau of Economic Research Working Paper 28698. Series: Working Paper Series.

Gregory, Jesse. 2017. "The Impact of Post-Katrina Rebuilding Grants on the Resettlement Choices of New Orleans Homeowners."

Hastings, Justine, and Jesse M. Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review*, 108(12): 3493–3540.

Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger. 2005. "Parental Preferences and School Competition: Evidence from a Public School Choice Program." National Bureau of Economic Research Working Paper 11805.

Hausman, Jerry A. 1996. "Valuation of New Goods under Perfect and Imperfect Competition." In *The Economics of New Goods.* , ed. Timothy F. Bresnahan and Robert J. Gordon, 207–248. University of Chicago Press.

Kamat, Vishal, and Samuel Norris. 2020. "Estimating welfare effects in a nonparametric choice model: The case of school vouchers." *arXiv preprint arXiv:2002.00103*.

Keane, Michael P., and Kenneth I. Wolpin. 2001. "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment." *International Economic Review*, 42(4): 1051–1103. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E Waugh. 2023. "The welfare effects of encouraging rural–urban migration." *Econometrica*, 91(3): 803–837.

Larroucau, Tomás, Ignacio Rios, Anaïs Fabre, and Christopher Neilson. 2024. "College Application Mistakes and the Design of Information Policies at Scale."

Muralidharan, Karthik, and Venkatesh Sundararaman. 2015. "The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India." *The Quarterly Journal of Economics*, 130(3): 1011–1066.

Neal, Derek, and Joseph Root. 2024. "The Provision of Information and Incentives in School Assignment Mechanisms." National Bureau of Economic Research.

Neilson, Christopher. 2013. "Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students."

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69(2): 307–342.

Pathak, Parag A, and Peng Shi. 2021. "How well do structural demand models work? Counterfactual predictions in school choice." *Journal of Econometrics*, 222(1): 161–195.

Petrin, Amil. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, 110(4): 705–729.

Petrin, Amil, and Kenneth Train. 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research*, 47(1): 3–13.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics*, 133(2): 553–609.

Rothstein, Jesse M. 2006. "Good Principals or Good Peers? Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition among Jurisdictions." *The American Economic Review*, 96(4): 1333–1350.

Sahai, Harshil. 2023. "The Welfare Implications of School Voucher Design: Evidence from India."

Schorfheide, Frank, and Kenneth I. Wolpin. 2012. "On the Use of Holdout Samples for Model Selection." *American Economic Review*, 102(3): 477–481.

Schorfheide, Frank, and Kenneth I. Wolpin. 2016. "To hold out or not to hold out." *Research in Economics*, 70(2): 332–345.

Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong. 2014. "Micro-loans, Insecticide-Treated Bednets, and Malaria: Evidence from a Randomized Controlled Trial in Orissa, India." *American Economic Review*, 104(7): 1909–1941.

Todd, Petra E., and Kenneth I. Wolpin. 2006. "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility." *The American Economic Review*, 96(5): 1384–1417.

Todd, Petra E, and Kenneth I Wolpin. 2020. "The best of both worlds: Combining rcts with structural modeling." *Journal of economic literature*.

Appendices

A Unified Model Estimation

A.1 Likelihood

We estimate the unified model using a modified EM algorithm where latent eligibility (of first graders) and the structural parameters are estimated in separate maximization steps. Let θ represent the structural parameters underlying search, the ability-to-pay constraint, and utility. At the θ-maximization step, we maximize:

$$\tilde{L} = \sum_i w_i[\tilde{e}_i \ln \mathsf{L}_{1i}(\theta) + (1 - \tilde{e}_i) \ln \mathsf{L}_{0i}(\theta)]$$

where w_i is a vector of weights and $\mathsf{L}_{1i}(\theta)$ is i's likelihood contribution given they are eligible (and $\mathsf{L}_{0i}(\theta)$ is analogously defined). For kindergartners, $\tilde{e}_i$ is their observed AP voucher eligibility status; for first graders, $\tilde{e}_i$ is their conditional or posterior probability of eligibility. This posterior eligibility probability is given by:

$$\tilde{e}_i = \frac{e_i \mathsf{L}_{1i}(\theta)}{e_i \mathsf{L}_{1i}(\theta) + (1 - e_i)\mathsf{L}_{0i}(\theta)}$$

where e_i is the (logit) probability i is eligible. The algorithm iterates until the parameters converge.

