

Experimentally Validating Welfare Evaluation of School Vouchers - Part I*

Peter Arcidiacono[†]
Karthik Muralidharan[‡]
Eun-young Shim[§]
John D. Singleton[¶]

July 2021

Abstract

In this paper, we use a unique two-stage experiment that randomized access to school vouchers across both markets and students in rural India to estimate the revealed preference value of school choice. In the first step of the research design, we develop an empirical model of school choice subject to liquidity and credit constraints that is estimated using data from only the control markets. Based on this exercise, we estimate that the voucher generated welfare gains exceeding four times the average private school's annual tuition on average to the students induced into private schooling. The second step of the research design will validate the estimated welfare impacts by comparing model predictions for a simulated voucher program in control markets with the data from the treatment group. The results in this paper are based on the first step (using only control data) and this draft serves as a pre-commitment to the model estimates and predictions before examining the experimental data.

Keywords: welfare effects, school choice, vouchers, experimental validation

JEL codes: D12, H42, I21, I28, O15

*We thank Joseph Altonji, Dennis Epple, Adam Kapor, Kala Krishna, Robert Miller, Christopher Neilson, Ronni Pavan, Cristián Sánchez, Kegen Tan, Petra Todd, Christopher Walters, and seminar participants at ASSA, ASU, the Banff Empirical Micro Workshop, Barcelona GSE, Clemson, Cleveland FRB, Cornell, Duke, ITAM, NYU, Penn State, SOLE, SUNY Buffalo, UCLA, UC San Diego, and Wisconsin for helpful comments and discussions. Paulo Lins provided excellent research assistance.

[†]Duke University & NBER (psarcidi@econ.duke.edu)

[‡]UC San Diego, NBER & JPAL (kamurali@ucsd.edu)

[§]Amazon (eunyoung.shim@gmail.com)

[¶]University of Rochester (john.singleton@rochester.edu)

1 Introduction

Governments routinely 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 and growing 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 paternalistic taxpayers and policy makers who care about the most cost-effective way to achieve outcomes of interest.

Yet, this default approach ignores the question of what program beneficiaries themselves may prefer. For instance, vouchers may increase beneficiary welfare by increasing choice and improving match quality on unobserved outcomes. This implies that policy evaluations of voucher programs should account for both impacts on those outcomes that a paternalistic policy maker may care about as well as on the welfare of program beneficiaries. More generally, beneficiary valuation of publicly-provided benefits should be a key parameter for policy evaluation, but is often ignored in the impact evaluation literature, in large part because it is not easy to estimate.²

In this paper, we complement an existing evaluation of the test score impacts of private school vouchers in rural India by also quantifying program effects on welfare based on revealed preference. The data are drawn from the Andhra Pradesh School Choice project, a randomized controlled trial conducted in the Indian state of Andhra Pradesh (AP) whose test score impacts are reported in Muralidharan and Sundararaman (2015). Our research design leverages the project’s two-stage randomization across both markets and individuals: After eliciting initial interest in participation, the experiment randomized villages (markets) into treatment and control groups. The program then paid the tuition and fees at private schools for randomly-selected students in the treatment villages for the duration of primary school. The data collected for households and schools that were randomized out of the program at the village-level thus creates a control sample uncontaminated by the voucher offers. We use this market-level randomization to pursue a research design aimed at credibly validating the estimated welfare impacts of private school vouchers.

In the first step of our research design, we estimate empirical models of school choice solely using control group data from the trial (i.e. from villages randomized-out in the first stage). We consider two classes of discrete choice demand models. The first are random coefficient-style logit

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

²For instance, in their work on distributional national accounts, Piketty et al. (2018) value public goods at the *cost* of providing them because there is no easy way of estimating their value.

models that are standard for welfare analysis in the industrial organization literature (Berry et al., 1995; Nevo, 2001; Petrin, 2002). These models have been applied to several other contexts of school choice (e.g. Neilson 2013; Carneiro et al. 2016; Abdulkadiroglu et al. 2017a). Second, we develop an empirical model that incorporates a constraint on households’ ability-to-pay for private schooling absent a voucher. The constraint reflects the reality that liquidity and access to credit are often limited in low-income settings, such as rural India.³ Neglecting the influence of ability-to-pay on choice can have serious consequences for assessing welfare gains and for successfully predicting voucher take-up, especially among poor or lower wealth households.

In the research design’s planned second step, we will use the treatment arm of the voucher trial to experimentally validate the empirical models. Specifically, we simulate a voucher trial in the models that mimics the AP School Choice project’s offer to treated households. The simulation generates predictions for take-up of the voucher that we will compare with analogous moments to be computed directly from the treatment village data. We will also estimate treated households’ revealed preferences in a hypothesis testing framework to understand the sources of model misspecification. These comparisons and tests to come will allow us to validate the choice models in terms of how well they reproduce experimentally-generated choice patterns. We will use the validated model to evaluate the voucher program implemented by the AP School Choice project and a counterfactual universal voucher in welfare terms.⁴

In this paper, we present results from the first step of the research design and predictions for the experimental validation. Estimates of the model that allows for credit constraints suggest that households, particularly those targeted by the AP project’s voucher eligibility rules, lack the ability to pay for private schooling. We estimate that around 25% of households who attend a government school in the control markets are unable to pay for (and hence choose) any private school in their village. This constraint in turn has significant implications for demand and for estimates of willingness-to-pay. For instance, while the two classes of models estimate similar preferences for school quality and English instruction among wealthier and better educated households, the choice models that neglect differences in ability-to-pay assign less willingness-to-pay for these characteristics among asset-poor households. Our baseline constrained choice model estimates willingness-to-pay for school value-added among this group that is around 50% greater on average.

Simulations of the models show aggregate welfare gains from the AP project voucher from both

³For instance, Tarozzi et al. (2014), present experimental evidence that micro consumer-loans substantially raised ownership and use of insecticide-treated bednets in rural India, while demand was instead highly elastic when households had to pay upfront. Other evidence for the salience of credit constraints from similar contexts includes Rosenzweig and Wolpin (1993); Townsend (1994); Banerjee and Duflo (2014). In our data, 41% of households whose child attends a government school cite “economic reasons” as the explanation for their choice.

⁴The experimentally-validated demand model will also allow us to study the incidence and fiscal impacts of alternative targeting rules.

classes of models.⁵ However, the models differ significantly in the estimated magnitude of the welfare gains from the program. The difference stems from compliers: In the standard unconstrained demand models, those induced into private schooling by the offer value the voucher at a little more than half of its expected cost (and less than always takers). This contrasts sharply with the ability-to-pay models, wherein those constrained may value the voucher by more than its consumption equivalent. And this is what we find: in the baseline constrained model, those induced into private schooling realize higher consumer surplus gains on average than the inframarginal students. We estimate that the AP project voucher generated average welfare value to compliers over four times the average private school’s annual tuition (or about 33% of median annual expenditure per capita).

Consistent with marginal households’ greater willingness-to-pay, however, the ability-to-pay constrained models also predict that more households will comply with the experimental voucher offer. We find that the baseline ability-to-pay constrained model predicts that 63% of treated students will use the voucher to attend a private school. The unconstrained demand models, in contrast, predict take-up of around 10 percentage points less. This moment is a key point for comparison in the research design’s second step. With respect to aggregate welfare, these differences regarding both compliance and compliers’ valuations mean that the baseline model with ability-to-pay constraints estimates around an 80% larger improvement in aggregate welfare.

Our application contributes to a growing literature that uses experimental data to test and validate structural econometric models (e.g. Wise 1985; Todd and Wolpin 2006).⁶ A longstanding criticism of structural models in policy analysis is lack of credible identification of parameters of interest. LaLonde (1986), for example, showed that model-based estimates of training program impacts derived from observational data could not reliably reproduce experimental treatment effects. Schorfheide and Wolpin (2012, 2016) propose that holdout samples, wherein researchers reserve a portion of the available data (or use forthcoming data) for external validation of a model, can potentially provide corrective incentives that guard against “structural data-mining.” Nonetheless, this approach, which trades off efficiency against credibility, has rarely been taken by structural researchers.⁷

⁵These aggregate welfare gains are driven primarily by the much lower cost per student in private schools, which is on average 20-25% of the cost per student in government schools. Because the expected cost of each voucher is much lower than average per student expenditures on government schools, inducing more households to switch to private schools generates greater reductions in the social cost of education. The large fiscal externality more than offsets the deadweight loss of raising tax revenue to finance vouchers for inframarginal always takers even under a universal non-targeted voucher program.

⁶There is also a related body of work using experimental data to fit structural models, e.g. Attanasio et al. (2012); Lagakos et al. (2017); Salz and Vespa (2017). See Galiani and Pantano (2021) for a valuable recent summary and discussion of combining random (or quasi-random) variation with structural modeling.

⁷An important example is Misra and Nair (2011), who use a structural model of salesperson effort to provide recommendations for designing compensation contracts. The firm then implemented these suggestions, permitting validation of the model against the observed responses.

Our research design, which combines the validation step with data from a randomized control trial, makes two design choices: First, we hold out all of the treatment data from the estimation. While alternative choices might combine validation with using some experimentally-induced variation (by holding out just a portion of the treatment data or instead holding out the control data), our approach aims to test the performance of model-based approaches to welfare analysis with the kind of observational data environment frequently confronted in applied policy analysis.⁸ The second design choice is *pre-commitment* to the model estimates and the predictions generated in the first step (i.e. using only control data) and presented in this paper. This pre-commitment is to prevent model re-fitting *ex post* to the experimental validation to come.⁹

The control data environment of the first research step presents two main challenges to estimating willingness-to-pay for a private school voucher. These data include household-level survey information on each sampled child’s primary school, household and student socio-demographic characteristics, as well as data characterizing the available school alternatives – private and governmental – in each village. The data also include geocoded school and household locations. The first challenge is ubiquitous to identifying demand in observational settings: the simultaneity of supply and demand. In discrete choice frameworks, this is frequently represented by choice alternative-specific characteristics that remain unobserved to the econometrician but which are known to firms when making pricing and other decisions. Failing to properly account for the role of these unobserved characteristics is liable to lead to own-price elasticities that are too inelastic. As this bias would also translate into underpredictions for take-up of the voucher offer, this challenge represents a significant threat to our research design. Recent advances in other school choice contexts leverage quasi-experimental variation or policy changes to help overcome this endogeneity (Neilson, 2013; Allende S.C., 2019). In contrast, we rely on instrumental variable strategies that aim to isolate exogenous variation in private schools’ mark-ups (e.g. Berry et al. 1995; Hausman 1996). While we show that the instruments are important for producing plausible demand elasticities, the validity of these instruments rests on strong assumptions. One contribution of this paper is thus testing the ability of these instruments to reproduce experimentally-generated elasticities.

The second challenge is that households in this setting may likely face unobserved credit and liquidity constraints on their choice of school. Although government schools are tuition-free (and

⁸For example, Duflo et al. (2012), who study incentives to reduce teacher absence in India, use the treatment data from a randomized control trial for estimation and use the control data for model validation. Galiani et al. (2015) take advantage of multiple treatment arms to use randomized variation for both estimation and validation of a model of neighborhood choice. Our choice is similar to Keane and Wolpin (2007), who hold out a sample “that faces a policy regime well outside the support of the data” from the estimation of a female life-cycle model to provide confidence in its forecasts for policy changes.

⁹In this respect, our approach is similar to Pathak and Shi (2017), who validate school choice models fit prior to a policy change in Boston.

provide free mid-day meals), private schools are fee-charging, with average annual tuition and fees corresponding to over 7% of median per capita consumption. Such constraints may imply that some households effectively do not have the option to choose private schooling. While the first class of models assume that all private schools are available to households, the constrained model we develop instead models ability-to-pay as a latent stochastic variable that depends on observed asset information. This structure generates expressions for the probability that each private school belongs to a given household’s choice set, which is decreasing in its tuition. To support identification, we apply the exclusion restriction that household assets are unrelated to differences in indirect utilities between schools. Though preferences are allowed to depend on correlated characteristics (e.g. parental education), this assumption may be overly strong and the experimental validation is a key research step in assessing the model’s fitness and validity. The 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 that similarly quantifies the salience of credit constraints using choice models (Cameron and Heckman, 2001; Keane and Wolpin, 2001; Gregory, 2017).

Previous work on voucher programs, including the AP School Choice project, has focused mainly on impacts on student outcomes (Epple et al., 2017).¹⁰ Our focus on complementing the test-score impacts reported in Muralidharan and Sundararaman (2015) by also evaluating program’s welfare impacts connects with the design of efficient school choice mechanisms (e.g. Abdulkadiroğlu and Sönmez 2003; Abdulkadiroğlu et al. 2015) and relates with prior work that studies school preferences (e.g. Hastings et al. 2005; Rothstein 2006; Bayer et al. 2007; Abdulkadiroglu et al. 2017a).¹¹ How and why households choose schools is key to understanding incentives in education markets at scale (MacLeod and Urquiola, 2013, 2015). In this regard, our focus on the role of liquidity and credit constraints complements prior empirical work on the equilibrium effects of lower income students’ responsiveness to value-added (Bau, 2017) and to information (Andrabi et al., 2017a; Allende S.C. et al., 2019) when choosing schools.

The remainder of this paper is organized as follows. In the next section, we describe the background and data, collected from AP School Choice project. In Section 3, we describe our empirical model of constrained primary school choice in detail. Identification and the set of specifications

¹⁰Evidence from international settings tends to show positive effects, including in the longer run (e.g. Angrist et al. 2006), while recent U.S.-based evidence reveals strikingly negative impacts on students receiving a voucher to attend private school (e.g. Abdulkadiroglu et al. 2017b and Mills and Wolf 2017). Milwaukee’s voucher program has produced evidence of positive effects in the past (Rouse, 1998). Epple et al. (2017) summarize also the international and U.S. evidence regarding spillover effects, including via competition, on non-recipients.

