Experimentally Validating Welfare Evaluation of School Vouchers*

Peter Arcidiacono†
Duke University, IZA & NBER

Karthik Muralidharan‡
UC San Diego, NBER & JPAL

John D. Singleton§
University of Rochester, Hoover Institution & NBER

March 2024

Preliminary and incomplete; please do not cite or circulate

Abstract

We leverage a unique two-stage experiment that randomized access to private school vouchers across markets as well as students to estimate the revealed preference value of school choice. To do this, we first estimate several choice models on data only from control markets and without access to the treatment data, which are used for validation. This exercise reveals that a model where school choice absent the voucher is constrained by ability-to-pay achieves better out-of-sample fit but nonetheless underpredicts true voucher take-up. We then present evidence from treatment markets that suggests: a) the voucher offer also induced search; and b) that private schools passed through the surplus as kickbacks. Further, we show that a model incorporating each of these features and that is estimated on all of the data successfully explains the take-up patterns. Estimates from that model imply that each dollar spent on the program raised the average recipients's welfare by around 84 cents.

*We thank Chao Fu, Adam Kapor, Petra Todd, and seminar participants at Barcelona, Federal Reserve Bank of Minneapolis, Princeton, Toulouse, University of Pennsylvania, and Yale for helpful comments and discussions.

†psarcidi@econ.duke.edu
‡kamurali@ucsd.edu
§john.singleton@rochester.edu

1 Introduction

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

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

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

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

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

In the first step of our research design, we estimate empirical models of school choice solely using control group data from the trial (i.e. from villages randomized-out in the first stage). We consider two classes of discrete choice demand models. The first are random coefficient-style logit models that are standard for welfare analysis in the industrial organization literature (Berry, Levinsohn and Pakes, 1995; Nevo, 2001; Petrin, 2002). These models have been applied to several other contexts of school choice (e.g. Neilson 2013). Second, we develop an empirical model that incorporates a constraint on households' ability-to-pay for private schooling absent a voucher. The constraint reflects the reality that liquidity and access to credit are often limited in low-income settings, such as rural India.[3]

The second step of the research design experimentally validates the choice models using the data from the treatment markets. We do this by comparing model predictions for take-up of the voucher offer (and other outcomes) with analogous moments computed directly from the data collected from students randomly offered a voucher. A feature of our design is that the estimation was completed while blind to the treatment group data and the model predictions previously reported and pre-committed to in working paper Arcidiacono et al. (2021). Motivated by longstanding concerns with the credibility of structural econometric estimates, our design—by holding out the treatment data from estimation for the purpose of validation—aims to test the performance of model-based approaches to welfare and policy analysis using observational data.

Both of our first-stage models predict substantial takeup of the voucher. Off a base of 27% private school attendance absent the voucher, the random coefficients model predicts and additional 29 percentage point increase in private school attendance; the similar number for the credit constrained model is 38 percentage points. Yet both are substantial underpredictions: the actual increase was 58 percentage points. While one might expect this to be due to underestimating price sensitivity, both models especially underpredict the rate at which voucher winners attend *low tuition* private schools. And while both models do a good job predicting behavior of non-applicants in treatment villages, both models underpredict private school attendance of *voucher losers* by about 15 percentage points.

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

We show that adding two features to our model, which supplemental evidence from treatment markets suggests were important mechanisms at play, can reconcile these findings. First, because the voucher was set a level substantially higher than the tuition at most private schools and they received this higher amount for enrolling voucher students, schools had strong incentives to enroll voucher students—all the more so for schools with low tuition. This raises the question of whether private schools passed through the surplus as kickbacks in some way and we present evidence indicating siblings of voucher children received scholarships. Motivated by this, we adjust the model to allow voucher households and private schools to split any surplus generated by the program. Second, we adjust the model to allow for search, consistent with evidence from the treatment data of elevated private school attendance among households who *ex post* did not receive the voucher for unanticipated reasons. In this unified model, voucher losers expect that they may get a voucher and search accordingly, as do voucher winners. In contrast, those who did not apply for the voucher may not search because the expected benefits of doing so—given their financial situation—is such that it is not worth the effort. Incorporating these two added mechanisms into the ability-to-pay constrained model—an additional three model parameters—and estimating on all of the data (control + treatment markets) allows us to reasonably match the patterns in the treatment villages.