We assume ineligible households already paid the search cost, so their likelihood contribution is equivalent to the control ability-to-pay constrained model. The likelihood contribution of eligible households reflects both ability-to-pay and search, however:

$$\mathsf{L}_{i1}(\theta) = \sum_{j_i^*} \phi_{ij_i^*} \prod_{j \in \mathcal{V}_i} [\frac{1}{R} \sum_r P_{ij}^r(j_i^*)]^{d_{ij}}$$

where the numerical integration over the private school random effects is represented by the r superscript and the choice probability is:

$$P_{ij}^r(j_i^*) = \begin{cases} 0, & j \text{ private } \& \ p_j > p_{j_i^*+1} \\ P_{ij|S}^r(j_i^*) P_{iS}^r(j_i^*), & j \text{ private } \& \ p_j \leq p_{j_i^*+1} \\ P_{ij|\neg S}^r(j_i^*)(1 - P_{iS}^r(j_i^*)) + P_{ij|S}^r(j_i^*) P_{iS}^r(j_i^*), & j \text{ government} \end{cases}$$

$P_{ij|\neg S}^r(j_i^*)$ is the probability i chooses (government) school j if they do not search for private schools and $P_{iS}^r(j_i^*)$ is the probability that i searches given choice set j_i^*. The probability of searching is given by:

$$P_{iS}^r(j_i^*) = 1 - \exp[\pi_l + \pi_s \ln P_{iG|S}^r(j_i^*)]$$

where $P_{iG|S}^r(j_i^*)$ is the probability that i chooses *any* government school after searching. π_l and π_s are the location and scale, respectively, of the exponentially distributed search cost shock. Note that for treatment market kindergarten applicants (i.e. voucher winners and losers), this probability

will embed their expectations that they will not have to pay tuition and fees at private schools and will receive an enrollment incentive. Ex-post, this expectation is not met for voucher losers.

A.2 First Stage

We estimate the first stage of private schools' tuition and fees on observed school characteristics and instruments on the full sample of private schools from both control and treatment markets. The first stage does not allow for heterogeneity by village treatment status, consistent with an assumption that the intervention did not impact tuition-setting. The estimates are shown in Table A8. The results in the text use the baseline IVs (column 3); the cost proxy instrument is a less meaningful predictor of tuition on the combined sample. The estimates in column (3) imply that the average treatment village private school is unobservably better than the average control private school.

A.3 Unified Model Sample and Weights

The estimation sample for the unified model combines all of the subgroups shown in Figure 1. Households in "flagged" treatment villages are included in the estimation sample, but voucher winners in these non-compliant villages are treated like voucher losers in the estimation. Likewise, for those households we code as intending to use the voucher but who did not actually use one, their (non-voucher) school choice observed in the data is matched with the one predicted by being a household that expects a voucher but ex-post does not receive one.

The weights used in estimation are the product of sampling weights and attrition weights. The weights to adjust for the AP project's sampling design are constructed in the same way discussed in Arcidiacono et al. (2021). For all kindergartner subgroups other than voucher winners, the attrition weights are constructed using a probit model of attrition. For voucher winners, the attrition weights are constructed as described in the text: by re-weighting actual voucher users in non-flagged villages such that the weighted size of the winner sample equals that expected based on the size of the control applicant sample; the probit attrition model estimates are used to reflect relative likelihoods of attritting between vouchers winners. Attrition weights are 1 for all first graders.

