Valuing School Choice: Using a Randomized Experiment to Validate Welfare Evaluation of Private School Vouchers*

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

October 2020

Abstract

In this paper, we use a unique two-stage experiment in rural India that randomized access to school vouchers across both markets and students to estimate the revealed preference welfare valuation of school choice. In the first step of the research design, we develop an empirical model of primary 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 school voucher generated welfare value exceeding 9,000 Rs. (about five times the average private school's annual tuition) on average to the students induced into private schooling by the program. This magnitude is twice as large as estimates obtained from demand models that instead ignore ability-to-pay constraints on households' choice. The second step of the research design will experimentally validate the estimated welfare impacts by comparing model predictions for a simulated voucher program in control markets with data from the treatment group. The results in this paper are based on the first step (using only control data) and the draft serves as a pre-commitment to the model estimates and predictions before examining the experimental data.

*We thank Joseph Altonji, Dennis Epple, Adam Kapor, Robert Miller, Ronni Pavan, Cristián Sánchez, Kegon Tan, Petra Todd, Christopher Walters, and seminar participants at Barcelona GSE, Clemson, Cleveland FRB, Cornell, Duke, ITAM, NYU, SUNY Buffalo, UC San Diego, ASSAs/Econometric Society, SOLE, and the Banff Empirical Micro Workshop 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

Interventions are typically evaluated in terms of impacts on outcomes alone. For example, student test scores or school completion rates are often examined to evaluate educational interventions, such as vouchers that defray the monetary cost of attending private schools. While more difficult to measure, program effects on economic well-being are of also of interest. Parents and students, for instance, view schools as differentiated products that represent bundles of attributes and amenities, a subset of which may represent anticipated impacts on student outcomes. A voucher may thereby facilitate better matches between households and schools. In the presence of ability-to-pay constraints, moreover, a private school voucher potentially improves welfare by allowing constrained households to choose their most preferred schooling alternative. Vouchers to attend private schools can also yield substantial reductions in the social costs of universal education, particularly where government provision is inefficient.

In this paper, we pursue a unique research design to estimate the welfare impacts of offering students in rural India voucher to attend private schools. The data are drawn from the Andhra Pradesh (AP) School Choice project, a randomized controlled trial of private school vouchers conducted in Andhra Pradesh. Our research design leverages the project's two-stage randomization: After eliciting initial interest in participation, the experiment randomized villages into treatment and control groups. The trial then paid the primary school tuition and fees at private schools for randomly selected students in the treatment villages. The data collection for households and schools that were randomized out 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 estimated welfare impacts of 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 model households' choice of primary school in this setting as potentially constrained by ability-to-pay. Although the data gathered for the AP School Choice project richly characterize households and schools, a key empirical challenge is separating households' willingness-to-pay from ability-to-pay, which is not observed in the data, in the observational control setting. Mispecification of choice sets is likely to underestimate the value of school choice. We therefore model latent ability-to-pay, placing restrictions on how information in the data about household wealth influences choice patterns. We compare this constrained model of school choice with flexible logit demand models, such as random coefficient, similar to those estimated in other school choice contexts (Neilson, 2013; Carneiro et al., 2016; Pathak and Shi, 2017; Abdulkadiroglu et al., 2017a).

In the research design's second step, we 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 students. The simulation generates predictions for take-up of the offer and the voucher elasticity of private schooling that we compare with analogous moments computed directly from the treatment village data. The simulation also generates predictions for which schools, in terms of their amenities and characteristics, that households offered a voucher will choose. These comparisons allow us to validate the choice models in terms of how well reproduce experimentally-generated choice patterns. We use the validated choice model to evaluate the voucher program implemented by the AP School Choice project and a counterfactual universal voucher in welfare terms.

In this draft, we present results from the first step of the research design and predictions for the experimental validation. The estimates provide evidence that households, particularly those targeted by the AP project, lack 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 choose any private school in their village. This constraint in turn has significant implications for demand and for of estimates of willingness-to-pay. For instance, while choice models that neglect differences across households in ability-to-pay estimate relatively lower preferences for school quality among asset-poor households (those that own less than three of six assets, e.g. a functional toilet), the empirical model we develop estimates few differences in willingness-to-pay across households. This suggests that unobserved liquidity and borrowing constraints can confound estimates of how lower socioeconomic households value school quality.

We estimate that the AP project voucher generated welfare value exceeding 9,000 Rs. on average to the students who are induced into private schooling by the program. This value represents around five times the average private school's annual tuition and about 40% median consumption per capita in this setting. This magnitude is also 100% larger than estimates obtained from choice models that do not consider constraints on households ability-to-pay. In fact, unconstrained choice models estimate that students who would have attended private school regardless ("always takers") place much greater value on the program than those who would not have otherwise attended a private school ("compliers"). As a result, the sole source of aggregate gains in models that neglect unobserved differences in ability-to-pay arises from potential reductions in spending on government schools. In contrast, the estimates of the constrained choice model that we develop indicate that – solely on the basis of the value to recipients created – a program that could be targeted to households who comply with a universal voucher would be socially justified.