With the results of this unified model, we calculate welfare impacts of the program. We first show that the program raised voucher recipients' welfare by essentially as much as its cost. Thus, any fiscal externality that can be realized by reducing spending on government schools—or impacts on outcomes a paternalistic policymaker additionally values or producer surplus generated—is gravy on top of that. In this calculation, the surplus of the voucher amount over private schools' tuition influences search and choices and enters both the benefit (as part of consumer surplus) and program cost sides of the equation. Thus, we can ask instead what the welfare impacts of a voucher program that did not allow kickbacks would look like. In that case, the ratio of total consumer surplus to the expected cost is lower at about 0.84. However, in this non-distorted scenario, the average student drawn into a private school—i.e. the average complier—actually values the voucher 25% more than the average student who attends a private school regardless because their choice of schools is otherwise more constrained. The present value of the voucher offer to these compliers is equal to about 7% of median annual household consumption.

2 Background and Research Design

Our data are drawn from a randomized controlled trial conducted in 180 villages in the Indian state of Andhra Pradesh (AP). The AP School Choice project was designed to study the impact of private school vouchers on student learning outcomes. This was motivated by evidence of large differences in skills measures between students at low-cost private schools and government schools.[4] Villages selected for the project had to have at least one private school that agreed to participate. Across project villages at baseline (2008), more than one of every two primary school students (57%) attended a private school.[5]

Students randomized into the treatment condition were offered a voucher covering the costs of tuition and associated fees or expenses (e.g. books and uniforms) at government-recognized private schools in their village for the duration of primary schooling. At the average private school in the project data, tuition and fees are otherwise about 1,900 Rs. per year, as shown in Table A2. This amount represents nearly 8% of median annual consumption per capita in this setting.[6] Expenses for transportation, however, were not covered by the voucher and, unlike government schools, private schools do not provide free mid-day meals. Beyond costs of attendance, the bundles of characteristics associated with private and government school also differ on average in notable ways. Table A2 highlights that private schools tend to have lower absenteeism (of less qualified teachers on average) and that most feature some English instruction.[7]

The program was intentionally targeted to students likely to otherwise attend government schools. For an older cohort at baseline, this meant they were currently attending a government for first grade; for a younger cohort, which we term "kindergartners," eligibility for the program was conditioned on attending a government daycare (Anganwadi). Table A1 compares the characteristics of this group with the characteristics of the private-attending and government-attending

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

[5]This figure is obtained from survey data based on those households with children in the sample age range in the AP project districts (ASER, 2018). The high penetration largely reflects the requirement that villages have at least one private school—the unconditional private school market share among primary school-going students in AP in 2008 was around 33%.

[6]Median household expenditure was 86,000 Rs. per the 2011-12 India Human Development survey in comparable rural villages (with a private school) of Andhra Pradesh.

[7]Many private schools also allocate class time to teaching Hindi (the national language). Government school instruction is entirely in the local language, Telugu.

populations, showing that students eligible for the program are more similar in background demographics and socioeconomic status to those in a government school at baseline.[8]

Participation in the project was voluntary at the school level, but participating private schools were not allowed to screen or selectively admit voucher students.[9] The annual voucher value was set at around the 90th percentile of the fee distribution of private schools in the study sample. For each voucher recipient verifiably enrolled, two years' worth (5,200 Rs.) was paid up front and directly to schools' bank accounts.

2.1 Research Design

An important feature of the AP School Choice project is its two-stage randomization: At baseline, parents of eligible students were invited to apply for the program with the knowledge that the voucher would be allocated by lottery and that applying would not guarantee receipt. After eliciting interest from eligible households, the project first randomized 180 villages into into 90 treatment and and 90 control markets. Applicant households in treatment villages were then randomized into or out of the voucher treatment group in the second stage.[10] Data collection at baseline and post-intervention was consistent across treatment and control markets.