¹¹A related connection is to work that uses housing prices to estimate willingness-to-pay for local public goods, including schools. Cellini et al. (2010), for example, find evidence from narrowly-decided bond referendums in California that such investments raise local housing prices. This unmet willingness-to-pay is indicative of underinvestment.

we estimate are discussed in Section 4. In Section 5, we present the results, which include model estimates, fit, and welfare calculations. We discuss the experimental validation and present model predictions for treated students' take-up of the voucher in Section 6. We conclude in Section 7.

2 The Andhra Pradesh School Choice Project

Our data are drawn from a randomized controlled trial conducted in 180 villages in the Indian state of Andhra Pradesh (AP) beginning in 2008. The AP School Choice project was designed to study the impact of school vouchers on student learning outcomes at both the student and market level. Students randomized into treatment were offered a voucher covering the costs of tuition and associated fees or expenses (e.g. books and uniforms) at private schools in their village for the duration of primary schooling (grades one through five). Expenses for transportation were not covered by the voucher. Villages selected for the project had to have at least one private school that agreed to participate in the project. Participation was voluntary at the school level, but schools were not allowed to screen or selectively admit voucher students.¹² The voucher value was set at around the 90th percentile of the fee distribution of private schools in the study sample, and was paid directly to schools' bank accounts.

Our research design uses the AP School Choice project's 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. The program was targeted to students likely to otherwise attend government schools. After eliciting interest, the project then randomized villages into treatment and control markets. Applicant households in treatment villages were then randomized into or out of the voucher treatment group. This double randomization design was important for estimating spillover effects on non-participants in the program (Muralidharan and Sundararaman, 2015). For this paper, the data collection for households and schools that were randomized out at the village-level provides a control group free of contamination from the experimental treatment. We use the data from the 90 control markets to fit empirical models of school choice, while we use the exogenous variation induced by the experiment to validate those models (and the estimated welfare impacts).

The project covered two cohorts: a cohort of first graders who had already chosen their primary school at baseline and a younger cohort, whom we term kindergartners, that chose their school after the baseline survey. We use both cohorts' choices in estimation, but focus on the kindergarten cohort for the validation of the empirical models because this is the cohort where, in the treatment

¹²The design stipulated that, similar to charter schools in the US, lotteries would be held to allocate places in oversubscribed schools. In practice, all applicants were accepted.

condition, students make their *initial* primary school choices with the AP project voucher.¹³ The project attempted to track all students who applied for a voucher, but only representative samples of those who did not or who were not eligible for the voucher at the start of the project.¹⁴ We also do not always observe a choice of primary school for kindergarten students in the control sample due to attrition. We therefore re-weight households to account for these elements of the project’s sampling design. We discuss the sample and weighting in expanded detail in Appendix A.

2.1 Sample and Data Summaries

Our estimation sample, drawn solely from the control data, contains detailed information for 4,251 households and 645 primary schools.¹⁵ Households were surveyed at baseline, while schools were surveyed beginning the first year of the program. The surveys elicited rich information important for modeling school choice, including the demographic and socioeconomic characteristics of students and households as well as data regarding each school’s amenities, tuition and fees, characteristics of teachers, curriculum, and finances.¹⁶ Geographic GPS locations were also collected, facilitating the mapping of travel distances between households and schools in their village. Students were assessed at baseline and during the third and fourth years of the program.

Private schooling has an aggregate market share of 57% across the project villages.¹⁷ This high penetration largely reflects the fact that the AP School Choice project was conducted only in villages that had at least one private school. The unconditional private school market share among primary school-going students in AP in 2008 was around 33% (ASER, 2018).

Students attending private schools come from higher socioeconomic status (SES) households on average than those attending government schools (Table 1). They are more likely to have parents who have both completed primary school (37% vs 9%), and have at least one parent who has completed secondary school (31% vs 5%). They are also less likely to have parents who both work as laborers (16% vs 42%). Their households are also wealthier on average as measured by asset ownership. Private schools students are more likely to live in a pucca (brick or stone) house,

¹³The treated first graders, on the other hand, chose whether or not to use the voucher to switch from their baseline primary school for the next grade.

¹⁴Kindergartners were eligible for the AP voucher if they were attending an Anganwadi (government daycare). First graders were eligible if they were attending a government primary school. One limitation of the data is that whether a first grader attended an Anganwadi (i.e. would’ve been eligible for the voucher *at the time of their primary school choice*) is not known. Because our models of primary school choice control for AP voucher eligibility, we treat it as a latent type for first grader observations in the estimation.

¹⁵We restrict the sample to kindergarten students at least four and at most seven years of age at baseline and to first grade students at least five and at most eight.

¹⁶When tuition and fees are not observed in the initial survey, we use subsequent survey years to impute them based on the school’s percentile in the overall distribution.

¹⁷This figure is obtained from 2008 survey data based on those households with children in the sample age range living in villages with at least one private school in the AP project districts (ASER, 2018).

have a water facility in the home, and to have a household toilet. These differences in wealth are suggestive of differences in ability-to-pay for private schooling. Students attending private schools also outperform students attending government schools on baseline exams. Students attending private schools score three fifths of a standard deviation above the government school student average scores in math.

Table 1: Characteristics of Households

	Attend		Average
	Government	Private	
Female	0.52	0.46	0.48
Lower caste	0.32	0.13	0.21
Muslim	0.06	0.11	0.08
# siblings	2.35	2.25	2.29
Older sibling in gov't school	0.50	0.13	0.29
Both parents completed primary school	0.09	0.37	0.25
≥ 1 parent completed secondary	0.05	0.31	0.20
Both parents laborers	0.42	0.16	0.27
Owens home	0.74	0.72	0.73
Pucca house	0.72	0.92	0.84
Water facility in home	0.41	0.62	0.53
Household toilet	0.26	0.61	0.46
Asset level < 3	0.39	0.12	0.23
Asset level = 3	0.28	0.21	0.24
Asset level = 4	0.19	0.30	0.26
Asset level > 4	0.13	0.36	0.26
Math score	-0.01	0.61	0.35
N	3,056	1,195	4,251

Notes: Table reports population average characteristics of households stratified by sector of school choice. Asset level refers to total of six possible asset indicators (owns home, pucca house, ≥ 2 covered rooms in home, water facility in home, household toilet, owns land)

Table 2 presents comparisons of key characteristics of government and private schools in our sample. The total tuition and fees for attending government schools is essentially zero, but averages 1,800 Rs. per year for private schools. For comparison, median household expenditure in comparable rural villages (with a private school) of Andhra Pradesh is around 86,000 Rs. (24,000 per capita) per the 2011-12 India Human Development survey. Thus, the average private school's tuition would account for over 7% of per-capita consumption for the median household, and an even higher share for poorer households. There is also considerable variation across private schools in tuition (the sample standard deviation is about 1,000 Rs.). English medium private schools are around 1,000 Rs. more expensive on average.¹⁸

¹⁸Appendix Table A20 provides additional information about the distribution of characteristics among private

Government and private schools also differ in terms of amenities and teacher characteristics. A larger share of government schools are pucca structures and nearly all include a library, whereas a greater share of private schools have a functioning toilet and a staffroom for teachers. Government schools also provide mid-day meals for students, which private schools do not. A much greater share of teachers in government schools have a bachelor’s degree or a formal teacher training credential. Private school teachers, on the other hand, are more likely to be female, from the village, and present in the classroom during enumerator visits. Private schools also have less multi-class teaching. Government school teacher salaries are also much higher (over 5 times greater) than those of private schools. Overall, government schools have better qualified and paid teachers; whereas private schools hire more teachers and have greater teacher effort.

Government and private schools likewise differ in instructional medium and subjects taught. While nearly all government schools teach in the local language (Telugu), about 61% of private schools feature instruction in English. Moreover, many private schools teach Hindi and computer skills (whereas few government schools do). Though not shown in the table, private schools also have a longer school year (2 weeks additional instructional time) and a longer school day (45 minutes more per day). Further, there are significant differences in allocation of instruction time with private schools spending less time on math and Telugu and using the extra time to teach additional subjects (see Muralidharan and Sundararaman (2015) for details).

In addition to the school characteristics gathered by project enumerators, Table 2 compares government and private schools in terms of school value-added – a proxy for schools’ human capital return. We estimate school value-added using baseline and two follow-up scores in math.¹⁹ Appendix B provides full details regarding the estimation of value-added. The table shows that government schools on average are around 0.05 standard deviations (on the student test score distribution) below average, while average private school value-added is 0.06 standard deviations above average. Notably, the differences in value-added between private and government schools on average in Table 2 are considerably narrower than the average differences in students’ test scores (shown in Table 1). In addition, Appendix Table A20 shows that, while English instruction, facilities, and teacher qualifications tend to be positively related to tuition, value-added is very weakly correlated with tuition among private schools.

schools.

¹⁹We use math scores so as not to confound school quality with differences across schools in language of instruction.

Table 2: Characteristics of Primary Schools

	Government	Private
Tuition and fees (Rs.)	1.52	1,814
English medium	0.02	0.61
Mid-day meals	0.99	0.04
Full pucca building	0.90	0.48
Library	0.94	0.77
Functioning toilet	0.64	0.82
Separate toilet for girls	0.31	0.59
Staffroom for teachers	0.20	0.70
Has secondary school	0	0.25
Average teacher salary (Rs. / month)	13,051	1,612
Multi-class teaching	0.65	0.27
Pupil-teacher ratio	26.05	16.07
Share teachers absent	0.23	0.10
Share teachers with BA	0.78	0.57
Share teachers with formal credential	0.89	0.16
Share teachers female	0.53	0.71
Share teachers from village	0.23	0.47
Offers Hindi	0	0.42
Teaches computer skills	0.01	0.13
School value-added	-0.05	0.06
N	352	293

3 Empirical Models

We describe our empirical models of household school choice in this section. 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 the influence of an 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 for identification places structure on how observed measures of household wealth influence choices, with flexible logit demand models that are similar to models of school choice that have been applied in other contexts (e.g. Neilson 2013; Carneiro et al. 2016; Pathak and Shi 2017; Abdulkadiroglu et al. 2017a).

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, captured in the model through a constraint on their choice problem:

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

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 Y_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.²¹

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:

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

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

²¹This “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.

D_{ij} is the distance between school j and household i 's home, while X_j represents school characteristics, such as whether a school is government or private, is English medium, the facilities, characteristics or the quality of teachers, and other school amenities observed in the data. Among these is school j 's value-added.²² ξ_j is regarded as an index of commonly-valued amenities of school j that remain unobserved to the econometrician. ϵ_{ij} represents idiosyncratic preferences that drive school choice and is assumed to follow a Type 1 extreme value distribution.

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

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

W_i includes demographics, such as gender, caste, religion, as well as parental education. These characteristics mediate the valuation households place on school amenities, capturing systematic heterogeneity across households in willingness-to-pay. This allows, for instance, that better educated households may place greater value on English instruction or on test score value-added.²³

A fundamental empirical challenge for estimating the choice problem described by equation (1) is that households' ability-to-pay, Y_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. As a result, we specify latent ability-to-pay as a function of observed household wealth factors, given by:

$$\ln Y_i = I_i' \lambda + v_i \tag{3}$$

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 v_i . We maintain the assumption that v is distributed normally and independent of the choice shocks. We discuss the identification of this model, which relies in part on placing structure on the way that observed household characteristics influence choice patterns, in the next section.

²²The estimation of school value-added assumes that it is common across all students. Note, however, that this does not imply that the choice models assume that households *believe* that value-added is homogeneous. The parameters of the choice models will capture systematic differences in both preferences over school attributes as well as in beliefs or information about those attributes.

²³We do not include measures of students' abilities in the models. This is due mainly to data limitations (the baseline test scores collected for first graders are endogenous to their choice), but we expect any bias from this to be limited given the characteristics that are included (e.g. parental education). There is also limited evidence that parents have good information about the abilities of their children (e.g. Dizon-Ross 2019). In any case, the experimental validation will allow us to test for ability sorting as a source of mis-specification, as discussed further in Section 5.

3.2 Logit Demand

We compare the latent ability-to-pay model with more familiar logit demand models of school choice. In this class of models, the underlying choice problem is unconstrained – households are able to choose from any primary school in their village:

$$\max_{j \in \mathcal{V}_i} U_{ij} \geq U_{ij'} \quad (4)$$

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

$$U_{ij} = -\alpha_i p_j + X_j' \beta_i + \gamma_i \ln D_{ij} + \xi_j + \epsilon_{ij}$$

While similar, this indirect utility differs from the ability-to-pay constrained model in two ways. First, the function allows for heterogeneity across households in their sensitivity to higher tuition and fees. In the specifications we estimate, for example, we allow this coefficient to depend on observed indicators of household wealth. Second, the logit demand models accommodate greater flexibility in how households value school characteristics. We compare two alternative specifications for preference heterogeneity, described below.

3.2.1 Clustered Logit

The first logit demand model that we estimate, which we term the clustered logit, takes a semi-parametric approach to modeling preference heterogeneity. The model is similar to the approach taken in Abdulkadiroglu et al. (2017a) to estimating household preferences over New York high schools.

We bin the households in the dataset into M groups of observably similar households based on their characteristics (e.g. demographics, assets, and parental variables). The clustered multinomial logit model then allows the preferences for school characteristics to differ for each group:

$$U_{ij} = -\alpha(I_i) p_j + X_j' \beta_{m(i)} + \gamma_{m(i)} \ln D_{ij} + \xi_j + \epsilon_{ij} \quad (5)$$

where U_{ij} is the indirect utility of school j for household i , a member of group m . The model flexibly captures heterogeneity in preferences across households by allowing each group m to have its own γ and β vector.²⁴ Under the assumption that observably similar households also have correlated unobserved characteristics, this approach to modeling demand also absorbs unobserved heterogeneity that influences choice patterns.