For the research design's second step, we generate and report predictions from the empirical models for the experimental validation. To do this, we simulate a voucher trial in the models mimicking the actual AP project offer. This allows us to compare model-generated moments with corresponding ones obtained from the treatment data. We find that the ability-to-pay constrained model predicts the largest overall take-up of the AP voucher, with 57% of applicants using it to attend a private school. The unconstrained models predict take-up of 49%. The sharpest differences in predictions between models appears for lower socioeconomic households. The ability-to-pay constrained model predicts a 150% percent change in private schooling for asset-poor households due to the voucher. This elasticity is about 50% greater than the ones estimated by the demand models that do not account for differences across household in ability-to-pay.

Our findings are important for the evaluation of programs that expand school choice. Recent experimental and lottery-based estimates of voucher effects on students outcomes have produced limited evidence for beneficial impacts. For example, voucher lottery winners in the AP School Choice Project, from which we also draw our data, performed no better than losers on exams in math, English, Telugu, and science/social studies on average (Muralidharan and Sundararaman, 2015).[1] In addition, recent evidence from U.S. voucher programs shows strikingly negative impacts (Abdulkadiroglu et al., 2017b; Mills and Wolf, 2017).[2] Taken together, these findings thus raise the question of how parents and students value school alternatives and school choice programs. Our work speaks to this by applying revealed preference to evaluate private school vouchers in terms of welfare.[3] Our focus on estimating preferences over schools, which builds upon discrete choice models of demand for differentiated products (e.g. Berry et al. 1995; Nevo 2001; Petrin 2002), also connects with the design and evaluation of efficient school choice mechanisms (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005, 2015).

Our application also contributes to a growing literature that uses experimental data to test structural econometric models (Wise, 1985; Todd and Wolpin, 2006; Keane and Wolpin, 2007).[4] In contrast with other papers that take this approach, however, we estimate our model without access to the treatment data or foreknowledge of the experimental moments.[5] In this respect, our design is closest to Pathak and Shi (2017), who validate school choice models fit prior to a policy

[1]The lottery winners performed better in Hindi, a subject generally not taught by the government schools, however.

[2]Evidence from other international settings tends to show positive effects (Rouse and Barrow, 2009; Bettinger, 2011). Angrist et al. (2002) and Angrist et al. (2006) find that vouchers improve student outcomes in Colombia. In the U.S., Milwaukee's voucher program has produced evidence of positive effects in the past (Rouse, 1998). A large literature has also examined competitive effects of vouchers. Epple et al. (2017) provides a recent survey.

[3]This motivation relates to other work that examines beneficiaries' responses to and valuation of in-kind transfers (e.g. Moffitt 1989; Currie and Gahvari 2008; Cunha 2014).

[4]There is also a related body of work using experimental data to fit structural models, e.g. Duflo et al. (2012); Attanasio et al. (2012); Galiani et al. (2015); Lagakos et al. (2017).

[5]The primary moments that we use for validation, voucher take-up rates by subgroups of one cohort, are not reported in Muralidharan and Sundararaman (2015), which examines the impact of the AP School Choice Project on test scores. Muralidharan and Sundararaman (2015) report overall take-up (for both cohorts combined) in Figure II and differences in application and acceptance rates for select subgroups (though not by household asset level) in Table A.1.

change in Boston, with two important distinguishing elements: First, as is commonplace to demand applications, households in our setting make choices that depend in part on an endogenous variable – the tuition and fees charged by private schools. Second, rather than observed directly, credit constraints generate unobserved heterogeneity in households' choice sets. These features of the environment introduce challenges for empirically separating ability-to-pay and willingness-to-pay in the observational data.

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 we estimate are discussed in Section 4. In Section 5, we present the results, which include model estimates, fit, welfare calculations, and validation against the experimental outcomes. We conclude in Section 6.

2 Background and Data

We consider the welfare gains from offering vouchers to attend private primary schools in the context of rural India. While primary school enrollment is high across the country (over 96% of primary school-aged children enrolled), recent assessments reveal that more than 60% of children in rural areas aged six to fourteen read at just a second-grade level (ASER 2013). These low outcomes accompany wide perceptions of waste and inefficiency in government schools, characterized by high teacher absenteeism and multi-class teaching (Chaudhury et al., 2006; Kremer et al., 2005; Muralidharan et al., 2014). At the same time, private schooling has grown rapidly in recent years, with as many as 29% of students even in rural India attending a private school (ASER 2013).

One major policy response to these trends has been national support for private school vouchers. India's 2008 Right to Education Act, for example, included a provision requiring private schools to reserve 25% of seats for disadvantaged students, whom the government would reimburse fees and tuition for. In contrast with government schools, which are free to attend in India and provide mid-day meals to students, private schools are fee-charging – though typically low-cost and cater to students from low-income backgrounds. As private school students significantly outperform their counterparts in public schools on standardized assessments on average (Muralidharan and Kremer, 2007; Desai et al., 2009), a voucher removes any financial barriers that may impede access to higher quality education.

In this setting, such barriers may be significant, as indicated by the direct and indirect evidence for credit constraints (Rosenzweig and Wolpin, 1993; Townsend, 1994; Ravallion and Chaudhuri, 1997; Chaudhuri and Cherical, 2012; Banerjee and Duflo, 2014). For instance, Tarozzi et al. (2014) present evidence from a randomized controlled trial that micro consumer-loans substantially raised ownership and use of insecticide-treated bednets while demand was instead highly elastic when households had to instead pay upfront with cash. In our sample, 41% of students attending government primary school point to "economic reasons" as the determinant of their choice.