Our research design leverages this market-level randomization and is visually represented in Figure 1. We first fit several alternative empirical models of primary school choice to only the data from control markets (shown by the light blue shading in the figure). The control models are detailed in the next section. The control market data are purely observational; we observe school choices made by households, characteristics of those households, and attributes of the school options (including the tuition and fees). As shown in the figure, control market estimation combines information from first graders, whose retrospective primary school choice is used, and kindergartners who chose subsequent to the baseline survey. The kindergarten sample includes three subgroups: those who were eligible and applied for the voucher, those who who were eligible and did not apply, and those who were ineligble.[11]

[8]As Table A1 shows, private school students are, among other things, more likely to have parents who both completed primary school; more likely to have a parent who completed secondary school; more likely to live in a pucca (brick or stone) house, have a water facility in the home, and to have a household toilet.

[9]The design stipulated that, similar to charter schools in the US, lotteries would be held to allocate places in oversubscribed schools.

[10]This double randomization design facilitated estimating spillover effects on non-participants in the program.

[11]Combining all these subsamples in estimation presents three practical challenges, which are discussed some later

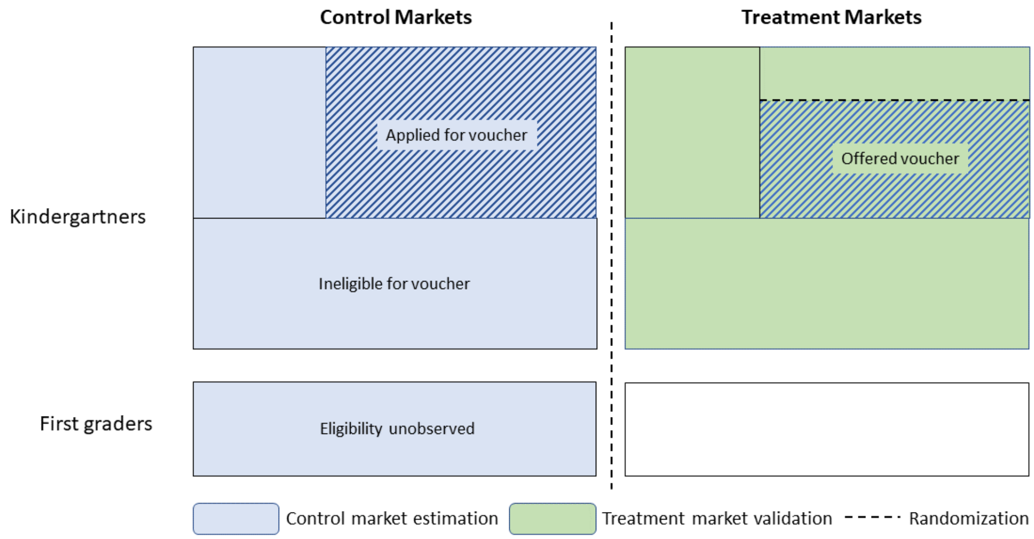

Figure 1: Control Market Sample and Treatment Market Validation

We then use data from the treatment markets for validation of the control models. For kindergartners in treatment villages that did not receive a voucher (e.g., those ineligible), we evaluate how the models fit out-of-sample under the assumption the program did not impact their choices. For households that instead received a voucher, we evaluate the models based on their predictions for choice patterns—e.g., what share of applicants would take-up the voucher offer. We do this by counterfactually setting tuition and fees at private schools to zero in the models. This experimental validation step is visually represented in the figure by the boxes overlayed with upward-sloping diagonal lines. Key to our design is that these model predictions were pre-committed to and generated without access to treatment market data.

3 Control Models and Results

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

(and fully in Arcidiacono et al. 2021): 1. the trial's endogenous sampling design; 2. sample attrition of kindergartners; and 3. unobserved eligibility status among first graders. We constructed weights to deal with the first two and handled the third by treating it as a latent type in estimation.

as well as human capital gains.