²⁴We make the restriction that the variance scale of the choice shocks is the same across clusters.

To limit the computational demands of estimating this model and to guard against overfitting, we use k-means to group the households into a finite number of clusters, chosen using the AIC criterion which adjusts the maximized likelihood for degrees of freedom.

3.2.2 Random Coefficient Logit

We also estimate random coefficient models similar to classic demand estimation applications (e.g. Berry et al. 1995; Nevo 2001; Petrin 2002) and the models of school choice in Neilson (2013) and Carneiro et al. (2016).

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

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

where W_i includes observed household characteristics. However, the random coefficient model we estimate 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:

$$\beta_i^P = \beta_1^P + \beta_2^P W_i + \nu_i \tag{6}$$

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

4 Identification and Estimation

This section discusses identification and estimation of the empirical models. For identifying households' willingnesses-to-pay, we face a challenge commonplace to differentiated demand applications using observational data: unobserved school quality, ξ_j . We describe the approach we take, which includes instruments for tuition, in the first subsection. In the second subsection, we discuss empirically separating the effect of ability-to-pay on choice from willingness-to-pay in the constrained model. We outline the details of the empirical specifications and estimation in the third subsection.

4.1 Instrumenting for Private School Tuition and Fees

As there may remain qualities that differentiate schools that remain unobserved, we instrument for private school tuition and fees as part of a control function approach (Petrin and Train, 2010).²⁵ This strategy regresses tuition and fees on school characteristics and a set of instruments in a “first stage.” The residuals plus a random effect term, which jointly represent the unobserved school characteristic, ξ_j , are then included in the indirect utility function. This is described in expanded detail in the Appendix.

To implement this, we assume that the tuition and fees, p_j , at private school j at time t are given by:

$$p_j = X_j' \Gamma + f(Z_j) + \mu_j \quad (7)$$

where X_j are observed school characteristics (including the estimated value-added) and Z_j are instruments for tuition and fees. μ_j represents the unobserved school characteristics that influence school tuition-setting, where $E[\xi_j \mu_j] > 0$. This approach maintains that μ_j is uncorrelated with p_j conditional on X_j and Z_j . Identification also requires that Z_j only influences school choice through the impact on private schools’ setting of tuition and fees.

We use two instruments. First, assuming Bertrand-Nash competition and that non-tuition school characteristics are fixed, each private school’s location in “product space” provides an exogenous source of variation in mark-ups (Berry et al., 1995). We thus total the characteristics of *other* schools in the village for each private school, e.g. the number of other English-medium schools. For efficiency, however, we use factor analysis to construct a summary measure from these totals as the first instrumental variable.²⁶ The second instrument we use in our baseline specifications 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. In implementation, we match private schools within medium of instruction and focus on other schools not in nearby villages.²⁷ The “first stage” F statistic with these instrumental variables, reported in Appendix Table A19, is 12.51.

²⁵We take a control function approach as the sampling design prevents us from computing market shares for all schools.

²⁶Specifically, we use the first underlying factor constructed from the standard “BLP” instruments. This is to address the well-known weakness of these instruments when used simultaneously (e.g. Gandhi and Houde 2016).

²⁷This is to minimize the confounding influence of spatially-correlated demand shocks. As an alternative to the “Hausman” predicted tuition instrument, we also estimate models that include the product space IV and a cost index instrument that we construct from private schools’ reported costs. The disadvantage of the cost index IV is that it does not vary across private schools in the same village. Appendix Table A23 compares key estimates and predictions obtained with the different sets of instruments.

4.2 Identification of Ability-to-Pay

The constrained choice model that we estimate models primary school choice as subject to households' ability-to-pay. In this subsection, we discuss restrictions placed on how observed characteristics of households influence choice patterns that help identify the effect of this constraint.

It is helpful to consider the probability that a household can afford to attend private school j^* in their village but not $j^* + 1$, whose tuition exceeds their ability-to-pay (where j^* indexes schools in the village in ascending order of tuition and fees). Under the maintained assumptions, this probability can be expressed as:

$$\begin{aligned}\phi_{ij^*} &= P(p_{j^*} \leq Y_i < p_{j^*+1}) \\ &= \Phi\left(\frac{\ln p_{j^*+1} - I'_i \lambda}{\sigma}\right) - \Phi\left(\frac{\ln p_{j^*} - I'_i \lambda}{\sigma}\right)\end{aligned}$$

This expression highlights that two data elements shift these state probabilities across households: First, some variation arises from differences in the tuition and fees distribution across villages. Second, I_i , the wealth factors in equation (3), will influence the likelihood a household is unable-to-pay for a given private school. As a result, it is important that I_i and W_i do not overlap completely to empirically separate the effect of ability-to-pay on choice behavior from that of willingness-to-pay.

For this reason, we apply the following exclusion restriction: we assume that assets (such as home ownership) are ability-to-pay factors that, all else equal, do not shift preferences across schools. These variables thus only enter I_i , being used exclusively to proxy for household wealth and ability to borrow (Filmer and Pritchett, 2001). In practice, we use the first principal factor obtained from the six household assets in our data. Intuitively, this restriction ascribes differences in choice patterns between households that are similar in their non-asset characteristics (e.g. demographics and parental education) but who differ in assets to differences in ability-to-pay. In our baseline specification, demographics and parental education only shift preferences; however, we estimate several alternative specifications, outlined in the next subsection, for sensitivity checks. Our baseline specification allows both preferences and ability-to-pay to depend on AP voucher eligibility and the student's number of siblings, a measure of household size.

4.3 Estimation and Specifications

This subsection describes the estimation of the models, including specification details. For all of the models, X_j includes whether a school is private or not, the tuition and fees, whether English

medium, school value-added, whether Hindi is offered, and facility and teacher characteristics.²⁸ For the latter two, we use first principal factors, one for the facility variables and one each for teaching quality (e.g. multi-class) and teacher characteristics (e.g. share female) to keep the models parsimonious. In addition to travel distance, we include an indicator for whether a school is the closest government school to the household in all of the models.

On the household side, we model heterogeneity in preferences according to whether the student is female, lower caste, Muslim, parental education attainment, whether an older sibling attends government school, and their number of siblings. We also allow for preference heterogeneity based on whether the student was eligible for the AP program voucher in all models. Kindergartners were eligible if attending a government daycare facility (Anganwadi) at baseline, but this variable is unknown for students in the first grade cohort. To address this, we treat their eligibility as a latent variable. We allow the probability of eligibility to depend on observed household characteristics and use the EM algorithm for estimation. Appendix D presents the likelihood functions and describes the estimation in greater detail, while Appendix Table A22 summarizes the school characteristics and heterogeneity in the baseline model specifications.

We focus on comparing the estimates and predictions for three baseline choice models: the clustered logit, a random coefficient demand model, and the ability-to-pay constrained choice model. The clustered logit model bases groupings of households on principal factor similarity using k-means, allowing preference parameters to be group-specific. This algorithm for grouping households treats all of the observed household information (i.e. demographics, parental education, as well as household assets) symmetrically. The number of groups that minimizes the AIC criterion is 51. For the clustered and random coefficient models, we specify $\alpha(I_i)$ as a function of discrete household asset levels (e.g. household owns three assets, etc.), total siblings in the household, and voucher eligibility.²⁹

As noted earlier, the baseline constrained model specifies ability-to-pay as a function of the first principal component of household assets as well as total siblings. We also include an indicator for eligibility for the AP voucher in the function, allowing ability-to-pay to be correlated with whether a student attended a government daycare. In addition to this specification, whose results we present and discuss in the text, we also estimate several alternatives. These are summarized in Appendix Table A23 and focus on robustness to the assumptions regarding how household characteristics influence preferences as opposed to ability-to-pay. For example, we estimate a specification that

²⁸We also include an indicator variable for whether a private school is officially recognized by the government and indicators flagging when distance and value-added are missing.

²⁹Note that the baseline random coefficient specification, similarly to the constrained model, restricts asset information to interact only with tuition in the utility function. In contrast, the clustered model uses all information (including assets) to define household groupings.

includes whether a household is lower caste in I_i (as well as in W_i). We also examine sensitivity of the results and predictions to the instruments and to the weighting for attrition.

5 Results

We present the results from estimating the empirical models on the control market data in this section. We begin with estimates of the price elasticity of demand for private schooling, willingness-to-pay estimates, and goodness-of-fit comparisons. We then present estimates of welfare impacts for the Andhra Pradesh project voucher and for a universal voucher.

5.1 Parameter Estimates and Elasticity of Demand

Estimates of the parameters, elasticity of demand, and households' ability-to-pay for private schooling are presented in this subsection. Table 3 reports estimates and standard errors for the coefficient on tuition and fees in the utility function. The columns report estimates for baseline specifications of the clustered logit ("CMNL"), random coefficient ("RC"), and the ability-to-pay constrained ("CC") models. The coefficient on tuition and fees depends on the household's asset level (and household size) in the non-constrained logit demand models, but is constant across all households in the constrained model (in which assets influence ability-to-pay).

The estimates in Table 3 show that the parameter on tuition and fees is negative for all households across each of the baseline model specifications. For the clustered multinomial logit and random coefficient models, sensitivity to fees tends to decrease with the household's asset level; households with fewer assets are more sensitive to the tuition and fees at private schools in their choice. The estimates also indicate that, all else held equal, students who were eligible for the AP voucher are also less sensitive to tuition and fees. The results in Table 3 also underscore the role of instrumenting for endogenous tuition: unobserved school quality (ξ_j) is controlled for by the combination of the first stage residual and private school-specific random effect. The estimate on the residual is significant and positive across all models, indicative of positive correlation between unobserved school quality and tuition and fees.

We present estimates of the average elasticity of demand for private schooling by subgroup in Table 4. Elasticities of demand are computed by calculating the percent change in private school attendance due to a simulated 1% reduction in private school tuition. Overall, demand for private schooling is inelastic (i.e. the elasticity is less than 1) according to all three sets of estimates. The estimates reveal an elasticity of 0.65 for the clustered multinomial logit model, indicated by the CMNL column. The random coefficient model and ability-to-pay constrained models estimate a

Table 3: Estimates: Coefficient on Tuition and fees

	CMNL	RC	CC
Tuition and fees (1000s of Rs.)	-2.00 (0.24)	-2.35 (0.28)	-1.28 (0.58)
× Asset level = 2	0.22 (0.14)	0.45 (0.20)	
× Asset level = 3	0.45 (0.15)	0.74 (0.20)	
× Asset level = 4	0.68 (0.16)	1.12 (0.20)	
× Asset level > 4	0.39 (0.17)	0.81 (0.21)	
× Eligible for AP voucher	0.37 (0.10)	0.07 (0.12)	
First stage residual	1.48 (0.18)	1.60 (0.20)	1.77 (0.63)
Random effect σ	1.73 (0.22)	2.23 (0.22)	2.66 (0.27)

Notes: Table presents select parameter estimates as estimated by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Estimates correspond to baseline specifications described in the text. CMNL and RC models also include interactions of total siblings in the household with tuition and fees, estimates for which are not reported in the table.

lower elasticity at 0.47.

Table 4: Estimates: Price Elasticity of Demand for Private Schooling

	CMNL	RC	CC
Overall	0.65	0.47	0.47
Female	0.73	0.52	0.52
Muslim	0.40	0.27	0.25
Lower caste	1.03	0.75	0.74
Older sibling in gov't school	1.15	0.84	0.92
Both parents completed primary school	0.22	0.17	0.15
≥ 1 parent completed secondary	0.18	0.13	0.11
Both parents laborers	1.04	0.72	0.69
Asset level < 3	1.40	1.04	0.79
Asset level = 3	0.69	0.49	0.54
Asset level = 4	0.33	0.24	0.36
Asset level > 4	0.25	0.15	0.21

Notes: Table presents average elasticity of demand for private schooling (% increase in private school attendance for a 1% reduction in tuition and fees) by subgroup as estimated by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Estimates correspond to baseline specifications described in the text.

The elasticity estimate reported in Table 4 exhibit several notable subgroup patterns as well. For one, demand for private schooling by households in which both parents completed primary school or at least one parent completed secondary is highly inelastic per all three models. For female students, the elasticity is somewhat higher than average. The model estimates are also in agreement that elasticity of demand decreases with household wealth (as indicated by assets). In fact, the clustered and random coefficient models estimate that the elasticity of demand exceeds 1 for households with fewer than three assets on average. At the same time, the table reveals that the elasticities estimated by the ability-to-pay constrained model tend to be smaller. This finding is consistent with a marginal tuition reduction (such as embedded in the demand elasticity calculations) doing little to relax constraints on choices.

Table 5 sheds more direct light on the role of constraints on ability-to-pay by presenting estimated probabilities that households are unable to pay for private schooling.³⁰ The ability-to-pay constrained model estimates that about 10% of all households are unable to pay for *any* private school in their village and that about 14% cannot pay for the private school nearest to their residence. The third column reports that around 83% of households cannot pay for the highest-tuition private school, suggestive that, although most households have the option of choosing private schooling, ability-to-pay still constrains choice to some degree. Table 5 also reports notable dif-

³⁰Appendix Table A24 reports the parameter estimates for the ability-to-pay function in the baseline specification.

ferences according to household demographics (e.g. lower caste households are more likely to be ability-to-pay constrained). The sharpest contrasts across households appear for different levels of assets: just 1% of households with at least five assets are estimated to be unable to pay for private schooling, for example. On the other hand, around 25% of asset-poor households (those with fewer than three assets) are estimated to lack the ability-to-pay for private school.

Table 5: Estimates: Share with No Ability-to-Pay for...