2.1 The Andhra Pradesh School Choice Project

Our data are drawn from a randomized controlled trial of private school vouchers conducted in 180 villages in Andhra Pradesh beginning in 2008. 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 and voluntary participation by private schools, who were not allowed to screen voucher students, was ensured by paying the voucher directly to schools' bank accounts.[6]

Our research design is facilitated by 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 relatively likely to attend government schools by conditioning eligibility on attendance in a government daycare (Anganwadi). After eliciting interest, the project then randomized villages into treatment and control market. 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).[7] However, 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 treatment variation to validate those models (and the estimated welfare impacts) against the experimental results.

Two "cohorts" of students were sampled for the project: a cohort of first graders whose primary school choice was collected at baseline and a younger cohort, whom we term kindergartners, that began primary schooling in years subsequent to the baseline survey. Though both cohorts' choices

[6]The 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.

[7]The village-level randomization provides control groups of private and public school students not exposed to voucher students. Muralidharan and Sundararaman (2015) find no evidence of spillovers on either group on students.

are used in estimation, one consequence of this is that we focus on the kindergarten cohort for the validation of the empirical models – those who, in the treatment condition, make their *initial* primary school choices with the AP project voucher.[8] A second consequence is that we do not always observe a choice of primary school for kindergarten students due to attrition. In addition, most effort by project enumerators was paid to tracking students (in treatment and control villages) who applied for the voucher. 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.

2.2 Sample and Data Summaries

Our estimation sample, drawn solely from the control data, contains detailed information for 4,251 households and 645 primary schools.[9] 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 (if any), characteristics of teachers, curriculum, and finances.[10] Geographic GPS locations were also collected, facilitating the mapping of travel distances between households and schools in their village.[11] Students were assessed at baseline and during the third and fourth years of the program.

Table 1 compares population characteristics of students that attend a government primary school with those that attend a private school, revealing several notable patterns.[12] First, while there are demographic differences across the sectors (e.g. a higher share of students in government schools are female), differences in the educational attainment of parents are especially dramatic. For example, 37% of both parents of students attending private schools completed primary school, with at least one parent having completed secondary school for 31% of private school students. In contrast, just 9% and 5%, respectively, of parents of students attending government schools have similar attainment. 42% of parents of government school students work as laborers (as compared with 16% for private school). 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 also shows

[8] The treated first graders, on the other hand, chose whether to use the voucher to switch from their baseline primary school for the next grade.

[9] We 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.

[10] When 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 distribution.

[11] This information was collected in year three of the program.

[12] The averages are representative of the population attending primary school (not of the estimation sample).

that households of students attending private schools have greater wealth on average, as indicated by asset ownership. Private schools students are more likely to live in a pucca house, have a water facility in the home, and to have a household toilet. 39% of students attending a government school have no more than two (of up to six) household assets. These differences in wealth and earning potential are suggestive of differences in ability-to-pay for private schooling. Finally, the table shows that private schooling has a high penetration rate across the 90 project villages, with an aggregate market share of 57% of primary school-going students.

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
$\geq$ 1 parent completed secondary	0.05	0.31	0.20
Both parents laborers	0.42	0.16	0.27
Math score	-0.01	0.61	0.35
Owns 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
Market share	0.43	0.57	
N	3,056	1,195	4,251

Notes: Table reports population average characteristics of households stratified by sector of school choice.

We present comparisons of the characteristics of government and private schools in our sample in Table 2. The total tuition and fees for attending government schools is essentially zero, but averages 1,800 Rs. per year for private schools. As a point of comparison, median total household consumption in comparable villages of Andhra Pradesh is around 84,000 Rs. (23,000 per capita) per the 2011-12 India Human Development survey. There is also considerable variation across private schools in tuition (the sample standard deviation is about 1,000 Rs.), with English medium private schools being around 1,000 Rs. more expensive on average. Government schools, unlike almost all private schools, provide mid-day meals for students. Government and private schools also differ in 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. The table also shows that a much greater share of teachers in government schools have a bachelor's degree. 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 and private schools also differ in instructional medium and their allocation of instruction time. 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 science (whereas few government schools do). These differences in medium and instructional time are indicative of horizontal differentiation.

Table 2: Characteristics of Primary Schools

	Government	Private
Tuition and fees (Rs.)	1.52	1814.42
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
Secondary school	0	0.25
Multi-class teaching	0.65	0.27
Share teachers absent	0.23	0.10
Share teachers with BA	0.78	0.57
Share teachers female	0.53	0.71
Share teachers from village	0.23	0.47
Offers Hindi instruction	0	0.42
Offers computer science	0.01	0.13
School value-added	-0.05	0.06
N	352	293

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

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 age 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, 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 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 alternatives that belong to their village. 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:

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

For any school, j, the household's consumption and tuition and fees, denoted p_j, must not exceed the household's ability-to-pay, which we denote by 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.[13]

Households rank the available schooling alternatives according to a quasi-linear indirect utility function. Letting α represent household i's marginal utility of consumption, the indirect utility to

[13]This "reduced-form" constraint also captures the possibility that households may be unable to commit to the schedule of private school tuition and fees due to uncertain income streams.

household i of school choice j can be written as:

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

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. ξ_j is regarded as an index of amenities of school j 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 may include household demographics, such as gender, caste, religion, as well as parental education. These characteristics therefore 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.