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

3.1 Ability-to-Pay Constrained Choice

In selecting a primary school, households weigh the utility of the school alternatives that belong to their village.[12] This set is denoted by $\mathcal{V}_i$ for household i. However, the tuition and fees may exceed the household's ability-to-pay. This is captured in the model through a constraint on their choice problem:

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

For any school, j, the household's consumption and tuition and fees, denoted p_j, must not exceed the household's ability-to-pay, which we denote by ω_i. For government schools, p_j is zero (or nearly so). The ability-to-pay constraint represents the combination of a household's income and any liquid wealth, such as accumulated savings, with their ability to borrow against future income to finance private schooling.[13]

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

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

D_{ij} is the distance between school j and household i's home, while X_j represents school charac-

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

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

teristics. $Closest_{ij}$ allows that the closest school, if a government school, is especially salient.[14] In estimation, we include in X_j whether a school is government or private, is government recognized (if private), is English medium, offers Hindi classes, is connected to a secondary school, and three indices respectively capturing the quality of facilities, of teachers, and the characteristics of teachers.[15] Also contained in X_j is school j's value-added in math, which we estimate from the panel of student test scores.[16] ξ_j represents an index of commonly-valued amenities of school j unobserved to the econometrician. ϵ_{ij} 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:

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

The household characteristics in W_i mediate the valuation households place on school amenities, capturing systematic heterogeneity across households in willingness-to-pay. W_i includes eligiblity status and indicators for gender, whether belongs to a scheduled caste, is Muslim, whether an older sibling attends government school, whether both parents completed primary school, and whether one parent completed secondary school.[17] Note we do not include assets in W_i; this information enters the model via the ability-to-pay constrained (discussed next).

3.1.1 Instrumenting for Private School Tuition and Fees

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

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

[14] We include an indicator for cases when distance is missing.

[15] We also include indicators for imputation of tuition and fees and for whether value-added information is missing.

[16] Appendix B of Arcidiacono et al. (2021) details the value-added estimation.

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

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

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

where $\hat{\mu}_j$ is the first stage residual for private school j and e_j is a normally-distributed random effect.[18]

Our baseline specification uses two instruments. First, we use a summary measure of each private school's location in "product space" (Berry, Levinsohn and Pakes, 1995). We do this using factor analysis applied to totals of characteristics of *other* schools in the village for each private school, e.g. the number of other English-medium schools. The second instrument uses the spatial environment to isolate exogenous cost differences across private schools (Hausman, 1996; Nevo, 2001). We construct the predicted tuition for each private school based on the average tuition chosen by similar private schools that are located in *other* villages.[19] Arcidiacono et al. (2021) provides additional details on construction of the instruments, implementation of the control function, and reports first stage estimates in Table A19.

3.1.2 Identifying and Estimating Ability-to-Pay

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

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

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

In this equation, the household's log ability-to-pay at the time of choosing a primary school depends

[18] Note that both terms are zero for government schools.

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

on the wealth factors, I_i, and unobservable household-specific v_i. We assume that v is distributed normally, with variance σ, and independent of the choice shocks. Our baseline model specification includes the household asset factor, an indicator for eligibility for the voucher program, and household size in I_i.

As this feature of the model is new, we briefly discuss estimation via maximum likelihood. Interested readers are referred to Arcidiacono et al. (2021) for more details. The basic insight is to recognize that each i can fall into one of a finite number of possible choice sets. Let j_i^* index schools in i's village in terms of ascending tuition and fees (such that J_i^* is the most expensive). Then denote by ϕ_{ij^*} the probability that household i is in choice state of being able to afford at most: $p_{j_i^*} \leq Y_i \leq p_{j_i^*+1}$. We can write this as:

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

where the state probability is a difference between points on the normal CDF that depend on data (tuitions and I_i) and parameters (λ and σ). $\Phi\left(\frac{\ln p_1 - I_i'\lambda}{\sigma}\right)$ is the probability of not being able to choose *any* private school in their village.[20] Combining logit expressions for choice probabilities with the state probabilities, the likelihood function takes the form:

$$\mathsf{L}_i(\theta) = \sum_{j_i^*} \phi_{ij_i^*} \prod_{j \in \mathcal{V}_i} P_{ij}(j_i^*)^{d_{ij}} \tag{5}$$

where $P_{ij}(j_i^*)$ is the probability that i chooses private school j in their village given they belong to choice state j_i^* and is zero for js whose tuition and fees exceed i's ability-to-pay.[21]