	Any	Closest	Highest-tuition
	Private School		
Overall	0.10	0.14	0.83
Female	0.10	0.14	0.83
Muslim	0.08	0.10	0.79
Lower caste	0.14	0.20	0.90
Older sibling in gov't school	0.12	0.17	0.87
Both parents completed primary	0.04	0.07	0.77
≥ 1 parent completed secondary	0.04	0.05	0.70
Both parents laborers	0.14	0.20	0.88
Asset level < 3	0.23	0.31	0.88
Asset level = 3	0.10	0.14	0.80
Asset level = 4	0.04	0.06	0.80
Asset level > 4	0.01	0.02	0.84
Choose gov't school	0.24	0.25	0.92

Notes: Table presents posterior estimated share of households by subgroup unable to pay for any, the closest, and the highest-tuition private school in their village, per the ability-to-pay constrained model (CC). Estimates correspond to baseline specification described in the text.

The bottom row of Table 5 summarizes (in)ability-to-pay among those specifically targeted by the AP voucher’s design: those who (*ex post*) chose to attend a government school in the control markets (i.e. in the absence of a voucher offer). As the table shows, 24% of households among this group are estimated to be unable to choose any private school, with 25% lacking the ability-to-pay for their closest private school according to the constrained choice model. Nearly all of these households are estimated to be unable to choose the most expensive private school in their village. One useful point of comparison for these estimates are households’ stated reasons for the school choice: As noted earlier, 41% of government school students in the data indicate “economic reasons” – which would include credit and liquidity constraints – for their choice of school.

5.2 Willingness-to-Pay Estimates

The model estimates allow us to compute willingness-to-pay (in terms of 1000s of Rs.) for characteristics of primary schools. In addition to revealing household preferences over schooling charac-

teristics, differences across households in willingness-to-pay will translate into different predictions regarding take-up of a voucher offer (and where students will use a voucher, if they do indeed take it up).³¹

Table 6 reports willingness-to-pay estimates by subgroup for English medium instruction, school value-added, and private schooling.³² Overall, the ability-to-pay constrained model estimates households’ average willingness-to-pay for English instruction for the duration of primary school (all else held equal) at around 3,350 Rs. The estimated willingness-to-pay is lower for female students and considerably greater among highly educated households. Willingness-to-pay for English medium schooling also increases with household assets. Though the subgroup patterns are qualitatively similar, the random coefficient model yields generally smaller estimates of willingness-to-pay for English instruction than the constrained choice model. The random coefficient model estimates indicate an average willingness-to-pay of just under 3,000 Rs. The clustered multinomial logit estimates differ from the other models in several ways. In particular, the average overall is lower still (around 2000 Rs.) and willingness-to-pay is actually negative on average among students with siblings in government schools, those whose parents are laborers, and lower asset households.

Table 6: Estimates: Willingness-to-Pay (1000s of Rs.) for Primary School Characteristics

	English Medium			School Value-Added			Private Schooling		
	CMNL	RC	CC	CMNL	RC	CC	CMNL	RC	CC
Overall	2.02	2.92	3.35	1.24	1.29	1.49	10.47	14.36	15.96
Female	0.48	1.54	1.74	1.08	1.39	1.57	8.69	11.33	14.11
Muslim	6.28	6.72	7.69	0.58	1.01	1.14	17.48	24.42	24.05
Lower caste	1.40	2.05	2.73	0.53	1.37	1.48	2.00	1.57	5.22
Older sibling in gov’t school	-0.77	2.29	2.64	1.65	1.30	1.50	-2.06	-0.61	1.11
Both parents completed primary	6.23	6.17	6.96	0.66	1.38	1.60	23.09	31.40	31.71
≥ 1 parent completed secondary	6.65	7.03	8.00	0.35	0.50	0.71	25.27	35.40	35.07
Both parents laborers	-1.55	1.61	1.99	0.99	1.36	1.57	0.01	3.41	5.86
Asset level < 3	-0.91	1.47	2.29	0.98	1.04	1.55	1.33	1.40	4.94
Asset level = 3	0.36	2.43	2.99	1.16	1.30	1.53	7.11	9.83	12.97
Asset level = 4	3.96	4.02	3.72	1.67	1.62	1.48	14.43	20.52	19.35
Asset level > 4	4.32	3.61	4.28	1.14	1.19	1.41	17.93	24.22	25.35

Notes: Table presents average willingness-to-pay (1000s of Rs.) by subgroup for English medium instruction, one standard deviation increase in school value-added, and private schooling as estimated by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Willingness-to-pay for private schooling calculated as change from average government to average private school. Estimates correspond to baseline specifications described in the text.

The middle three columns of Table 6 report households’ estimated marginal willingness-to-pay

³¹Appendix E provides additional detail regarding calculating willingness-to-pay.

³²Appendix Table A21 reports willingnesses-to-pay for a 1 mile reduction in travel distance to school by model and subgroup. Note that these disutilities of travel will implicitly “price” schooling options under the simulated voucher treatment, wherein tuition and fees of private schools are effectively set to zero.

for school value-added (above and beyond other school characteristics). Willingness-to-pay for a 1 standard deviation increase (on the student test score distribution) in value-added during primary school is around 1,300 Rs. on average according to the clustered and random coefficient models. This compares with 1,500 Rs per the ability-to-pay constrained model. The constrained model estimates also indicate important differences from the other models for particular subgroups. The clustered logit model, for example, estimates that lower caste household are only willing to pay 530 Rs. In comparison, the ability-to-pay constrained model estimates this subgroup’s willingness-to-pay as nearly three times larger (1,480 Rs.). Similarly, the clustered and random coefficient models estimate that asset-poor households generally have lower willingness-to-pay than do higher wealth households on average. In contrast, the ability-to-pay constrained model estimates that low asset households have more or less an equivalent (if not higher) preference for school quality as higher asset households. These results suggest that unobserved ability-to-pay constraints may confound estimates of how lower income households value school quality.

Estimates of the willingness-to-pay for private schooling are presented in the last three columns of Table 6. We calculate willingness-to-pay for private schooling as the value of switching from the *average* government school in terms of (non-travel) amenities, X , to the *average* private school. In this way, the estimates reflect the differences on average between private and government schools summarized in Table 2. The clustered multinomial logit model estimates an average willingness-to-pay for private schooling of around 10,500 Rs. at the population level. The random coefficient model estimates average willingness-to-pay for private schooling of more than 14,000 Rs. Willingness-to-pay per the ability-to-pay constrained model, however, is largest at about 16,000 Rs. These estimates are economically large (representing up to 10 times annual private school tuition), corresponding to the average for a population most of which does (and can) choose to attend private school. The subgroup level estimates in Table 6 (and their differences across models) are therefore of interest for understanding willingness-to-pay for private schooling by those households targeted by the AP voucher. For example, the clustered logit model estimates zero willingness-to-pay among households in which both parents are laborers, whereas the ability-to-pay constrained model estimates a positive 6,000 Rs. for this group. Similarly, the clustered and random coefficient models estimate little willingness-to-pay for private schooling among asset-poor households. The ability-to-constrained model, in contrast, indicates a willingness-to-pay for private schooling of about 5,000 Rs. for this group. These differences across models have important implications for predicted take-up of the voucher offer (as well as for estimated welfare impacts).

5.3 Goodness-of-Fit

Before turning to welfare impacts and the experimental validation, this subsection assesses and compares the models in terms of goodness-of-fit to the control markets data. The maximum log-likelihood and AIC for each of the models are presented in Table 7. The statistics show that the clustered multinomial logit model achieves a much higher likelihood value and even the AIC (which adjusts for the many extra parameters) is substantially smaller than the goodness-of-fit statistics for the other models. In addition, the random coefficient model achieves a better fit to the control data than does the ability-to-pay constrained model according to these statistics.

Table 7: Goodness-of-Fit Statistics

	CMNL	RC	CC
$\ln L$	-5,332	-5,806	-5,821
AIC	11,640	11,741	11,767

Notes: Table presents value of log likelihood and AIC for the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Statistics correspond to baseline specifications described in the text.

We also assess goodness-of-fit by presenting predictions from the models side-by-side with corresponding moments from the control data. To focus these comparisons on moments not specifically “targeted” by the first order conditions of the maximum likelihood estimation, we examine predicted private school attendance of kindergarten students eligible for the AP project voucher. The empirical models will be validated against the choices of those among this group who applied for the AP voucher at baseline. Table 8 reports that 32% of eligible kindergartners choose to attend a private school in the control data. The models and specifications match this moment reasonably closely, though the clustered model underpredicts by four percentage points. Looking at household demographics, all of the models underpredict private school attendance by Muslim students (58% in the data). The models also tend to underpredict private attendance by laborer households (24% in the data) and asset-poor eligible households (23%) and, with the exception of the clustered logit model, overpredict private attendance by the highest asset households. As compared with the 39% observed in the data, the random coefficient and ability-to-pay constrained models estimate that, respectively, 42% and 45% of high asset eligible students attend private schools.

Table 8: Goodness-of-Fit: Private Schooling of AP Voucher Eligible Students

	Data	CMNL	RC	CC
Overall	0.32	0.28	0.31	0.30
Female	0.28	0.27	0.28	0.28
Muslim	0.58	0.46	0.54	0.53
Lower caste	0.20	0.17	0.20	0.20
Older sibling in gov't school	0.16	0.16	0.14	0.14
Both parents completed primary school	0.49	0.47	0.50	0.48
≥ 1 parent completed secondary	0.56	0.51	0.56	0.55
Both parents laborers	0.24	0.20	0.24	0.23
Asset level < 3	0.23	0.19	0.22	0.20
Asset level $= 3$	0.33	0.24	0.28	0.28
Asset level $= 4$	0.39	0.38	0.40	0.38
Asset level > 4	0.39	0.39	0.42	0.45

Notes: Table presents private school attendance for AP Project voucher-eligible kindergarten students by subgroup in the data and as predicted by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Estimates correspond to baseline specifications described in the text.

5.4 Welfare Impacts

We present estimates of the aggregate welfare impacts of private school vouchers and their sources in this subsection. We focus on two voucher programs: 1) the voucher actually offered to applicants in treatment markets by the AP School Choice project; and 2) a universal voucher that would pay primary school tuition and fees for all households.

5.4.1 The AP Project Voucher

The total social welfare generated by a voucher program is the sum of three component parts: 1) the gain in consumer surplus to recipients of vouchers; 2) less the expected cost of financing the program; 3) plus any fiscal gain that arises from re-allocating students out of government schooling. Note that this latter part accrues only for voucher recipients who would have otherwise attended a government school, whereas the program cost must account also for students who would have attended a private school regardless. To calculate the fiscal externality, we assume that two thirds of per pupil spending in government schools in Andhra Pradesh (8,390 Rs.) – the share of spending allocated to teachers (Dongre, 2012) – could be cut. Consumer surplus change is given by the added inverse of the estimated compensating variation.³³

³³Compensating variation is the amount of income that each household would need to be compensated to keep their utility level under the voucher the same. Mechanically, this is obtained by the Rs.-valued difference between expected utility with the voucher (when tuition and fees at private schools are set to zero) and without the voucher (as in the data). Appendix D provides additional detail. This exercise thus abstracts from any equilibrium adjustments by schools, consistent with the scale of the actual trial.

Table 9: Estimates: Welfare Impacts (1000s of Rs.) per AP Voucher Applicant

	CMNL	RC	CC
Gain in consumer surplus (A)	3.45	3.37	4.88
Cost of voucher (B)	4.83	4.63	5.53
Fiscal externality (C)	5.46	4.60	6.75
Net welfare change (A-B+C)	4.07	3.34	6.10

Table 9 reports estimated welfare impacts from the AP Project voucher, a program that was targeted to students likely to otherwise attend a government school. The first row reports the average gain in consumer surplus (in 1000s of Rs.) per applicant according to the three empirical models. The second row reports the expected voucher cost. The last row adds the fiscal externality to the difference between the first two lines to arrive at the average social welfare generated by the program. The total welfare per applicant generated by the AP voucher according to the logit demand models is 3,300 (the random coefficient model) to 4,000 Rs. (the clustered model). The ability-to-pay model instead estimates the gain in social welfare at over 6,000 Rs. Thus, while all three empirical models indicate that the AP voucher improves social welfare in the aggregate, the ability-to-pay model indicates that the magnitude of the gain is up to 80% larger in magnitude. Table 9 also suggests one important source of this difference in aggregate welfare impacts: consumer surplus generated per applicant is around 40% larger according to the constrained model estimates.

While Table 9 reports impacts per applicant to quantify the total welfare change, not all applicants use the voucher and, moreover, the models differ in what fraction of households they predict will use it (a prediction that we will return to for the experimental validation). Table 10 instead reports estimates of average consumer surplus generated by the AP Project voucher *per recipient*. The recipients represent households “treated” by the offer in that they use the voucher to attend a private school. On average, the ability-to-pay constrained model estimates that each voucher recipient gains almost 8,000 Rs. in consumer surplus. The logit demand models estimate that each recipient gains a little less than 6,500 Rs. on average. In addition to differences in magnitude, Table 10 also reveals an important difference between models in the slope with respect to assets: the logit demand models generally show that the surplus gain increases with assets, whereas this pattern is exactly reversed when unobserved differences in ability-to-pay across households are accounted for.