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. Misspecifying 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:

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

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 υ_i. We maintain the assumption that υ 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.

3.2 Logit Demand

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

$$\max_{j \in \mathcal{V}_i} U_{ij} \geq U_{ij'} \tag{4}$$

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 baseline 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:

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

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.[14] 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.

[14] We 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' willingesses-to-pay, we face a challenge commonplace to differentiated demand applications using observational data (Berry et al., 1995; Nevo, 2001): 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).[15] 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:

$$p_j = X_j'\Gamma + f(Z_j) + \mu_j \tag{7}$$

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.

For instruments, we take advantage of the spatial environment in our setting to construct predicted tuition for each private school based on the average tuition chosen by similar schools that are located in *other* villages. These predictions isolate cost shifters common to the schools (Hausman, 1996; Nevo, 2001). In implementation, we match private schools within their medium of instruction and focus on other schools not in nearby villages (to minimize the confounding influence of spatially-correlated demand shocks). In the Appendix, we describe the construction of these instruments and the first stage estimation in more detail.

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

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

probability can be expressed as:

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

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 two exclusion restrictions: First, 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. The second restriction we impose is that the non-asset characteristics of households (e.g. demographics and parental education) only enter W_i. Our baseline specification, however, allows the student's number of siblings, a measure of household size, to influence choices through either channel. For sensitivity checks, we also estimate several alternative specifications, outline in the next subsection.

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.[16] 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. This variable is

[16] We 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.

inherently unknown for students in the first grade cohort, however. To address this, we treat their eligibility as a latent variable in the estimation. 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 A4 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.[17]

As noted earlier, the baseline constrained model specifies ability-to-pay as a function of the first principal component of household assets and total siblings in the household. 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 A5 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 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 instrument construction 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. In the last subsection, we present predictions for the experimental validation of the models.

[17]Note 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.

5.1 Parameter Estimates and Elasticity of Demand

We present estimates of the parameters, elasticity of demand, and households' ability-to-pay for private schooling 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 specification 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).

Table 3: Estimates: Parameter on Tuition and fees

	CMNL	RC	CC
Tuition and fees (1000s of Rs.)	-1.76	-2.18	-0.99
	(0.21)	(0.26)	(0.11)
× Asset level = 2	0.24	0.47	
	(0.14)	(0.21)	
× Asset level = 3	0.42	0.74	
	(0.15)	(0.20)	
× Asset level = 4	0.63	1.13	
	(0.16)	(0.21)	
× Asset level > 4	0.35	0.82	
	(0.17)	(0.21)	
× Eligible for AP voucher	0.42	0.14	
	(0.10)	(0.12)	
First stage residual	1.27	1.41	1.46
	(0.14)	(0.16)	(0.12)
Random effect σ	1.67	2.33	2.64
	(0.20)	(0.22)	(0.13)

Notes: Table presents select parameter estimates as estimated by the clustered multinomial logit (CMNL), random coefficient (RC), and ability-to-pay contrained 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.

The estimates in Table 3 show that the parameter on tuition and fees is negative for all households across each of the models. 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 0.58 for the clustered multinomial logit model, indicated by the CMNL column. The random coefficient model estimates a lower elasticity at 0.40, while the ability-to-pay constrained specification estimates a similar overall elasticity of 0.36.

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

	CMNL	RC	CC
Overall	0.53	0.40	0.36
Female	0.59	0.44	0.40
Muslim	0.33	0.23	0.19
Lower caste	0.84	0.64	0.58
Older sibling in gov't school	0.94	0.72	0.72
Both parents completed primary school	0.18	0.14	0.12
≥ 1 parent completed secondary	0.15	0.11	0.08
Both parents laborers	0.84	0.62	0.54
Asset level < 3	1.15	0.90	0.63
Asset level $= 3$	0.56	0.42	0.42
Asset level $= 4$	0.25	0.19	0.28
Asset level > 4	0.21	0.13	0.16

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 contrained 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 model estimates 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 contrained model tend to be smaller on average. This finding is consistent with a marginal tuition reduction, which represents movement along the budget constraint, having little impact on the choices of constrained households.

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.[18] 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 13% cannot pay for the closest private school. 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 differences 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 2% of households with at least five assets are estimated to be unable to pay for private schooling, for example. On the other hand, around 23% and 30% of asset-poor households (those with fewer than three assets) are estimated to lack the ability-to-pay for any private school and their closest private school, respectively.