B Additional Tables

Table A1: Summary Statistics: Household Characteristics by Subgroup

	First Graders				Kindergartners					
	Attend Gov't		Attend Private		Applicants		Non-applicants		Ineligible	
	Mean	Diff	Mean	Diff	Mean	Diff	Mean	Diff	Mean	Diff
Female	0.52	0.02	0.47	0.02	0.58	-0.02	0.55	0.07	0.47	-0.00
Lower caste	0.34	0.01	0.12	-0.01	0.32	0.03	0.36	-0.02	0.11	-0.02
Muslim	0.06	-0.00	0.09	-0.01	0.07	0.02	0.07	-0.06*	0.08	0.02
Christian	0.07	0.01	0.04	-0.01	0.08	0.01	0.11	-0.02	0.04	0.02*
# siblings	2.37	0.01	2.18	-0.12**	2.23	0.05	2.29	-0.08	2.13	-0.03
Older sibling in gov't school	0.50	0.01	0.11	-0.06***	0.37	-0.00	0.48	0.02	0.10	-0.03
Both parents completed primary	0.09	-0.00	0.34	-0.03	0.17	0.01	0.15	-0.02	0.35	-0.01
$\geq$ 1 parent completed secondary	0.06	0.00	0.25	-0.04	0.10	0.00	0.07	-0.01	0.25	-0.05
Both parents laborers	0.45	-0.01	0.18	0.04*	0.39	0.00	0.43	-0.05	0.19	-0.03
Math score σ (baseline)	0.02	0.01	0.64	0.14**						
Telugu score σ (baseline)	0.03	0.07**	0.72	-0.03	0.00	0.04	-0.04	-0.42***	0.39	-0.15**
Owns home	0.75	0.01	0.76	0.05*	0.76	-0.01	0.76	-0.00	0.77	0.00
Pucca house	0.72	0.01	0.92	-0.02	0.75	0.01	0.65	0.03	0.91	-0.00
Water facility in home	0.41	-0.01	0.60	-0.04	0.44	-0.07***	0.45	-0.05	0.61	-0.08**
Household toilet	0.24	-0.02	0.58	-0.00	0.28	-0.03	0.23	0.04	0.57	0.05
Owns land	0.18	0.02**	0.31	-0.02	0.19	-0.01	0.17	0.09*	0.33	0.02
Asset level < 3	0.39	-0.02	0.13	0.02	0.36	0.04	0.40	-0.06	0.12	0.01
Asset level $= 3$	0.27	0.00	0.21	-0.02	0.26	-0.02	0.26	-0.01	0.20	-0.03
Asset level $= 4$	0.20	0.02	0.29	-0.03	0.23	-0.01	0.23	0.04	0.27	0.00
Asset level > 4	0.13	0.00	0.37	0.02	0.15	-0.01	0.11	0.04	0.40	0.01
First principal asset factor	-0.13	0.01	0.43	-0.06	-0.05	-0.04	-0.15	0.06	0.44	-0.01
N households	4439		975		1915		258		787	

Notes: Table reports summaries of household characteristics by subgroups as well as treatment-control balance checks. Means refer to all households (in control and treatment markets); columns labeled "Diff" report differences in means (and their statistical significance) between households in the subgroup in treatment markets and in control markets. *** p<0.01, ** p<0.05, * p<0.1

Table A2: Summary Statistics: Characteristics of Primary Schools

	Government		Private	
	Mean	Diff	Private	Diff
Tuition and fees (Rs.)	0.81	-1.45	1924	226**
English medium	0.02	0.00	0.57	-0.08*
Unrecognized	0	.	0.23	-0.04
Mid-day meals	0.99	0.00	0.03	-0.01
Kitchen facility	0.26	0.04	0.01	-0.00
Full pucca building	0.89	-0.01	0.52	0.08**
Library	0.94	-0.01	0.77	-0.01
Functional water tap	0.42	0.05	0.62	0.02
Functioning toilet	0.65	0.01	0.84	0.05
Separate toilet for girls	0.34	0.07*	0.60	0.02
Staffroom for teachers	0.20	0.00	0.72	0.03
Playground	0.52	0.00	0.70	0.04
Has secondary school	0	.	0.27	0.05
Total school enrollment	74.28	-1.88	286.18	8.69
Multi-class teaching	0.70	0.10***	0.24	-0.06*
Pupil-teacher ratio	26.53	1.00	16.68	1.20
Share teachers absent	0.21	-0.04***	0.09	-0.01
Share teachers with BA	0.78	-0.00	0.54	-0.05
Share teachers with formal certificate	0.90	0.01	0.16	0.01
Share teachers female	0.50	-0.07***	0.71	-0.01
Share teachers lower caste	0.24	-0.02	0.12	0.01
Share teachers Muslim	0.02	-0.01	0.07	-0.01
Share teachers from village	0.25	0.03	0.48	0.02
Offers Hindi instruction	0	.	0.44	0.02
Offers computer skills	0.01	0.01	0.13	-0.00
School value-added	-0.04	0.02	0.04	-0.05
N	686		570	