In Table 11, we further decompose the aggregate welfare impacts per recipient into two sub-groups of particular interest: a) for those who would have attended private school regardless of the voucher offer (“always takers”); and b) for those recipients who would have attended government

Table 10: Estimates: Consumer Surplus (1000s of Rs.) per AP Voucher Recipient

	CMNL	RC	CC
Overall	6.39	6.37	7.81
Female	6.45	6.34	7.63
Muslim	8.05	8.18	9.77
Lower caste	6.02	6.01	7.48
Older sibling in gov't school	5.65	5.48	6.46
Both parents completed primary school	7.41	7.32	8.73
≥ 1 parent completed secondary	7.70	7.66	9.09
Both parents laborers	5.99	5.99	7.66
Asset level < 3	5.44	5.61	8.36
Asset level = 3	6.13	6.16	7.49
Asset level = 4	7.21	7.23	7.71
Asset level > 4	7.14	6.78	7.48

school otherwise (“compliers”). Intuitively, the voucher offer is valued similarly to its cash value by always takers. This is reflected in Table 11 by the fact that the gain in consumer surplus less the voucher cost is essentially zero. These values are not identically zero, however, because in some cases always takers attend a more expensive private school than they otherwise would have and, in the case of the ability-to-pay constrained model, there is some surplus generated from an expanded choice set. Each always taker also generates no fiscal externality. For compliers – those induced into private schooling by the voucher offer – the contrast between models in their normative implications is especially sharp: the logit demand models estimate that the average surplus to each complier is less than 5,000 Rs. This value is much less than the surplus that accrues to the average always taker. On the other hand, the ability-to-pay constrained model estimates the surplus gain to the average complier to be 8,000 Rs. (60% larger in magnitude). The average value to compliers is estimated to be larger than the average gain to the always takers by the ability-to-pay model, highlighting the impact of the constraint on their choice in the absence of the voucher.

Table 11: Welfare Impacts (1000s of Rs.) per AP Voucher Recipient

	Always Takers			Compliers		
	CMNL	RC	CC	CMNL	RC	CC
Share of applicants	0.30	0.32	0.32	0.24	0.21	0.30
Gain in consumer surplus (A)	7.75	7.54	7.81	4.78	4.66	7.99
Cost of voucher (B)	8.17	8.04	8.03	9.98	10.04	9.82
Fiscal externality (C)	0	0	0	22.42	22.44	22.35
Net welfare change (A–B+C)	-0.42	-0.50	-0.22	17.22	17.06	20.52

5.4.2 A Universal Voucher

In this subsection, we consider the aggregate welfare impacts of a voucher offered universally to all households. As with the AP voucher in practice, this counterfactual assumes that private schools would be unable to screen or otherwise select voucher students and abstracts from general equilibrium adjustments, in particular supply responses by schools.³⁴

Table 12 reports welfare impacts per person from a universal voucher. The pattern of findings across models mirrors findings regarding the AP voucher: the total surplus change from a universal voucher is estimated to be positive according to all three empirical models, but the quantitative magnitude of the gain is roughly twice as large according to the ability-to-pay constrained model (4,000 Rs. vs. 2,000 Rs. per person).

Table 12: Estimates: Welfare Impacts (1000s of Rs.) per Household from Universal Voucher

	CMNL	RC	CC
Gain in consumer surplus (A)	5.24	5.04	6.30
Cost of voucher (B)	6.31	6.08	6.66
Fiscal externality (C)	5.46	4.60	6.75
Net welfare change (A–B+C)	2.16	1.98	4.01

To understand the sources of aggregate gains, we decompose the welfare impacts from a universal voucher between always takers and compliers in Table 13. Around three fifths of households attend private schools regardless, but the voucher offer may induce some always takers to attend more expensive private schools than otherwise. On net, the result is about negative 600 Rs. for the average always taker according to the logit demand models. In the constrained model, the effect of re-allocating some always takers' choice among private schools is mitigated by the value of expanding the choice sets of others. Turning to the complier households, the logit demand models estimate their value from the voucher as far less than the expected cost and less than the average value to the always takers. The ability-to-pay constrained model finds the opposite relationship: it is the compliers who place relatively more value on the voucher offer. The constrained model estimates an average willingness-to-pay for the voucher by compliers that is around 80% greater than the unconstrained demand models' estimates.

³⁴Such responses to a universal voucher at scale are likely to include adjustments to tuition, school quality, peer composition, program participation, and entry/exit. Several papers focused on diverse settings have examined these and other dimensions of school competition (e.g. Bau 2017; Andrabi et al. 2017b; Singleton 2017; Sánchez 2018; Allende S.C. 2019).

Table 13: Estimates: Welfare Impacts (1000s of Rs.) per Universal Voucher Recipient

	Always Takers			Compliers		
	CMNL	RC	CC	CMNL	RC	CC
Share of population	0.59	0.58	0.58	0.14	0.13	0.20
Gain in consumer surplus (A)	7.40	7.22	7.61	5.36	5.39	9.28
Cost of voucher (B)	7.99	7.82	7.90	10.26	10.40	10.23
Fiscal externality (C)	0	0	0	22.42	22.44	22.35
Net welfare change (A-B+C)	-0.59	-0.60	-0.28	17.51	17.43	21.40

5.5 Sensitivity

Beyond the baseline set of estimates discussed above, we estimate model specifications that instead use a cost index instrument (in place of the cost proxy). We also specifically estimate several ability-to-pay constrained specifications. The full suite of models we estimate are described in Appendix Table A23, including comparisons along key dimensions with the baseline specifications.

The alternative instrument set generates very similar estimates of the overall price elasticity of demand to the baseline specifications. For example, the elasticity is 0.65 in the baseline clustered logit model, but 0.70 using the cost index instrument instead. At the same time, accounting for tuition endogeneity is important for producing sensible results. As Appendix Table A23 shows, the ability-to-pay constrained model otherwise generates an opposite-signed elasticity of demand, suggesting the private enrollment declines when prices decrease (this is reflected also in the parameter estimates shown in column (3) of Table A24). Table A23 also shows that controlling for AP voucher eligibility in the estimation is important for achieving a good fit to the control data; the constrained model specification that ignores this heterogeneity overpredicts private school enrollment of eligible students by 10 percentage points.³⁵

The ability-to-pay constrained model estimates are not very sensitive to alternative assumptions to the baseline specification regarding how household characteristics influence choice patterns. For example, models that allow the constraint to depend also on lower caste status or allow preferences and the ability-to-pay constraint to depend on whether the parents are laborers produce nearly identical values for the share of students fully constrained in their school choice. Correspondingly, the estimates for welfare and fit are likewise essentially the same as for the baseline specification. One modeling choice that does affect the salience of the ability-to-pay constraint is whether the model includes cohort-level heterogeneity, but the elasticity and surplus estimates from this model

³⁵Though we do not report estimates from the corresponding specifications, these points regarding the roles of the control function in producing plausible elasticities and of controlling for eligibility for fit apply equally to the unconstrained logit demand models.

are nonetheless largely unchanged.

6 Experimental Validation

In this section, we present predictions used to validate the empirical models against the experimental outcomes of the AP School Choice project. To do so, we simulate a private school voucher that mimics the eligibility and conditions of the program in the control markets using the models. The predictions are generated by setting tuition and fees at private schools for applicants to zero in the models and simulating their choice of school in that counterfactual. In this simulated treatment condition, travel costs (relative to school quality) “price” the alternative schooling options. This allows us to generate model-based “treatment” moments to compare directly with analogous moments calculated from the treatment group.

In addition to the aggregated predictions, we present plans to validate the models against the specific choices made by those households randomly assigned the voucher in the treatment villages. We plan to test for model mis-specification in a hypothesis testing framework. This framework will allow us to understand how and why the models may fail to reproduce the experimentally-generated choice patterns. For instance, we develop tests to examine whether mis-specification may be due to voucher recipients gathering additional information about schools before choosing or to treated households’ excitement at being randomly selected.

As an initial step to the experimental validation, the treatment village sample is likely to require re-weighting for attrition. Analogous to the adjustment for attrition of kindergartners from the control sample, we will estimate the probability that each treatment village household attrits as a function of their observed characteristics.³⁶ We will then ensure that the re-weighted treatment and control samples are balanced on observed characteristics, consistent with the village-level randomization.

6.1 Predictions

We focus on model predictions for take-up of the voucher offer overall and across subgroups of households. The models facilitate generating predictions in one of two ways: First, simulating the voucher trial in the models directly generates predictions for the share of households who take-up the offer to attend a private school. These predictions are reported in Table 14. The table includes a column labeled “RCT” that will report the corresponding moments obtained from the treated data. Second, voucher use can alternatively be forecast based on the models’ implied voucher

³⁶Including whether or not they randomly received the voucher offer.

elasticities. This method corrects for the fact that the models' fit to the control sample is not perfect and differs. We report voucher elasticities of private schooling by subgroup in Table 15

As Table 14 reports, the simulations for voucher take-up predict that 54% and 53% of applicant households overall will choose to use the voucher according to the clustered and random coefficient logit models predict, respectively. In contrast, the ability-to-pay constrained model predicts that 63% will take-up the offer. Differences in predicted take-up between models tend to be smaller among higher educated and higher wealth households and somewhat more pronounced for low asset households. The ability-to-pay constrained model predicts that 60% of asset-poor households will use the voucher, a share that is 9 to 11 percentage points higher than the other predictions made by the logit demand models. The corresponding treatment moments will allow for comparing the mean squared errors of the models across the subgroups reported in Table 14.

Table 14: Validation: Predicted Take-Up of Voucher Offer

	RCT	CMNL	RC	CC
Overall		0.54	0.53	0.63
Female		0.53	0.51	0.60
Muslim		0.69	0.76	0.81
Lower caste		0.43	0.42	0.52
Older sibling in gov't school		0.38	0.31	0.41
Both parents completed primary school		0.71	0.70	0.76
≥ 1 parent completed secondary		0.71	0.74	0.80
Both parents laborers		0.46	0.46	0.57
Asset level < 3		0.51	0.49	0.60
Asset level = 3		0.50	0.52	0.62
Asset level = 4		0.57	0.54	0.62
Asset level > 4		0.63	0.61	0.69

Notes: Table presents average voucher take-up by eligible households by subgroup in the treatment data (RCT), and as predicted by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Predictions correspond to baseline specification described in the text.

Table 15 reports voucher elasticities of private schooling. The overall voucher elasticity is highest in the ability-to-pay constrained model, which predicts that private schooling will increase by 93% under the voucher. Mirroring the pattern for use, the overall elasticity is lower for the other models, but is relatively higher according to the clustered multinomial logit model. This reflects differences in the models' fit: though the clustered and random coefficient models predict the same overall use, that level reflects a bigger increase per the clustered model because that model underpredicts private schooling in the control markets. The elasticities in Table 15 also reveal a pattern not immediately apparent from the take-up predictions: the voucher elasticity decreases with household assets in all of the empirical models, but the slope is the steepest in the ability-

to-pay constrained model. For instance, the models predict relatively similar elasticities for high asset households, but the ability-to-pay constrained elasticity for asset-poor households is 20-50% larger in comparison. The ability-to-pay model estimates that asset-poor households will increase their private schooling over 170% under the voucher. We will use these elasticity estimates to also calculate mean squared errors for predicted voucher take-up.³⁷

Table 15: Validation: Voucher Elasticity of Private Schooling

	RCT	CMNL	RC	CC
Overall		82	63	93
Female		83	67	98
Muslim		44	28	38
Lower caste		125	93	138
Older sibling in gov't school		140	128	196
Both parents completed primary school		45	38	55
≥ 1 parent completed secondary		41	31	43
Both parents laborers		111	82	125
Asset level < 3		141	114	173
Asset level = 3		92	67	101
Asset level = 4		57	40	69
Asset level > 4		48	38	42

Notes: Table presents average voucher elasticity (percent change in private schooling due to the voucher offer) of eligible households by subgroup in the treatment data (RCT), and as predicted by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Predictions correspond to baseline specification described in the text.

The empirical models also make predictions for the characteristics or the qualities of schools that households will choose under the voucher. Table 16 reports intent-to-treat effects of the voucher offer on four characteristics of households' chosen school: the tuition and fees charged, whether the school teaches in English, the travel distance to school, and school value-added. The table reveals that the average tuition and fees otherwise charged at chosen schools will increase by 590 Rs. under the voucher on average according to the random coefficient model. The clustered model predicts a 660 Rs. increase, while the ability-to-pay model predicts 810 Rs. The table also reports the predictions in term of elasticities, which deflate the intent-to-treat effects by the baseline level to adjust for differences in fit to the control data. The ability-to-pay constrained model predicts that the share of students attending an English medium schools will more than double (112%) due to the voucher and that travel distance to school will increase by 20%. Similar to the pattern for the elasticity of private schooling, these elasticities are largest per the ability-

³⁷Elasticity-based predictions for use can be generated from the equation $P_T^m(X) = P_C(X) \times [1 + \hat{\epsilon}^m(X)]$, where $\hat{\epsilon}^m(X)$ is the percent change in private schooling due to the voucher according to model m and $P^C(X)$ are the (true) control data moments.

to-pay constrained model and generally smallest according to the random coefficient model. The clustered and random coefficient models predict that the value-added of chosen schools will increase by 0.03 and 0.02 standard deviations, respectively, whereas the ability-to-pay constrained model predicts an increase of 0.04.

Table 16: Validation: Voucher Intent-to-Treat Effects and Elasticities on Characteristics of Chosen School

	RCT		CML		RC		CC	
	ITT	ϵ	ITT	ϵ	ITT	ϵ	ITT	ϵ
Mean tuition and fees (1000s of Rs.)			0.66	122	0.59	104	0.81	145
Share choosing English medium school			0.18	92	0.16	79	0.22	112
Mean distance to school (mi.)			0.14	13	0.14	14	0.20	20
Mean school value-added			0.03		0.02		0.04	

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 eligible households in the treatment data (RCT), and as predicted by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay constrained models (CC). Predictions correspond to baseline specifications described in the text.

6.1.1 Sensitivity

The predictions for the experimental validation are generally robust across model specifications. For example, the alternative IV set yields a predicted take-up of 63% (as compared to 61% for the baseline specification) for the ability-to-pay constrained model. The overall voucher take-up and elasticity predictions of the random coefficient model is likewise similar across the two sets of instruments. However, the alternative IV set generates a voucher take-up of 56% – 2 points higher than the baseline IV set – for the clustered multinomial logit model. This translates to a voucher elasticity of 88, which is somewhat higher than the predicted voucher elasticity of the ability-to-pay constrained model with these instruments.