3.2 Random Coefficient

We compare the latent ability-to-pay model with random coefficient models similar to classic demand estimation applications (e.g. Berry, Levinsohn and Pakes 1995; Nevo 2001; Petrin 2002) and the models of school choice in Neilson (2013) and Carneiro, Das and Reis (2016).[22] In this class of

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

[21] The full model we estimate includes two additional sources of unobserved heterogeneity that are not reflected in equation (5): the private school-specific random effects v_j (which we use Monte Carlo integration for) and unobserved eligibility status among first graders (which we model as a latent type and iterate on using the EM algorithm).

[22] Note that in Part I we also estimated a "clustered" multinomial logit demand model. We do not consider that model here as its estimates and predictions are largely in-line with the random coefficient model.

models, the underlying choice problem is unconstrained—households are able to choose from any primary school in their village:

$$\max_{j \in \mathcal{V}_i} U_{ij} \geq U_{ij'} \tag{6}$$

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

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

$$\begin{aligned} U_{ij} &= -\alpha_i p_j + X_j' \beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \xi_j + \epsilon_{ij} \\ &= -\alpha_i p_j + X_j' \beta_i + \gamma_i \ln D_{ij} + \delta Closest_{ij} + \kappa \hat{\mu}_j + e_j + \epsilon_{ij} \end{aligned} \tag{7}$$

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

Like the ability-to-pay constrained model, the random coefficient model specifies a parametric relationship between observed household characteristics 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 again includes observed household characteristics. However, the random coefficient model includes an additional stochastic component on household preferences for private schooling. Letting β_i^P indicate the marginal utility to household i of attending private school, this parameter can be expressed as:

$$\beta_i^P = \beta_1^P + \beta_2^P W_i + \nu_i \tag{8}$$

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

3.3 Control Estimation and Results

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

Figure 1 presents a visual representation of our research design. Estimation on the control data pools observations from several subgroups of students, shown with the light blue shading: kindergartners who were eligible (by virtue of attending an Anganwadi at baseline) and who applied for the voucher program; eligible kindergartners who did not apply; ineligible kindergartners; and first graders (whose retrospective choice of primary school we use in estimation). Though the model specifications allow for preferences (and ability-to-pay, in the constrained model case) to depend on eligibility, we do not model application status.[23] However, since eligibility status is unknown for this older cohort, we model latent eligibility of these students (and use the EM algorithm in estimation). Full control market estimation details, including how we construct weights to account for the endogenous sampling design, are provided in the Appendices of Arcidiacono et al. (2021).

Table 1: Estimates: Selected Parameters—Control Models

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

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

4 Treatment Data

Table 2: Private Schooling and Tuition and fees by Subgroup

	Attend Private		Tuition|Private	
	Control	Treat	Control	Treat
First graders	0.57	0.58	1.71	1.82
Ineligible for voucher	0.99	0.99	1.79	1.87
Eligible non-applicants	0.22	0.16	1.58	1.95
Applicants not offered voucher	0.32	0.43	1.85	2.12
Voucher winners		0.80		2.12

4.1 Data Cleaning

In this subsection, we describe cleaning the treatment market data. Data collection in treatment markets mirrored the data collection in control markets (with one exception, discussed below). We therefore process the data in the exact same way, producing a cleaned dataset that connects each student to a village-specific set of primary schools.[24] The data contain observable characteristics of students and household (e.g. whether belongs to a scheduled caste) as well as characteristics of primary schools in the village (e.g. whether English is the medium of instruction). For latent factor variables, such as the asset index, we impute their values in the treatment data using the relationships between characteristics and factors in the control markets. Similarly, we use the first stages for tuition and fee endogeneity estimated on the control data to impute residuals for treatment market private schools.