Notes: Table reports summaries of school characteristics by government and private as well as treatment-control balance checks. Means refers to all schools; columns labeled "Diff" report differences in means (and their statistical significance) between schools in treatment markets and in control markets. ** p<0.01, ** p<0.05, * p<0.1

Table A3: Estimates: Selected Parameters—Control Models

	RC	CC
Tuition and fees (1000s of Rs.)	-2.35	-1.28
	(0.28)	(0.58)
× Eligible for AP voucher	0.07	
	(0.12)	
× Asset level = 2	0.45	
	(0.20)	
× Asset level = 3	0.74	
	(0.20)	
× Asset level = 4	1.12	
	(0.20)	
× Asset level > 4	0.81	
	(0.21)	
First stage residual	1.60	1.77
	(0.20)	(0.63)
Private random effect σ	2.23	2.66
	(0.22)	(0.27)
Ability-to-pay constraint		
Intercept		2.96
		(0.55)
Eligible for AP voucher		-1.29
		(0.41)
Asset factor		1.09
		(0.23)
σ		1.34
		(0.28)

Notes: Table reports selected parameter estimates (and standard errors in parentheses) of control random coefficient model (RC) and control ability-to-pay constrained model (CC). Coefficient on total siblings in ability-to-pay constraint excluded from the table. The estimation sample contains 4,251 households and 35,796 household-school observations. All indirect utility estimates for both models are reported in Arcidiacono et al. (2021).

Table A4: Coding Voucher Use

Voucher code	Tracking	N	N*	Use
accepted and admitted	Private	416	410	yes
	Government	9	9	no
rejected voucher	Private	8	8	yes
	Government	49	49	no
migrated	Private	1	1	yes
	Government	9	9	no
own private admission	Private	31	22	yes
	Government	12	11	no
under age	Private	7	6	yes
	Government	14	14	yes
admitted, dropped out	Private	29	27	yes
	Government	7	7	no
waiting list not used	Private	0	0	.
	Government	1	1	yes
school rejected	Private	9	0	.
	Government	27	0	.
	Total	629	574	489

Notes: Table displays our coding of voucher Use based on information from project team (Voucher code) and tracking data. N represents counts of households in each cell; N* reports counts excluding households residing in nine "flagged" treatment villages where, collectively, very few students were actually able to use a voucher to attend a private school.

Table A5: Validation: Voucher Elasticity of Private Schooling

	RCT	RC	CC
Overall	221	116	148
Female	252	126	159
Muslim	110	72	85
Lower caste	328	158	209
Older sibling in gov't school	474	262	335
Both parents completed primary school	116	87	110
≥ 1 parent completed secondary	66	63	78
Both parents laborers	259	132	176
Asset level < 3	303	189	247
Asset level $= 3$	190	125	162
Asset level $= 4$	247	87	124
Asset level > 4	136	78	78

Notes: Table presents average voucher elasticity (percent change in private schooling due to the voucher offer) of applicant households by subgroup in the treatment data (RCT), and as predicted by the random coefficient (RC) and ability-to-pay constrained control models (CC). Predictions correspond to baseline specification described in the text and detailed in Arcidiacono et al. (2021).