The alternative constrained model specifications also produce similar predictions to the baseline regarding take-up of the experimental offer and voucher elasticity; predictions for take-up range from 0.62 to 0.64 across those specifications that control for AP voucher eligibility. The predictions of the ability-to-pay constrained model are thus not very sensitive to the alternative modeling choices examined. The constrained model specification that ignores voucher eligibility heterogeneity produces a much larger predicted take-up (78%), but this is an artifact of this specification’s poor fit to the data; the predicted voucher elasticity from this specification is only 76.

6.2 Model Fit and Hypothesis Tests

The estimated models not only can be used to generate aggregated moments, they make specific predictions for each treated household’s indirect utility from each choice alternative in their village. These predictions provide a way to empirically validate the models using the voucher recipient’s actual choices as well as to understand the sources of model mis-specification via a series of hypothesis tests.

For those kindergartners who were offered the voucher in the treatment villages (i.e. were treated), we plan to estimate the following choice model:

$$U_{ij}^m = \hat{U}_{ij}^m + \delta^m Private_j \quad (8)$$

for each empirical model m estimated on the control sample. \hat{U}_{ij}^m is treated household i ’s predicted indirect utility from choice j , according to the estimates from model m based on the control data:

$$\hat{U}_{ij}^m = X_j' \hat{\beta}_i^m + \hat{\gamma}_i^m \ln D_{ij} + \hat{\xi}_j^m + \epsilon_{ij}$$

Note that tuition and fees (p_j) do not appear in the indirect utility because they are zero at all schools for the treated students. $\hat{\xi}_j^m$ is formed for the treatment village private schools using the parameters of the control function estimated on the control sample.³⁸

The null hypothesis to test for model validation is that:

$$H_0 : \delta^m = 0$$

In words, this tests says that model m accurately captures treated students’ take-up of the voucher offer (i.e. their preferences over private schooling). Using the treated students’ choices, we plan to test this hypothesis for each of the models estimated on the control data. Note that this test assumes that the distributions of unobservables are the same between control and treatment villages, consistent with the village-level randomization.

This hypothesis testing framework further allows for understanding the sources of model mis-specification through the inclusion of additional terms in equation (8). For example, it may be that the models fail because of the instruments for tuition. To test this, we can insert p_j into the equation – if treated households, who in reality face no tuition by virtue of receiving the voucher, nonetheless appear to be sorting on private schools’ tuition, that indicates that the model fails to

³⁸Likewise, school value-added for treatment village schools is formed using the parameters estimated on the control sample.

adequately control for unobserved private school characteristics. Alternatively, it may be that the voucher recipients search out or seek better information about schooling alternatives before making their choice. If so, this would show up as a change in their “preferences” over private schools’ non-tuition characteristics, such as value-added. These variables can similarly be included in equation (8) to test for this source of mis-specification. We can also examine whether the models fail for specific subgroups via interactions with household characteristics. Table 17 details specific sources of mis-specification we will examine using this framework. By including applicants who did not receive the voucher in the estimation sample, we can test whether winning the voucher induced apparent “excitement” on the part of recipients (i.e. that the voucher offer itself has a treatment effect on recipients’ preferences for private schooling).³⁹

Table 17: Experimental Validation: Hypothesis Tests

Source of Mis-specification	Add to Indirect Utility
Private school unobservables	p_j
Voucher-induced info acquisition	$Private_j$ interactions with non-tuition school characteristics
Subgroup heterogeneity	$Private_j$ interactions with household characteristics
Omitted match quality	Interaction of school quality and baseline test score
“Excitement” for voucher	$Private_j$ for voucher recipients (include non-recipient applicants in estimation)
Selection into application	$Private_j$ for voucher applicants (include eligible non-applicants in estimation)

7 Conclusion

Alongside treatment effects on outcomes targeted by policymakers, effects on economic well-being are especially of interest for interventions that expand the choices available to beneficiaries. In this paper, we pursue a unique research design to estimate the welfare impacts of private school vouchers in rural India. Vouchers may create significant welfare gains by facilitating better matches between students and schools. Vouchers may also reduce the social costs of universal education when government provision is inefficient.

Using control group data drawn from a randomized controlled trial of private school vouchers in rural India, we first develop and estimate empirical models of how households choose among

³⁹Given that we assume that school value-added is homogeneous across students, systematic differences in returns *ex post* can be a source of mis-specification. However, because households choose schools based on (potentially mistaken) beliefs about match quality, the quantitative importance of this mis-specification will depend in practice on how correlated the school characteristics-student characteristics interactions included in the models are with expected match quality. In this regard, our models do not allow households to value schools differently according to their baseline measured abilities. If parents’ beliefs depend on their child’s ability, hypothesis tests using this information – e.g. interacting baseline test score with school value-added – can help assess this source of bias.

government and private primary schools in their village. We model households' choice of school in this setting as potentially ability-to-pay constrained and place structure on how observed household assets influence choice patterns to empirically separate ability-to-pay from willingness-to-pay. We compare this model with estimates and predictions from flexible logit models of demand, including a random coefficient model.

In this paper, we present results from the first step of the research design, including predictions for the experimental validation. The model estimates provide evidence that households, particularly those targeted by the AP project voucher, lack ability-to-pay for private schooling. The constrained model estimates that about a quarter of households that choose a government school are unable to choose any private school in their village. This constraint in turn has significant implications for the estimates of willingness-to-pay, including for English medium instruction and school value-added. For instance, while non-constrained models estimate lower relative preferences for school quality for asset-poor households, the ability-to-pay constrained model estimates that willingness-to-pay for value-added is similar across households of varying socioeconomic status. Moreover, the ability-to-pay constrained model estimates a substantially larger consumer surplus from the voucher, particularly among those marginal students that are induced into private schooling because of the offer.

We generate and report predictions from the estimated models to be compared with the experimental outcomes in treatment villages in the research design's second step yet to come. In particular, we simulate a voucher trial mimicking the experimental treatment offer and compare take-up and voucher elasticities. We find that the baseline ability-to-pay constrained model predicts that around 63% of households will use the voucher to attend a private school. The models that do not accommodate ability-to-pay differences instead estimate take-up to be around 10 percentage points lower. The sharpest differences between models in predictions appear for lower socioeconomic households, though. The ability-to-pay constrained model predicts a voucher elasticity among asset-poor households that is up to 50% larger than the predictions generated by unconstrained demand models.

References

- Abdulkadiroglu, A., Pathak, P. A., Schellenberg, J., and Walters, C. R. (2017a). Do Parents Value School Effectiveness? Working Paper 23912, National Bureau of Economic Research.
- Abdulkadiroglu, A., Pathak, P. A., and Walters, C. R. (2017b). Free to Choose: Can School Choice Reduce Student Achievement? *American Economic Journal: Applied Economics*.
- Abdulkadiroglu, A., Agarwal, N., and Pathak, P. A. (2015). The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match. Working Paper 21046, National Bureau of Economic Research.
- Abdulkadiroglu, A. and Sönmez, T. (2003). School Choice: A Mechanism Design Approach. *The American Economic Review*, 93(3):729–747.
- Allende S.C., C. (2019). Competition Under Social Interactions and the Design of Education Policies.
- Allende S.C., C., Gallego, F., and Neilson, C. (2019). Approximating the Equilibrium Effects of Informed School Choice.
- Andrabi, T., Das, J., and Khwaja, A. I. (2017a). 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.
- Andrabi, T., Das, J., Khwaja, A. I., Ozyurt, S., and Singh, N. (2017b). Upping the Ante: The Equilibrium Effects of Unconditional Grants to Private Schools.
- Angrist, J., Bettinger, E., and Kremer, M. (2006). Long-Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia. *The American Economic Review*, 96(3):847–862.
- ASER (2018). *Annual Status of Education Report*. Aser Centre.
- Attanasio, O. P., Meghir, C., and Santiago, A. (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, A., Hanna, R., Olken, B. A., Satriawan, E., and Sumarto, S. (2021). Food vs. Food Stamps: Evidence from an At-Scale Experiment in Indonesia. Working Paper 28641, National Bureau of Economic Research. Series: Working Paper Series.
- Banerjee, A. V. and Duflo, E. (2014). Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program. *The Review of Economic Studies*, 81(2):572–607.
- Barseghyan, L., Coughlin, M., Molinari, F., and Teitelbaum, J. C. (2021). Heterogenous Choice Sets and Preferences.
- Bau, N. (2017). School Competition and Product Differentiation.
- Bayer, P., Ferreira, F., and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, 115(4):588–638.

- Ben-Akiva, M. and Boccara, B. (1995). Discrete choice models with latent choice sets. *International Journal of Research in Marketing*, 12(1):9–24.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.
- Cameron, S. V. and Heckman, J. J. (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, P., Das, J., and Reis, H. (2016). The Value of Private Schools: Evidence from Pakistan.
- Cellini, S. R., Ferreira, F., and Rothstein, J. (2010). The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design*. *The Quarterly Journal of Economics*, 125(1):215–261.
- Currie, J. and Gahvari, F. (2008). Transfers in Cash and In-Kind: Theory Meets the Data. *Journal of Economic Literature*, 46(2):333–383.
- Deming, D. J. (2014). Using School Choice Lotteries to Test Measures of School Effectiveness. *American Economic Review*, 104(5):406–411.
- Dizon-Ross, R. (2019). Parents’ Beliefs about Their Children’s Academic Ability: Implications for Educational Investments. *American Economic Review*, 109(8):2728–2765.
- Dongre, A. (2012). What is the Per Child Expenditure in Government Schools? In *Accountability Initiative*. New Delhi.
- Duflo, E., Hanna, R., and Ryan, S. P. (2012). Incentives Work: Getting Teachers to Come to School. *The American Economic Review*, 102(4):1241–1278.
- Epple, D., Romano, R. E., and Urquiola, M. (2017). School Vouchers: A Survey of the Economics Literature. *Journal of Economic Literature*, 55(2):441–492.
- Filmer, D. and Pritchett, L. H. (2001). Estimating Wealth Effects Without Expenditure Data—Or Tears: An Application To Educational Enrollments In States Of India. *Demography*, 38(1):115–132.
- Galiani, S., Murphy, A., and Pantano, J. (2015). Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment. *The American Economic Review*, 105(11):3385–3415.
- Galiani, S. and Pantano, J. (2021). Structural Models: Inception and Frontier. Working Paper 28698, National Bureau of Economic Research. Series: Working Paper Series.
- Gandhi, A. and Houde, J.-F. (2016). Measuring Substitution Patterns in Differentiated Products Industries.
- Gregory, J. (2017). The Impact of Post-Katrina Rebuilding Grants on the Resettlement Choices of New Orleans Homeowners.
- Hastings, J. and Shapiro, J. M. (2018). How Are SNAP Benefits Spent? Evidence from a Retail Panel. *American Economic Review*, 108(12):3493–3540.

- Hastings, J. S., Kane, T. J., and Staiger, D. O. (2005). Parental Preferences and School Competition: Evidence from a Public School Choice Program. Working Paper 11805, National Bureau of Economic Research.
- Hausman, J. A. (1996). Valuation of New Goods under Perfect and Imperfect Competition. In Bresnahan, T. F. and Gordon, R. J., editors, *The Economics of New Goods*, pages 207–248. University of Chicago Press.
- Kane, T. J. and Staiger, D. O. (2008). Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation. Working Paper 14607, National Bureau of Economic Research.
- Keane, M. P. and Wolpin, K. I. (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].
- Keane, M. P. and Wolpin, K. I. (2007). Exploring the Usefulness of a Nonrandom Holdout Sample for Model Validation: Welfare Effects on Female Behavior. *International Economic Review*, 48(4):1351–1378.
- Koedel, C., Mihaly, K., and Rockoff, J. E. (2015). Value-added modeling: A review. *Economics of Education Review*, 47:180–195.
- Lagakos, D., Mobarak, M., and Waugh, M. E. (2017). The Welfare Effects of Encouraging Rural-Urban Migration.
- LaLonde, R. J. (1986). Evaluating the Econometric Evaluations of Training Programs with Experimental Data. *The American Economic Review*, 76(4):604–620. Publisher: American Economic Association.
- MacLeod, W. B. and Urquiola, M. (2013). Competition and Education Productivity: Incentives Writ Large. In *Education Policy in Developing Countries*, pages 243–284.
- MacLeod, W. B. and Urquiola, M. (2015). Reputation and School Competition. *American Economic Review*, 105(11):3471–3488.
- Manski, C. F. and Lerman, S. R. (1977). The Estimation of Choice Probabilities from Choice Based Samples. *Econometrica*, 45(8):1977–1988. Publisher: [Wiley, Econometric Society].
- Mills, J. N. and Wolf, P. J. (2017). Vouchers in the Bayou: The Effects of the Louisiana Scholarship Program on Student Achievement After 2 Years. *Educational Evaluation and Policy Analysis*, 39(3):464–484.
- Misra, S. and Nair, H. S. (2011). A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quantitative Marketing and Economics*, 9(3):211–257.
- Muralidharan, K. and Sundararaman, V. (2015). The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India. *The Quarterly Journal of Economics*, 130(3):1011–1066.
- Neilson, C. (2013). Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students.
- Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*, 69(2):307–342.