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

	Any	Closest	Highest-tuition
	Private School in Village		
Overall	0.10	0.13	0.83
Female	0.10	0.14	0.83
Muslim	0.09	0.10	0.80
Lower caste	0.15	0.19	0.90
Older sibling in gov't school	0.12	0.17	0.87
Both parents completed primary school	0.04	0.06	0.77
$\geq$ 1 parent completed secondary	0.04	0.05	0.70
Both parents laborers	0.15	0.20	0.88
Asset level < 3	0.23	0.30	0.88
Asset level $= 3$	0.10	0.14	0.80
Asset level $= 4$	0.04	0.06	0.80
Asset level > 4	0.02	0.02	0.84
Choose gov't school	0.24	0.25	0.94

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-

[18] Appendix Table A3 reports parameter estimates for the ability-to-pay function

to-pay for their closest private school according to the constrained choice model. Nearly all of these houesholds are estimate 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 characteristics, 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).[19]

Table 6 reports willingness-to-pay estimates by subgroup for English medium instruction, school value-added, and private schooling. 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 3,650 Rs. The estimated willingness-to-pay is lower for female students and considerably greater among higher 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 around 3,000 Rs. The clustered multinomial logit estimates differ from the other models in several ways. In particular, the average overall is much lower (around 780 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.

The middle three columns of Table 6 report households' estimated marginal willingness-to-pay for school value-added. Willingness-to-pay for a 1 standard deviation increase (on the student test score distribution) in value-added during primary school is around 1,600 Rs. on average according to the clustered and random coefficient models, but nearly 2,000 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 640 Rs. In comparison, the ability-to-pay constrained model estimates this subgroup's willingness-to-pay as three times larger (1,900 Rs.). Similarly, the clustered and random coefficient models estimate that asset-poor households tend to have

[19] Appendix D provides additional detail regarding calculating willingness-to-pay.

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	1.67	3.08	3.65	1.54	1.59	1.96	10.87	15.71	19.18
Female	-0.19	1.46	1.58	1.35	1.65	2.03	8.78	12.15	16.72
Muslim	6.60	7.52	9.08	0.98	1.59	1.75	18.36	26.46	28.34
Lower caste	0.95	2.01	2.79	0.67	1.63	1.90	0.99	0.58	5.01
Older sibling in gov't school	-1.67	2.38	2.78	2.15	1.63	1.99	-4.22	-1.94	-0.20
Both parents completed primary	6.35	6.76	8.08	0.82	1.68	2.03	25.60	35.54	39.50
$\geq$ 1 parent completed secondary	6.77	7.66	9.28	0.44	0.61	0.90	28.03	40.06	43.55
Both parents laborers	-2.61	1.59	1.97	1.27	1.65	2.06	-1.62	2.96	6.00
Asset level $<$ 3	-1.60	1.44	2.33	1.27	1.23	2.04	0.30	0.81	4.61
Asset level $=$ 3	-0.32	2.50	3.20	1.50	1.57	2.02	6.79	10.19	15.30
Asset level $=$ 4	3.90	4.42	4.11	2.08	2.09	1.94	15.72	23.17	23.76
Asset level $>$ 4	4.30	3.80	4.79	1.30	1.45	1.85	19.45	26.99	31.45

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 specification described in the text.

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 than 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 11,000 Rs. The random coefficient model estimates average willingness-to-pay for private schooling of more than 15,000 Rs. Willingness-to-pay per the ability-to-pay constrained model, however, is largest at 19,000 Rs. These estimates are economically large (representing up to 10 times annual private school tuition), but also correspond 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 a *negative* willingness-to-pay of over 4,000 Rs. among households in which both parents are laborers, whereas the ability-to-pay constrained model estimates a positive 5,000 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 exceeding 4,500 Rs. for this group. These differences across models have important implications for predicted take-up of a 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
$\ln L$	-5,322	-5,796	-5,814
AIC	11,619	11,723	11,755

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 specification 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 kindergarten cohort students eligible for the AP project voucher 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 for eligible students with an older sibling in government school (16% in the data), by Muslim students (58% in the data), and for those whose parents are laborers (24% in the data).

The models also tend to underpredict private attendance by asset-poor eligible households (16% in the data) 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	0.27
Muslim	0.58	0.47	0.54	0.52
Lower caste	0.20	0.17	0.20	0.20
Older sibling in gov't school	0.15	0.15	0.14	0.14
Both parents completed primary school	0.49	0.47	0.50	0.47
$\geq$ 1 parent completed secondary	0.56	0.52	0.55	0.54
Both parents laborers	0.24	0.20	0.23	0.23
Asset level < 3	0.23	0.19	0.21	0.20
Asset level $= 3$	0.33	0.24	0.29	0.28
Asset level $= 4$	0.39	0.37	0.39	0.37
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 specification 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 even 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.[20]

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

	CMNL	RC	CC
Gain in Consumer Surplus	3.20	3.19	4.74
− Cost of Voucher =	-1.11	-1.03	-0.20
+ Fiscal Externality =	3.19	2.75	5.39

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 surplus net of the expected voucher cost, while the last row adds the fiscal externality 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 about 3,000 Rs. The ability-to-pay model instead estimates the gain in social welfare at over 5,000 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 around 70% larger in magnitude. Table 9 also suggests one important source of this difference in aggregate welfare impacts: consumer surplus generated per applicant is around 50% larger according to the constrained model estimates.

While Table 9 reports impacts per applicant to quantity 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 of the models. 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 over 8,000 Rs. in consumer surplus. The logit demand models estimate that each recipient gains about 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

[20] Compensating 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 regarding calculating compensating variatin. This exercise thus abstracts from any equilibrium adjustments by schools, consistent with the scale of the actual trial.

exactly reversed when unobserved differences in ability-to-pay across households are accounted for.