Table A6: Validation: Voucher Intent-to-Treat Effects and Elasticities on Characteristics of Chosen School

	RCT		RC		CC	
	ITT	ϵ	ITT	ϵ	ITT	ϵ
Tuition and fees (Rs.)	1.08***	183	0.68	120	0.94	168
English medium	0.13***	54	0.08	42	0.14	72
Distance to school (mi.)	-0.25	-21	-0.15	-15	-0.15	-14
School value-added	0.01		0.00		0.01	
Offers Hindi	0.33***	206	0.11	59	0.17	90
Unobservable	0.25***		0.08		0.07	

Notes: Table presents voucher intent-to-treat effects (ITT) and elasticities (ϵ) – the percent change in the average value of the choice characteristic due to the voucher offer – for applicant households in the treatment data (RCT), and as predicted by the random coefficient (RC) and ability-to-pay constrained control models (CC). Predictions correspond to baseline specifications described in the text and detailed in Arcidiacono et al. (2021). *** p<0.01, ** p<0.05, * p<0.1

Table A7: Validation: Hypothesis Tests for Mis-specification of Ability-to-Pay Constrained Control Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Private voucher school		2.60		5.28	4.53	4.70	4.58	3.98
		(0.22)		(0.40)	(0.42)	(0.50)	(0.43)	(0.48)
Private voucher school × Asset factor					0.03	-0.76		
					(0.31)	(0.63)		
Private voucher school × Older sibling in gov't school					1.63	0.89	1.72	1.74
					(0.42)	(0.81)	(0.43)	(0.45)
Tuition and fees (@ voucher school) (1000s of Rs.)			0.52	-1.32	-1.34	-1.43	-1.37	-1.26
			(0.08)	(0.16)	(0.16)	(0.21)	(0.16)	(0.18)
Tuition and fees × Asset factor						0.39		
						(0.26)		
Tuition and fees × Older sibling in gov't school						0.36		
						(0.33)		
Private voucher school × Ability							0.47	0.27
							(0.26)	(0.29)
English medium × Ability							-0.37	-0.34
							(0.28)	(0.29)
Value-added × Ability							-1.08	-1.08
							(0.29)	(0.31)
Offers Hindi × Ability							0.00	-0.12
							(0.32)	(0.34)
Distance × Private voucher school								-0.70
								(0.19)
English medium × Private voucher school								-0.66
								(0.32)
Value-added × Private voucher school								-2.09
								(0.63)
Has Hindi × Private voucher school								0.63
								(0.37)
First stage residual × Private voucher school								-0.08
								(0.21)
$\hat{\text{Use}}$	0.65	0.84	0.74	0.84	0.84	0.84	0.84	0.84
AIC	1,400	1,235	1,360	1,164	1,153	1,153	1,137	1,105

Notes: Table reports hypothesis tests of model mis-specification that examine variables' predictive power for voucher winners' choice patterns conditional on the indirect utility of the alternative implied by the control ability-to-pay constrained model estimates. Standard errors reported in parentheses.

Table A8: First Stage: Private School Tuition and fees

	(1)	(2)	(3)	(4)
	Control Markets		Control + Treatment	
	Baseline IVs	Alternative	Baseline IVs	Alternative
Product space location	238.1***	288.9***	197.0***	210.1***
	(56.42)	(59.60)	(44.07)	(44.45)
Cost proxy	0.376***		0.155*	
	(0.135)		(0.080)	
Cost index		0.246***		0.418***
		(0.0947)		(0.1105)
Cost index^2		-0.000599***		-0.000717***
		(0.000132)		(0.000138)
First-stage F	12.51	17.63	11.39	20.64
Cragg-Donald stat	11.20	13.13	11.39	16.29
R^2	0.309	0.341	0.232	0.265
Observations	293		570	

Notes: Table presents first stage estimates that regress private school tuition and fees on school characteristics and instrumental variables on the control markets sameple (columns 1 and 2) and the entire sample (columns 3 and 4). Baseline IVs refers to instruments summarizing product space location (first factor of fixed characteristics of other private schools in *same* village) and proxying for school-level costs (predicted tuition and fees based on similar private schools in *other* villages), while Alternative replaces the cost proxy with a village-level cost index (and its square). Estimation and validation results of control models in the text pertain to column (1); unified model estimation uses column (3). Though not reported, regressions control for the school characteristics included in the choice models. Standard errors reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A9: Estimates: Indirect Utility Parameters—Control Ability-to-Pay and Unified Models