We restrict the sample to students for whom a specific choice of school in their village at baseline is recorded in the tracking data. This is to mirror restrictions made for the control sample and to match the fact that the control models restrict choices to schools in the student's baseline village. This yields a total sample of 629 younger cohort students (i.e. "kindergarteners") who were randomly offered a voucher in treatment villages.

The randomization and symmetric data collection (and processing) imply that there should be baseline balance on average between 1) control and treatment market schools; 2) control and treatment market households who applied for the voucher; and 3) treatment market applicant households randomly offered and randomly not offered a voucher. A concern for the household

[24] This processing is described in fuller detail in Part 1.

comparisons (and for the experimental validation to come), however, is sample attrition. We note below that students offered a voucher are less likely to attrit (than control applicants) and discuss how we address this issue.

4.1.1 Coding Voucher Take-Up

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

The project team collected information about actual take-up of the voucher offer as reflected in payments to private schools as well as mitigating circumstances in cases on non-use. About 66% of the 629 students in the cleaned sample who were offered a voucher actually used it. Note that this number closely matches the figure stated in Muralidharan and Sundararaman (2015).

However, we do not focus on this variable for purposes of the validation. Rather, we focus on whether a student *intended to take-up* the voucher offer (and whether this intention would be, or otherwise is, also reflected in the tracking data). This decision is motivated by keeping in mind the validation exercise, which takes the data and sets tuition to zero in the models estimated on control markets to predict to take-up. The predictions thus correspond to choices—as they would appear in the tracking data—under the voucher provided there were no extenuating circumstances.

We infer intention to use from the coding of actual voucher use and from tracking information. Appendix Table presents the full details for how individual observations are coded. Informally, we view our intended use variable as an upper bound. We code as "intended to use" students who appear in private schools in tracking data but who did not receive a voucher (implying they paid tuition).[25] However, our analyses to come do not presume that such students would have used the voucher at the private school they attend.

Because no (or very few) students actually used a voucher in several treatment market villages, likely due to private school non-compliance, we remove all offered students residing in nine of the 90 treatment data villages from the sample.[26] This leaves a sample of 574 individual households who

[25] Another subgroup of students were unable to use a voucher by virtue of being too young; we assume all of this group intended to use a voucher.

[26] The project team also indicated several private schools who reneged on participating; we thus do not set tuition and fees to 0 at these schools when generating model predictions.

were randomly offered a voucher in treatment villages. Of these, 489 (85%) intended to take-up the voucher offer.[27] The next section dives into these choice data for the experimental validation.

4.1.2 Selective Attrition

The attrition rate, calculated as the share of households at baseline with valid tracking data, is noticeably smaller for households in treatment markets offered a voucher (11%) than it is for control market applicants (19%). This suggests that the voucher offer, by attracting students to private schools in their village, induced households to be more likely to stay in the final sample. We thus adjust model predictions and estimates based on the treatment data to account for selective attrition.

To do this, we first solve for the number of households that would have attrited from the final offered student sample *in the absence of the voucher offer*. The calculation assumes that the attrition rates of applicants between treatment and control markets would be the same in the absence of the offer and comes to 70 of the 574 students. We then assume that the 70 students who otherwise would have attrited also belong to the subgroup of students who actually took-up the voucher offer.

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

5 Treatment Validation

We then turn to evaluating the empirical models out-of-sample performance using the held-out data from treatment markets. The treatment data allow for two kinds of validation: non-experimental

[27] *Actual* use as reflected by voucher payments, by virtue of removing students in the non-complying villages, is higher for the final sample at 72%.

and experimental. These can be understood visually from Figure 1. In the treatment data, we have several subgroups of kindergarten households who did not receive a voucher offer: those who were eligible and applied, but randomized out at the household level; those eligible who did not apply; and the ineligible. We can therefore ask how the models estimated on the control data do in explaining the choice patterns of households in treatment markets also in the "control" condition.

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

5.1 Non-Experimental Validation

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

Table 3: Non-Experimental Validation: Treatment Market Predictions

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

Table 3 shows the results of non-experimentally validating the models out-of-sample. Both the random coefficient model and the ability-to-pay constrained model match well the private school attendance rates for ineligibles and eligible households who didn't apply for the voucher program. However, both models significantly underpredict private school attendance of voucher losers by nearly 15 percentage points.