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

	CMNL	RC	CC
Overall	6.52	6.48	8.28
Female	6.60	6.45	8.05
Muslim	8.40	8.36	10.35
Lower caste	6.22	6.13	8.02
Older sibling in gov't school	5.78	5.66	6.93
Both parents completed primary school	7.58	7.43	9.19
≥ 1 parent completed secondary	7.93	7.80	9.66
Both parents laborers	6.12	6.05	8.11
Asset level < 3	5.46	5.61	8.87
Asset level $= 3$	6.36	6.30	8.11
Asset level $= 4$	7.37	7.39	8.09
Asset level > 4	7.25	6.92	7.80

In Table 11, we further decompose the aggregate welfare impacts per recipient into two subgroups 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 school otherwise ("compliers"). Intuitively, the voucher offer is valued similarly to its cash value by always takers.[21] This is reflected in Table 11 by the relatively negligible values reported in the second row. 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 about over 9,000 Rs. (100% larger in magnitude). This value to compliers is estimated to be larger than the gain to the always takers in the ability-to-pay model, highlighting the impact of the constraint on their choice in the absence of the voucher. Table 11 also reports that the constrained model estimates a large share of compliers, a key prediction we will return to in the next section.

[21] Note that the share of always takers is pinned down by the control data, though the models differ somewhat in their fit to this moment.

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	0.19	0.17	0.25
Gain in Consumer Surplus	7.75	7.53	7.82	4.69	4.61	9.13
− Cost of Voucher =	-0.55	-0.80	-0.10	-5.25	-5.33	-0.59
+ Fiscal Externality =	-0.55	-0.80	-0.10	17.16	17.11	21.85

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. A growing literature studies education markets at scale (Neilson, 2013; Bau, 2017; Andrabi et al., 2017; Singleton, 2017; Sánchez, 2018).

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

	CMNL	RC	CC
Gain in Consumer Surplus	5.03	4.87	6.20
− Cost of Voucher =	-0.87	-0.85	0.03
+ Fiscal Externality =	1.66	1.62	3.61

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 over 50% larger according to the ability-to-pay constrained model (3,600 Rs vs. 1,700 Rs. per person). Whereas it is the magnitude of the fiscal externality alone that leads to aggregate gain per the logit demand models, the constrained choice model estimates a larger gain in consumer surplus.

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 2,700 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 against by the value of expanding the choice sets of other. Turning to the complier households, the logit demand models estimater their value from the voucher as far less than the expected cost and less than the value to the average always taker. The ability-to-pay constrained model finds the opposite relationship: it is the compliers who place relatively more value on the voucher offer. Moreover, the constrained model estimates indicate that a program targeted exclusively to the 16% of households who comply with a universal voucher would generate over 600 Rs. per recipient in social welfare just from the value to those households.

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	0.11	0.11	0.16
Gain in Consumer Surplus	7.34	7.13	7.57	5.25	5.30	10.65
− Cost of Voucher =	-2.67	-2.67	-0.48	-4.90	-4.92	0.63
+ Fiscal Externality =	-2.67	-2.67	-0.48	17.51	17.52	23.06

5.5 Experimental Validation

In this subsection, 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 in the control markets using the models that mimics the eligibility and conditions of the program.[22] This allows us to generate model-based "treatment" moments to compare directly with analogous moments calculated from the treatment group.

We focus on three main sets of moments for the experimental validation: First, we compare overall use of the voucher, as measured by the share of households who take-up the offer to attend a private school. We examine take-up by subgroup as well. Second, we compare voucher elasticities – the predicted percentage increase in private schooling due to the voucher – overall and by subgroup. Lastly, we examine predictions for how the voucher offer influences which schools, in terms of their characteristics, households choose.

5.5.1 Voucher Take-up and Elasticity

Model predictions for take-up of the voucher offer are presented in Table 14. The table also includes a column labeled "RCT" that will report the corresponding moments obtained from the treated data. The clustered and random coefficient logit models both predict that 49% of applicant households overall will choose to use the voucher. In contrast, the ability-to-pay constrained model

[22] The predictions are generated by setting tuition and fees at private schools to zero in the models and simulating households' choice of school in that counterfactual.

predicts that 57% 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 54% of asset-poor households will use the voucher, a share that is 9 percentage points higher than the other predictions made by the logit demand models.

Table 14: Validation: Take-Up of Voucher Offer

	RCT	CMNL	RC	CC
Overall		0.49	0.49	0.57
Female		0.48	0.47	0.54
Muslim		0.66	0.74	0.78
Lower caste		0.38	0.39	0.47
Older sibling in gov't school		0.32	0.28	0.35
Both parents completed primary school		0.66	0.67	0.72
≥ 1 parent completed secondary		0.68	0.71	0.76
Both parents laborers		0.41	0.43	0.52
Asset level < 3		0.45	0.45	0.54
Asset level $= 3$		0.45	0.48	0.57
Asset level $= 4$		0.51	0.50	0.57
Asset level > 4		0.60	0.59	0.65

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.

In addition to take-up, we examine voucher elasticities of private schooling – the percent change in private schooling due to the voucher offer – by subgroup in Table 15. The elasticity comparisons facilitate assessing how goodness-of-fit to the control data drives differences in the model predictions. The overall voucher elasticity is highest in the ability-to-pay constrained model, which predicts that private schooling will increase by 77% 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 similar elasticities for high asset households, but the ability-to-pay constrained elasticity for asset-poor households is around 50% larger in comparison. The ability-to-pay model estimates that asset-poor households will increase their private schooling nearly 150% under the voucher.