- Pathak, P. A. and Shi, P. (2017). How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice. Working Paper 24017, National Bureau of Economic Research.
- Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, 110(4):705–729.
- Petrin, A. and Train, K. (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research*, 47(1):3–13.
- Piketty, T., Saez, E., and Zucman, G. (2018). Distributional National Accounts: Methods and Estimates for the United States. *The Quarterly Journal of Economics*, 133(2):553–609.
- Rosenzweig, M. R. and Wolpin, K. I. (1993). Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India. *Journal of Political Economy*, 101(2):223–244.
- Rothstein, J. 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.
- Rouse, C. E. (1998). Private School Vouchers and Student Achievement: An Evaluation of the Milwaukee Parental Choice Program. *The Quarterly Journal of Economics*, 113(2):553–602.
- Salz, T. and Vespa, E. (2017). Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation.
- Schorfheide, F. and Wolpin, K. I. (2012). On the Use of Holdout Samples for Model Selection. *American Economic Review*, 102(3):477–481.
- Schorfheide, F. and Wolpin, K. I. (2016). To hold out or not to hold out. *Research in Economics*, 70(2):332–345.
- Singleton, J. D. (2017). Incentives and the Supply of Effective Charter Schools.
- Sánchez, C. (2018). Understanding School Competition Under Voucher Regimes.
- Tarozzi, A., Mahajan, A., Blackburn, B., Kopf, D., Krishnan, L., and Yoong, J. (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, P. E. and Wolpin, K. I. (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.
- Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica*, 62(3):539–591.
- Wise, D. A. (1985). A behavioral model versus experimentation: The effects of housing subsidies on rent. *Methods of Operations Research*, 50:441–89.

Appendix

A Sample and Weights

We weight students in the estimation sample to account for three features of the AP School Choice project sampling design. First, students who did not apply for the AP project voucher (kindergartners and first graders, the latter asked regarding a voucher for second grade) are underrepresented in the sample. We adjust for this using the student counts reported in Table II of Muralidharan and Sundararaman (2015) to estimate sampling probabilities.

Second, first graders (whose primary school choices were collected at baseline) were sampled conditional on their choice of primary school. For consistency, we therefore re-weight to match the population market shares of students attending public and private schools (Manski and Lerman, 1977). To do so, we consult India’s Annual Status of Education Report (ASER) survey in 2008 (the same calendar year at the project baseline). We calculate population shares of private school attendance at the district level, restricting the ASER sample to households in Andhra Pradesh villages with at least one private school and students in the age range of our estimation sample (and excluding children not enrolled in school).

Finally, there is attrition of kindergartners (whose primary school choices are made in subsequent years) from the baseline sample. As summarized in Table A18, we observe primary school choices in the first wave (Y1) followup of 1,002 of the 1,783 kindergarten students surveyed at baseline. Most of the remaining students remained in daycare during the first wave. We therefore use the next wave of household data collection (Y3) to collect primary school choices for 483 of these students. This yields an estimation sample of 1,485 kindergarten students as well as 298 from the baseline sample who attrit (i.e. we do not observe choice information by Y3). We adjust for this for by estimating the probability that each household attrits. The estimated attrition probability depends on student and household sociodemographic characteristics, baseline Telugu score, district of residence, as well as whether they are eligible and/or applied for the AP project voucher.

Table A18: Kindergarten Subsample

	N
Baseline sample	1,783
Y1 choice observed	1,002
Y3 choice observed	483
Estimation sample	1,485

B School Value-Added

We estimate and control for each schools’ value-added to student learning in the choice estimation using baseline and follow-up exam scores in math. We assume that the achievement of student i in subject k at year t is a linear function of household i ’s characteristics, H_i , the quality of the school they attend $\omega_{j(i,t)}$, and their prior exam performance:

$$A_{ijt} = \rho(A_{it-1}) + H_i' \pi + \omega_{j(i,t)} + \zeta_{ijt} \quad (9)$$

A_{ijt} is the student’s exam performance in year t , which is normalized across students within year. School j ’s unobserved value-added to the learning process, ω_j , is assumed to be fixed within our panel. We include a cubic of prior exam performance, A_{ijt-1} , and control for student demographics, parental education, and household socioeconomic status in H_i .

We observe up to three math scores for each student with the first two coming in *non-consecutive* school years (the baseline year and year three) in the data. We therefore estimate school value-added by first estimating equation (9) separately by year. For each school, this step yields up to two fixed effects. We then shrink the fixed effect estimates using empirical Bayes techniques (Kane and Staiger, 2008; Deming, 2014; Koedel et al., 2015). We finally recover estimates of the value-added for each school from the shrunken fixed effects by imposing the assumption that value-added net of a depreciation parameter is the same in both years (i.e. we solve for δ such that $\hat{\omega}_{j4} = \delta \hat{\omega}_{j3}$).

C Control Function and Instruments

We use a control function approach in estimating the models to instrument for endogenous private school tuition and fees (Petrin and Train, 2010). We take a control function approach because the sampling design does not allow us to compute population market shares for each school. The “first stage” is given by:

$$p_j = X_j' \Gamma + f(Z_j) + \mu_j$$

where X_j are observed characteristics (including value-added) and Z_j consists of instruments for tuition and fees at private school j . ξ_j and μ_j are assumed jointly normal. $\hat{\mu}_j$ and η_j (a normally distributed random effect) then replace ξ_j in the indirect utility function.

Our baseline models use as instruments a summary measure of each private school’s location in product space as well as proxy for its costs. We also estimate models with a cost index (and its square) replacing the cost proxy for robustness. We construct the product space instrument by

first calculating:

$$Z_j^k = \sum_{l \in \mathcal{V}_j} X_l^k - X_j^k$$

for each non-tuition school characteristic k . We then take the first factor of the Z_j^k variables as the instrument. To construct the cost proxy instrument, we separate private schools by their language of instruction and then use k-means clustering to group private schools within medium of instruction according to similarity in their characteristics.⁴⁰ The cost proxy instrument is then the average tuition and fees of those private schools that belong to the same cluster, but which are located in other districts of Andhra Pradesh. The village cost index is estimated using village fixed effects in a regression of private schools' log per pupil expenditures on fixed school characteristics.⁴¹ Table A19 below presents first-stage estimates and diagnostics for both sets of instruments.

	(1) Baseline IVs	(2) Alternative
Product space location	238.1*** (56.42)	288.9*** (59.60)
Cost proxy	0.376*** (0.135)	
Cost index		0.246*** (0.0947)
Cost index ²		-0.000599*** (0.000132)
First-stage F	12.51	17.63
Cragg-Donald stat	11.20	13.13
R-squared	0.309	0.341

Notes: Table presents first stage estimates that regress private school tuition and fees on school characteristics and the instrumental variables. 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). Though not reported, regressions control for the school characteristics included in the choice models. $N = 293$ private schools. Standard errors reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

⁴⁰We form groups based on similarity in terms of facilities, teaching quality, and a factor summarizing horizontal differentiation (e.g. whether offers Hindi and/or teaches computer skills). We group English medium private schools into 10 clusters and Telugu schools into 8 clusters.

⁴¹We then use the average fixed effects of other villages in the same district as each village's cost index. We take this additional step to ensure that the index is not confounded by unobserved differences in ξ_j between villages.

D Likelihood Functions

We detail the likelihood functions for each model in this section. The log-likelihoods we estimate can be expressed by:

$$L^m = \sum_i w_i \ln[e_i L_{1i}^m(\theta) + (1 - e_i) L_{0i}^m(\theta)] \quad (10)$$

where m indexes the models and specifications (e.g. $m \in \{\text{CMNL, RC, CC}\}$). w_i are household-specific weights, described in Appendix A. e_i is a variable taking the value of 1 if household i is eligible for the AP voucher and $L_{1i}^m(\theta)$ is the sub-likelihood of choices given the household is eligible. We detail these model-specific likelihoods below. θ (for a given m) represents the parameter vector we aim to estimate. For kindergarten households in our dataset, we observe AP voucher eligibility, e_i . However, e_i is unknown for first graders. We thus specify for these households that

$$\ln \frac{\text{Pr}(e_i = 1)}{1 - \text{Pr}(e_i = 1)} = b' C_i$$

where C_i includes observed household characteristics (demographics, parental education, discrete household asset levels).

We estimate equation (10) using the EM algorithm. At the θ -maximization step, we maximize:

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

where \tilde{e}_i is the conditional or posterior probability that i is eligible for the AP voucher (if i is a first grader) given model m . This is given by:

$$\tilde{e}_i^m = \frac{e_i L_{1i}^m(\theta)}{e_i L_{1i}^m(\theta) + (1 - e_i) L_{0i}^m(\theta)}$$

The algorithm iterates on the expectation and maximization steps until the parameter vector θ converges.

D.1 Clustered and Random Coefficient Model Likelihoods

The likelihoods for these models are pretty standard expressions, so we do not write them out here. Both numerically integrate over the private school-specific random effects. We use 100 Halton draws for this integration. For the random coefficient model, we use Gauss-Hermite quadrature to integrate the random coefficient.

D.2 Ability-to-Pay Constrained Likelihood

The constrained choice model requires integration over households' unobserved ability-to-pay. This integration is simplified by the fact that there are a finite number of possible choice sets (or states) for each household. 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, state j_i^* corresponds to the choice set in which the household can choose j_i^* , but not $j_i^* + 1$: $p_{j_i^*} \leq Y_i \leq p_{j_i^*+1}$.

Given the assumption on the choice shocks, the probability that i chooses private school j in their village in choice state j_i^* can be written as:

$$P_{ij}(j_i^*) = \frac{\mathbf{1}\{p_j \leq p_{j_i^*}\} \exp U_{ij}}{\sum_{j \in \mathcal{V}_i} \mathbf{1}\{p_j \leq p_{j_i^*}\} \exp U_{ij}} \quad (11)$$

Note that this choice probability is 0 if j is not in the choice set and the summation in the denominator is only over those village alternatives for which i is able to pay.

If ability-to-pay were observed, these choice probabilities would be simply plugged into the log likelihood for estimation. With unobserved ability-to-pay, however, we integrate out the finite possible choice sets for each i . Let ϕ_{ij^*} denote the probability that household i is in choice state j_i^* . Given the distributional assumption for ε_i , we can write this as:

$$\begin{aligned} \phi_{ij^*} &= P(p_{j_i^*} \leq Y_i < p_{j_i^*+1}) \\ &= P(\ln p_{j_i^*} \leq \ln Y_i < \ln p_{j_i^*+1}) \\ &= \Phi\left(\frac{\ln p_{j_i^*+1} - I_i' \lambda}{\sigma}\right) - \Phi\left(\frac{\ln p_{j_i^*} - I_i' \lambda}{\sigma}\right) \end{aligned}$$

For the most expensive choice alternative, the expression is modified accordingly as:

$$\begin{aligned} \phi_{iJ_i^*} &= P(\ln p_{J_i^*} \leq \ln Y_i) \\ &= 1 - \Phi\left(\frac{\ln p_{J_i^*} - I_i' \lambda}{\sigma}\right) \end{aligned}$$

Combining the expressions for the choice probabilities with these state probabilities, the likelihood function takes the form:

$$L_i^{CC}(\theta) = \sum_{j_i^*} \phi_{ij^*} \prod_{j \in \mathcal{V}_i} P_{ij}(j_i^*)^{d_{ij}} \quad (12)$$

Though not represented in the above expression, the ability-to-pay constrained choice model also includes private school-specific error components. As with the clustered and random coefficient models, we numerically integrate these normal random effects using 100 Halton draws.

E Willingness-to-Pay and Compensating Variation Calculations

Calculating willingness-to-pay and the compensating variation of a voucher in money terms requires scaling the changes by marginal (flow) utility of consumption. To obtain an estimate of marginal utility of consumption for each model, we calculate:

$$\alpha_i^m = \frac{\hat{\alpha}_i^m}{1 + \delta + \delta^2 + \delta^3 + \delta^4}$$

where m indexes the models (e.g. $m \in \{\text{CMNL,RC,CC}\}$), $\hat{\alpha}_i^m$ corresponds to estimated coefficients on tuition and fees (presented in Table 3), and δ is the effective annual discount factor. For δ , we use the product of 0.95 (a 5% annual discount rate) and 0.94 (the annual probability that a voucher recipient remains in private school, calculated from Muralidharan and Sundararaman (2015)). Note that we also use δ for calculating expected costs and fiscal externalities of vouchers.

This calculation of marginal utility of consumption can be viewed as following from the assumptions that primary schooling is five periods (during which tuition remains constant) with future periods discounted by δ and that the post-primary schooling value of the primary school choice does not depend on primary school tuition and fees. This latter assumption may be violated if, for example, a voucher during primary school allows some households to finance private secondary schooling. In such a case, however, note that our estimates of compensating variation will be lower bounds.

F Additional Tables

Table A20: Characteristics of Private Schools by Quartile of Tuition and fees

	1st	2nd	3rd	4th
Tuition and fees (Rs.)	720	1,327	1,957	3,295
English medium	0.34	0.43	0.76	0.92
Full pucca building	0.45	0.40	0.44	0.64
Library	0.76	0.74	0.74	0.85
Functioning toilet	0.75	0.78	0.82	0.92
Separate toilet for girls	0.53	0.55	0.57	0.71
Staffroom for teachers	0.55	0.68	0.76	0.84
Has secondary school	0.22	0.32	0.24	0.22
Multi-class teaching	0.36	0.23	0.30	0.18
Share teachers absent	0.09	0.12	0.08	0.10
Share teachers with BA	0.47	0.46	0.63	0.73
Share teachers female	0.73	0.69	0.71	0.72
Share teachers from village	0.55	0.52	0.38	0.42
Offers Hindi	0.36	0.33	0.41	0.59
Teaches computer skills	0.01	0.09	0.13	0.29
School value-added	0.03	0.12	0.03	0.08
N	74	76	70	73

Table A21: Estimates: Willingness-to-Pay (1000s of Rs.) for 1 Mile Reduction in Travel Distance to School

	CMNL	RC	CC
Overall	6.77	6.11	7.54
Female	6.86	6.31	7.90
Muslim	3.94	4.47	5.41
Lower caste	6.90	5.43	6.56
Older sibling in gov't school	6.58	6.04	7.03
Both parents completed primary school	6.63	5.74	7.65
≥ 1 parent completed secondary	6.20	5.35	7.17
Both parents laborers	6.87	5.88	7.10
Asset level < 3	6.16	4.87	7.38
Asset level = 3	7.72	6.42	7.95
Asset level = 4	8.17	7.58	7.52
Asset level > 4	5.10	5.53	7.33

Notes: Table presents average willingness-to-pay (1000s of Rs.) by subgroup for 1 mile reduction in travel distance to school. Estimates correspond to baseline specifications described in the text.