	Control		Unified	
	Coef	SE	Coef	SE
Tuition and fees (1000s of Rs.)	-1.28	0.58	-1.54	0.07
First stage residual	1.77	0.63	1.60	0.07
Private random effect σ	2.66	0.27	1.77	0.09
Enrollment incentive			2.26	0.19
Log distance	-1.41	0.09	-0.69	0.06
× Eligible for AP voucher	0.29	0.15	-0.58	0.06
× Age > 5	0.15	0.08	0.02	0.05
× Female	-0.14	0.08	-0.04	0.05
× Muslim	0.13	0.14	0.19	0.08
× Lower caste	-0.05	0.09	0.03	0.05
Private school	11.35	2.35	9.74	0.39
× Eligible for AP voucher	-10.13	1.76	-5.51	0.45
× Female	-0.60	0.25	-0.20	0.13
× Muslim	0.16	0.46	-0.01	0.23
× Lower caste	-1.50	0.29	-0.74	0.14
× Both parents completed primary	0.17	0.42	0.76	0.20
× ≥ 1 parent completed secondary	0.58	0.53	0.10	0.25
× Older sibling in gov't	-2.59	0.49	-2.15	0.12
× Total siblings−2	-0.07	0.10	-0.17	0.07
English medium	0.90	0.40	0.95	0.10
× Female	-0.92	0.23	-0.64	0.11
× Muslim	1.42	0.47	1.07	0.18
× Lower caste	0.05	0.27	-0.03	0.14
× Both parents completed primary	0.85	0.29	0.68	0.15
× ≥ 1 parent completed secondary	1.35	0.60	1.31	0.17
Unrecognized private school	-0.61	0.16	-1.03	0.08
Value-added	0.51	0.17	0.33	0.20
× Female	0.03	0.19	0.05	0.10
× Muslim	-0.16	0.29	0.37	0.19
× Lower caste	-0.04	0.20	-0.37	0.05
× Both parents completed primary	0.31	0.25	0.28	0.17
× ≥ 1 parent completed secondary	-0.51	0.28	-0.32	0.19
Offers Hindi	0.03	0.33	-0.09	0.11
× Female	0.55	0.34	0.10	0.13
× Muslim	1.15	0.43	0.63	0.21
× Lower caste	1.15	0.38	0.16	0.16
× Both parents completed primary	0.39	0.30	0.13	0.17
× ≥ 1 parent completed secondary	0.18	0.33	0.09	0.19
Closest public school	0.78	0.12	0.59	0.06
Facilities factor	0.46	0.07	0.32	0.03
Teaching quality factor	-0.34	0.05	-0.14	0.04
Teacher characteristics factor	-0.06	0.04	-0.08	0.02
N households	4,251		8,374	
N observations	35,796		69,413	

Notes: Table reports point estimates (and standard errors) for indirect utility parameters of control ability-to-pay constrained model (Control) and unified model estimated on full dataset (Unified). Estimates on indicator for whether school serves secondary grades, whether value-added is missing, whether tuition is imputed, and whether distance is missing not reported.

Table A10: Comparing Universal Voucher Welfare Estimates Between Control Models and Unified Model

	Control		
	RC	CC	Unified
Outcomes			
Increase in private schooling	0.10	0.16	0.17
Average complier HH's CV	1.06	2.05	1.46
Average inframarginal HH's cost	1.81	1.85	1.69
Welfare metrics			
MVPF	1.28	2.14	3.05
Benefit/cost ratio	0.73	0.95	1.06

Notes: Table reports welfare impacts of voucher program universally available to all households that covers tuition up to Rs. 2,600 if implemented in control markets, as estimated by the control random coefficient (RC) and ability-to-pay constrained (CC) models and the unified model (Unified). 1000 simulations. Willingness-to-pay calculated by compensating variation. Welfare metrics calculations assume that, for every household induced to switch to a private school (i.e. complier), 2/3rds of government spending per pupil (8,400 Rs.) is cut; benefit/cost calculated assuming marginal excess burden (MEB) of 1.5.