Table 15: Validation: Voucher Elasticity of Private Schooling

	RCT	CMNL	RC	CC
Overall		64	52	77
Female		65	55	81
Muslim		33	24	34
Lower caste		97	76	112
Older sibling in gov't school		109	102	147
Both parents completed primary school		37	32	48
$\geq$ 1 parent completed secondary		32	26	37
Both parents laborers		83	67	101
Asset level $<$ 3		112	97	145
Asset level = 3		67	54	80
Asset level = 4		44	30	57
Asset level $>$ 4		39	31	35

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.

5.5.2 Voucher Elasticities of Choice Characteristics

The empirical models also make predictions for the characteristics or the qualities of schools that households will choose under the voucher. Table 16 reports voucher elasticities for four school characteristics: 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 tuition and fees otherwise charged by households' chosen school will increase by 84% under the voucher on average according to the random coefficient model. The clustered model predicts a 95% increase, while the ability-to-pay model predicts nearly 120%. The ability-to-pay constrained model also predicts that attendance at English medium schools will approximately double (91%) due to the voucher and that travel distance to school will increase by 15%. Similar to the pattern for the elasticity of private schooling, these elasticities are largest per the ability-to-pay constrained model and generally smallest according to the random coefficient model. The table also reports predictions for the average change (in levels, not percent) in school-value added due to the voucher. Both the clustered and random coefficient models predict that the value-added of households' chosen schools will increase by 0.02 standard deviations, whereas the ability-to-pay constrained model predicts an increase of 0.03.

Table 16: Validation: Voucher Elasticities of Choice Characteristics

	RCT	CMNL	RC	CC
Tuition and fees		95	84	118
English medium		71	64	91
Distance to school (mi.)		10	11	15
School value-added		0.02	0.02	0.03

Notes: Table presents average voucher elasticities (percent change in characteristic of chosen school 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). Note that "elasticity" reported for school value-added is level (not percent) change. Predictions correspond to baseline specification described in the text.

6 Conclusion

Effects on economic well-being are especially of interest for programs 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 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 multinomial logit models of demand, including a random coefficient model.

In this draft, 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 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 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 induced into private schooling by 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. In particular, we simulate a voucher trial mimicking the experimental treatment offer and compare take-up and voucher elasticities. We find that the ability-to-pay constrained model predicts that 57% 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 about 9 percentage points lower. The sharpest differences between models in predictions appear for lower socioeconomic households, though. The ability-to-pay constrained model predicts a 50% greater voucher elasticity for asset-poor households.

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*.

Abdulkadiroğlu, 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.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. (2005). The Boston Public School Match. *The American Economic Review*, 95(2):368–371.

Abdulkadiroğlu, A. and Sönmez, T. (2003). School Choice: A Mechanism Design Approach. *The American Economic Review*, 93(3):729–747.

Andrabi, T., Das, J., Khwaja, A. I., Ozyurt, S., and Singh, N. (2017). Upping the Ante: The Equilibrium Effects of Unconditional Grants to Private Schools.

Angrist, J., Bettinger, E., Bloom, E., King, E., and Kremer, M. (2002). Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment. *The American Economic Review*, 92(5):1535–1558.

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.

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. 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.

Bau, N. (2017). School Competition and Product Differentiation.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890.

Bettinger, E. (2011). Educational Vouchers in International Contexts. In *Handbook of the Economics of Education*, volume 4, pages 551–572. Elsevier.

Carneiro, P., Das, J., and Reis, H. (2016). The Value of Private Schools: Evidence from Pakistan.

Chaudhuri, K. and Cherical, M. M. (2012). Credit rationing in rural credit markets of India. *Applied Economics*, 44(7):803–812.

Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., and Rogers, F. H. (2006). Missing in Action: Teacher and Health Worker Absence in Developing Countries. *The Journal of Economic Perspectives*, 20(1):91–116.

Cunha, J. M. (2014). Testing Paternalism: Cash versus In-Kind Transfers. *American Economic Journal: Applied Economics*, 6(2):195–230.

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.

Desai, S., Dubey, A., Vanneman, R., and Banerji, R. (2009). The Rise of Private Schooling in India. *India Policy Forum*, 5.

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.

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. (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.

Kremer, M., Chaudhury, N., Rogers, F. H., Muralidharan, K., and Hammer, J. (2005). Teacher Absence in India: A Snapshot. *Journal of the European Economic Association*, 3(2-3):658–667.

Lagakos, D., Mobarak, M., and Waugh, M. E. (2017). The Welfare Effects of Encouraging Rural-Urban Migration.

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.

Moffitt, R. (1989). Estimating the Value of an In-Kind Transfer: The Case of Food Stamps. *Econometrica*, 57(2):385–409.

Muralidharan, K., Das, J., Holla, A., and Mohpal, A. (2014). The Fiscal Cost of Weak Governance: Evidence from Teacher Absence in India. page n/a.

Muralidharan, K. and Kremer, M. (2007). Public and Private Schools in Rural India.

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.