Table A22: Preference Heterogeneity in Baseline Model Specifications

CMNL	RC	CC	School characteristic	Interactions
X	X		Tuition and fees	Eligibility, total siblings, discrete asset levels (e.g. asset level = 3, etc.)
X	X	X	Distance	Eligibility, female, Muslim, lower caste
X			Closest Public	
X	X	X	Private	Eligibility, female, Muslim, lower caste, parental education, older sibling gov't, total siblings
X	X	X	English medium	Female, Muslim, lower caste, parental education
X	X	X	Value-added	Female, Muslim, lower caste, parental education
X	X	X	Offers Hindi	Female, Muslim, lower caste, parental education
X			Facilities	
X			Teaching Quality	
X			Teacher Char.	
X			Secondary	

Notes: Table details preference heterogeneity in baseline model specifications. For CMNL model, heterogeneity for each non-tuition attribute is captured by 51 group-specific coefficients (not interactions); tuition is interacted with listed household characteristics. Like other models, CMNL model also includes private school and distance interactions with AP voucher eligibility. Parental education variables are whether both parents completed primary and whether at least one completed secondary school. Facilities (e.g. pucca), teaching quality (e.g. multiclass teaching, share teachers absent), and teacher characteristics (e.g. share female, from village) attributes are first principal factors. Secondary is an indicator for whether school offers secondary grades. Variables included in the models but not listed in the table are indicators for imputed tuition and fees and for missing value-added, effects of which are not heterogeneous across households.

Table A23: Estimates and Predictions for Alternative Model Specifications

Specification	Demand ϵ	Constraint	Surplus	Fit	Take-up	Voucher ϵ
<i>A. Ability-to-pay Constrained (CC)</i>						
(1) Baseline	0.47	0.24	7.99	-0.02	0.63	93
(2) Instead uses cost index IV	0.45	0.22	7.89	-0.01	0.61	86
(3) No control function	-0.09	0.25	-76.68	-0.02	0.35	10
(4) Includes whether lower caste in preferences and ability-to-pay	0.47	0.23	7.96	-0.02	0.62	93
(5) Includes whether parents are laborers in preferences and ability-to-pay	0.46	0.25	8.13	-0.01	0.63	94
(6) Allows preferences for private schooling to be cohort-specific	0.45	0.14	7.29	0.00	0.64	86
(7) Does not include eligibility for AP voucher in the model	0.67	0.37	8.44	0.10	0.78	76
(8) Does not re-weight for sample attrition in estimation	0.48	0.27	7.98	-0.02	0.63	99
<i>B. Clustered Multinomial Logit (CMNL)</i>						
(1) Baseline	0.65	0	4.78	-0.04	0.54	82
(2) Instead uses cost index IV	0.70	0	4.87	-0.04	0.56	88
<i>C. Random Coefficient (RC)</i>						
(1) Baseline	0.47	0	4.66	-0.01	0.53	63
(2) Instead uses cost index IV	0.48	0	4.70	-0.01	0.53	64

Notes: Table reports select estimates and predictions for baseline and alternative model specifications. See Tables A24 and A25 for corresponding point estimates. Constraint column refers to probability that government-attending students cannot choose any private school; surplus column reports welfare gain to compliers. Surplus reported in 1,000s of Rs. Fit column is difference between model predicted private school attendance of eligible kindergartners and share observed in the data.

Table A24: Ability-to-pay Constrained Model Estimates: Baseline and Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tuition and fees	-1.28	-1.19	0.25	-1.29	-1.28	-1.25	-1.52	-1.29
First stage residual (baseline IVs)	1.77			1.77	1.78	1.67	1.94	1.73
First stage residual (alternative)		1.59						
Random effect σ	2.66	2.64	2.48	2.66	2.63	2.33	2.30	2.53
Log distance	-1.41	-1.38	-1.33	-1.42	-1.41	-1.00	-1.10	-1.38
x Eligible for AP voucher	0.29	0.28	0.23	0.30	0.28	0.04		0.22
x Age > 5	0.15	0.14	0.15	0.15	0.15	0.11	0.31	0.17
x Female	-0.14	-0.14	-0.15	-0.14	-0.14	-0.15	-0.13	-0.12
x Muslim	0.13	0.13	0.16	0.13	0.14	0.13	-0.01	0.17
x Lower caste	-0.05	-0.05	-0.06	-0.05	-0.06	-0.08	-0.10	-0.02
x Parents laborers					0.03			
x First grader						-0.31		
Closest public school	0.78	0.80	0.74	0.78	0.79	0.76	0.77	0.78
Private	11.35	10.80	9.76	11.40	11.52	10.06	3.50	11.14
x Eligible for AP voucher	-10.13	-9.66	-10.43	-10.18	-10.15	-8.70		-9.77
x Female	-0.60	-0.70	-0.58	-0.59	-0.60	-0.58	-0.59	-0.74
x Muslim	0.16	0.01	0.04	0.16	0.11	0.23	-0.33	0.00
x Lower caste	-1.50	-1.61	-1.54	-1.57	-1.37	-1.69	-2.19	-1.39
x Both parents completed primary	0.17	0.21	0.07	0.17	0.08	0.03	0.50	0.17
x ≥ 1 parent completed secondary	0.58	0.37	0.76	0.57	0.57	0.73	1.91	0.53
x Older sibling in gov't	-2.59	-2.52	-2.45	-2.59	-2.59	-2.12	-3.06	-2.53
x Total siblings - 2	-0.07	-0.09	-0.03	-0.07	-0.06	-0.34	0.38	0.08
x Parents laborers					-0.26			
x First grader						-1.59		
English medium	0.90	0.83	-0.12	0.90	0.90	0.89	0.87	0.79
x Female	-0.92	-0.93	-0.93	-0.92	-0.92	-0.91	-0.74	-0.85
x Muslim	1.42	1.46	1.41	1.42	1.42	1.33	1.76	1.56
x Lower caste	0.05	0.07	0.26	0.04	0.06	0.11	0.04	-0.02
x Both parents completed primary	0.85	0.83	0.90	0.85	0.85	0.81	0.79	1.07
x ≥ 1 parent completed secondary	1.35	1.33	1.06	1.36	1.34	1.19	1.18	1.24
Value-added	0.51	0.52	0.47	0.51	0.51	0.49	0.34	0.54
x Female	0.03	0.03	0.04	0.03	0.03	0.07	-0.02	-0.05
x Muslim	-0.16	-0.16	-0.20	-0.16	-0.17	-0.20	-0.19	-0.24
x Lower caste	-0.04	-0.04	-0.01	-0.05	-0.05	-0.07	0.07	-0.03
x Both parents completed primary	0.31	0.27	0.16	0.31	0.31	0.26	0.42	0.37
x ≥ 1 parent completed secondary	-0.51	-0.56	-0.57	-0.51	-0.51	-0.49	-0.45	-0.54
Offers Hindi	0.03	-0.03	-0.21	0.03	0.02	0.07	-0.03	-0.03
x Female	0.55	0.55	0.53	0.55	0.54	0.44	0.47	0.51
x Muslim	1.15	1.04	1.03	1.15	1.16	1.17	0.70	1.09
x Lower caste	1.15	1.13	1.09	1.15	1.16	1.08	1.07	1.28
x Both parents completed primary	0.39	0.36	0.25	0.39	0.40	0.31	0.47	0.31
x ≥ 1 parent completed secondary	0.18	0.11	0.11	0.18	0.18	0.18	0.14	0.36
Unrecognized private	-0.61	-0.67	-0.64	-0.62	-0.62	-0.51	-0.45	-0.49
Facilities factor	0.46	0.45	0.32	0.46	0.46	0.44	0.45	0.44
Teaching quality factor	-0.34	-0.33	-0.43	-0.34	-0.34	-0.30	-0.27	-0.32
Teacher characteristics factor	-0.06	-0.08	-0.02	-0.06	-0.06	-0.07	-0.01	-0.07
<i>Ability-to-pay function</i>								
Constant	2.96	3.27	3.08	2.93	2.96	2.60	2.44	3.29
Asset factor	1.09	1.21	1.08	1.09	1.01	0.93	2.23	1.17
Eligible for AP voucher	-1.29	-1.34	-1.42	-1.28	-1.23	-0.75		-1.61
Total siblings - 2	0.37	0.42	0.31	0.37	0.34	0.53	-0.12	0.21
Lower caste				0.11				
Parents laborers					-0.34			
σ	1.34	1.50	1.42	1.33	1.34	1.16	2.10	1.50
Share first graders eligible	0.60	0.60	0.59	0.60	0.60	0.49		0.59
Attrition weights	Y	Y ₅₀	Y	Y	Y	Y	Y	N

Notes: Table reports point estimates for baseline and alternative specifications of ability-to-pay constrained model. 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 A25: Logit Demand Model Estimates

	RC		CMNL			
	(1)	(2)	(1)		(2)	
			Coef	SD	Coef	SD
Tuition and fees	-2.35	-2.43	-2.00		-2.23	
x Eligible for AP voucher	0.07	0.16	0.37		0.50	
x Asset level = 2	0.45	0.46	0.22		0.24	
x Asset level = 3	0.74	0.73	0.45		0.46	
x Asset level = 4	1.12	1.11	0.68		0.69	
x Asset level > 4	0.81	0.80	0.39		0.38	
x Total siblings - 2	0.05	0.04	0.05		0.03	
First stage residual (baseline IVs)	1.60		1.48			
First stage residual (alternative)		1.62			1.64	
Random effect σ	2.23	2.26	1.73		1.79	
Log distance	-1.15	-1.12	-0.94	0.55	-0.95	0.55
x Eligible for AP voucher	-0.06	-0.07	-0.30		-0.28	
x Age > 5	0.15	0.14				
x Female	-0.11	-0.11				
x Muslim	0.10	0.10				
x Lower caste	-0.05	-0.05				
Closest public school	0.75	0.76	0.78	1.01	0.81	1.01
Private	11.75	11.60	6.84	5.67	6.64	5.65
x Eligible for AP voucher	-10.86	-10.63	-8.78		-8.51	
x Female	-0.72	-0.81				
x Muslim	0.96	0.79				
x Lower caste	-1.81	-1.94				
x Both parents completed primary	0.94	0.98				
x ≥ 1 parent completed secondary	0.97	0.68				
x Older sibling in gov't	-3.52	-3.42				
x Total siblings - 2	0.14	0.10				
Random coeff. σ	3.32	3.23				
English medium	0.91	0.95	-0.35	6.36	-0.23	6.37
x Female	-0.89	-0.89				
x Muslim	1.31	1.38				
x Lower caste	-0.07	-0.09				
x Both parents completed primary	0.88	0.86				
x ≥ 1 parent completed secondary	1.29	1.32				
Value-added	0.50	0.52	0.36	0.74	0.38	0.75
x Female	0.05	0.04				
x Muslim	-0.16	-0.12				
x Lower caste	0.00	-0.01				
x Both parents completed primary	0.32	0.30				
x ≥ 1 parent completed secondary	-0.57	-0.62				
Offers Hindi	-0.05	-0.06	1.13	8.92	1.11	8.89
x Female	0.53	0.51				
x Muslim	1.09	1.00				
x Lower caste	1.19	1.20				
x Both parents completed primary	0.21	0.18				
x ≥ 1 parent completed secondary	0.21	0.14				
Unrecognized private	-0.52	-0.61	-0.34		-0.42	
Facilities factor	0.45	0.45	0.52	0.44	0.53	0.44
Teaching quality factor	-0.32	-0.31	-0.20	0.70	-0.17	0.71
Teacher characteristics factor	-0.07	-0.10	-0.01	0.38	-0.05	0.38
Share first graders eligible	0.60	0.60	0.57		0.57	

Notes: Table reports point estimates for random coefficient and clustered multinomial logit models using (1) baseline instruments and (2) alternative cost index instrument. 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. Coef columns report across-cluster averages of estimates for CMNL model; SD columns report across-cluster standard deviations. CMNL models include 51 clusters.

Table A26: Factor variables

Asset factor Variables	Coef	Facilities factor Variables	Coef	Teaching quality factor Variables	Coef	Teacher characteristics factor Variables	Coef	Product space location IV Variables	Coef
Pucca house	0.65	Full pucca building	-0.04	Multi-class teaching	0.31	Average age of teachers	0.00	<i>Other, same-village...</i>	0.25
Water facility in home	0.41	Library	0.21	Share teachers with BA	0.32	Share teachers female	-0.22	Private schools	0.16
Household toilet	0.55	Full compound walls	0.25	Share teachers with ed. training	0.74	Share teachers lower caste	0.08	English medium schools	-0.02
Owens home	0.11	Staffroom for teachers	0.67	Share teachers absent	1.15	Share teachers from village	-0.24	Unrecognized schools	0.01
≥2 covered rooms in home	0.63	Functional water tap	0.28	Share teachers actively teaching	-0.90	Share teachers Hindu	2.46	Schools' value-added	-0.01
Owens land	0.28	Separate toilet for girls	0.35	Share teachers maintaining disc.	-0.52	Share teachers Muslim	-0.33	Schools' facilities	0.02
Constant	-1.40	Playground	0.33	Constant	-0.38	Share teachers Christian	-0.28	Schools' teaching quality	-0.01
		Firstaid	0.42			Constant	-1.86	Schools' teacher char.	-0.01
		Kitchen facility	0.00					Schools	-0.03
		Has electricity	0.40					Constant	-1.07
		Constant	-1.52						
N	5,255		645		645		645		293