Ravallion, M. and Chaudhuri, S. (1997). Risk and Insurance in Village India: Comment. *Econometrica*, 65(1):171–184.

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.

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.

Rouse, C. E. and Barrow, L. (2009). School Vouchers and Student Achievement: Recent Evidence and Remaining Questions. *Annual Review of Economics*, 1(1):17–42.

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 A1, 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 or applied for the AP voucher.

Table A1: 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:

$$A_{ijt} = \rho(A_{it-1}) + H_i'\pi + \omega_{j(i,t)} + \zeta_{ijt} \tag{8}$$

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 (8) 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 we cannot compute 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. $\kappa\hat{\mu}_j + \eta_j$ (where κ is the correlation between ξ_j and μ_j and η_j is a normally distributed random effect) then replaces μ_j in the indirect utility function.

We instrument private school tuition and fees using the tuition and fees of similar private schools in other villages. The main instrument we construct first separates private schools by their language of instruction. Then, the tuition and fees of the ten most similar same-language private schools in other districts in terms of spending per pupil are averaged. We weight the matched private schools by their similarity to the focal school when averaging. We match on spending per pupil on the assumption that this will capture common cost (though not demand) shocks. We examine sensitivity to this instrument construction by comparing results with an alternative identification of similar schools. The alternative uses k-means clustering to group same-language private schools based on their first principal facilities factor, teaching quality factor, and a factor summarizing horizontal differentiation (e.g. whether offers Hindi). Table A2 below presents first-stage estimates and diagnostics for both instrumental variables.

Table A2: First Stage: Private School Tuition and Fees

	(1) Baseline IV	(2) Alternative IV
Predicted tuition and fees	0.535***	0.469***
	(0.131)	(0.139)
F statistic	16.59	11.37
R-squared	0.363	0.351

Notes: Table presents first stage estimates that regress private school tuition and fees on school characteristics and an instrumental variable. Baseline IV refers to instrument constructed from similar per pupil spending same-language schools in other schools, while Alternative IV use k-means to group same-language schools based on facilities, teacher quality, and kinds of instruction offer. Though not reported, regressions control for school characteristics in choice models and include district intercepts. N = 293 private schools. Standard errors reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1

D Willingness-to-Pay and Compensating Variation Calculations

Calculating willingness-to-pay and compensating variation of a voucher requires scaling the change 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 the parameter estimates 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)). 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.

E Likelihood Functions

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

$$L^m = \sum_i w_i \ln[e_i \mathsf{L}^m_{1i}(\theta) + (1 - e_i)\mathsf{L}^m_{0i}(\theta)] \tag{9}$$

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 $\mathsf{L}^m_{1i}(\theta)$ 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

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

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

We estimate equation (9) using the EM algorithm. At the θ-maximization step, we maximize:

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

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:

$$\tilde{e}^m_i = \frac{e_i \mathsf{L}^m_{1i}(\theta)}{e_i \mathsf{L}^m_{1i}(\theta) + (1 - e_i)\mathsf{L}^m_{0i}(\theta)}$$

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

E.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.

E.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:

$$P_{ij}(j_i^*) = \frac{\mathbf{1}\{p_j \leq p_{j_i^*}\} \exp U_{ij}}{\sum_{j \in \mathcal{V}_i} \mathbf{1}\{p_j \leq p_{j_i^*}\} \exp U_{ij}} \tag{10}$$

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:

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

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

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

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

$$\mathsf{L}_i^{CC}(\theta) = \sum_{j_i^*} \phi_{ij^*} \prod_{j \in \mathcal{V}_i} P_{ij}(j_i^*)^{d_{ij}} \tag{11}$$

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.

F Additional Tables

Table A3: Estimates: Ability-to-pay function

	CC
Constant	3.32
	(0.68)
First asset factor	1.17
	(0.28)
Eligible for AP voucher	-1.51
	(0.44)
σ	1.55
	(0.34)

Notes: Table presents parameter estimates of ability-to-pay function for constrained (CC) model. Estimates correspond to baseline specifications described in the text. Baseline specification also includes total siblings in household, not reported in the table.

Table A4: Preference Heterogeneity in Baseline Model Specifications

CMNL	RC	CC	School characteristic	Interactions
X	X		Tution 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 A5: Estimates and Predictions for Alternative Specifications of Ability-to-Pay Constrained Model

Specification	Demand ϵ	Constraint	Surplus	Fit	Take-up	Voucher ϵ
Baseline	0.36	0.24	4.74	-0.02	0.57	77
Instead uses k-means to construct tuition and fees IV	0.42	0.25	4.85	-0.02	0.60	87
Includes whether lower caste in preferences and ability-to-pay	0.37	0.23	4.73	-0.02	0.57	77
Includes whether parents are laborers in preferences and ability-to-pay	0.36	0.25	4.81	-0.02	0.58	78
Allows preferences for private schools to be cohort-specific	0.36	0.22	4.68	-0.02	0.57	77
Does not include eligibility for AP voucher in the model	0.45	0.35	6.26	0.10	0.71	60
Does not re-weight for sample attrition in estimation	0.35	0.27	4.73	-0.03	0.57	79

Notes: Table reports select estimates and predictions for baseline ability-to-pay constrained model and several alternative specifications. Constraint column refers to probability that government-attending students cannot choose any private school. 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.