To further examine the fit of our models to these groups, we formulate the question of misspecification as hypotheses tests. To do so, we begin by fixing u_{ij} for each option j in treatment models to that predicted from the control model estimation. For empirical model m:

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

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

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

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

Table 4: Non-Experimental Validation: Hypotheses Tests

	RC		CC	
	(1)	(2)	(3)	(4)
Private school		-0.12		0.07
		(0.50)		(0.09)
Private school × Voucher loser		2.21		2.05
		(0.56)		(0.30)
Tuition and fees		0.00		-0.11
		(0.10)		(0.09)
AIC	2399	2260	2411	2265

The results are presented in Table 4. Consistent with Table 3, the estimates of the private dummy are not different from zero for ineligibles or for non-applicants, but it is significantly positive for voucher losers. At the same time, the coefficient on price is small and insignificant, suggesting that the control estimates are providing good estimates of the tuition gradient.

Overall, the non-experimental validation suggests that both control models provide a good fit for the treatment data for those who did not apply for a voucher—with the important exception of applicants who were randomized out from receiving a voucher offer at the household-level. That both models miss for this group raises the question, which we turn to in the next section, of whether and how the voucher intervention nonetheless influenced their choices in some way.

5.2 Experimental Validation

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

5.2.1 Voucher Take-up

Table 5 presents actual and model-predicted take-up of the voucher offer. The first column reports attendance at private schools by applicant kindergartners in control markets. Overall, 27% of applicants in control markets choose to attend a participating private school. Columns (3) and (4) then report model predictions for voucher take-up according to the random coefficient and ability-to-pay constrained model, respectively.[28]

The random coefficient model predicts that private school attendance under the voucher will increase 50%. The ability-to-pay constrained model predicts that private school attendance will more than double, increasing another 10 points. Across households, this gap is pretty uniform, though is only six points among those where both parents completed primary school. Among households where a parent completed secondary school, the ability-to-pay constrained model underpredicts by only 2 points. Columns (5) and (6) of Table 5 then adjust the model predictions for selective attrition—that the voucher offer induced fewer students to attrit. This correction raises the predictions to 56% and 65% take-up of the voucher offer, respectively. As Appendix Table A4 reports,

[28] Note these predictions do not exactly match those reported in Arcidiacono et al. (2021). This is because Table 5 re-computes the predictions on the students offered the voucher in the treatment markets, whereas the original predictions were computed for applicants in control markets. These predictions are thus adjusted for minor treatment-control differences in observables. They also account for non-participation in the program by some schools. There is a second reason the predictions are different (and generally a little lower), which is that Arcidiacono et al. (2021) simulated take-up allowing households to use a voucher at non-recognized private schools. In practice, the voucher could only be used at government-recognized private schools.

Table 5: Experimental Validation: Take-up of Voucher Offer

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

65% take-up (as estimated by the constrained model after adjusting for less attrition) represents a predicted elasticity of private schooling with respect to the voucher offer of about 150%.

Column (2) of 5 reports take-up of the voucher in the treatment markets—what actually happened. As reported earlier, 85% of applicants randomly offered the voucher used (or intended to use it) to attend a participating private school. Compared with the random coefficient model prediction, this represents a gap of 29 points. The ability-to-pay constrained model's prediction was also too low, but by 9 fewer points. The subgroups comparisons show that the models performed especially badly at predicting take-up of students with an older sibling in government school. The ability-to-pay constrained model was off by 30 points for this group. This reflects that the control market estimates of both models assign a significant disutility to attending a private school for this group of students and suggests model mis-specification. Table A3 presents comparisons restricting the moments to just those without an older sibling in government school. This reduces the data-prediction gap in the case of the constrained model to 15 points.[29]

Table A4 compares model predictions for elasticities of private schooling with respect to the voucher offer with those computed using the experimental variation. These comparisons likewise reveal that the models generally underpredict—albeit the ability-to-pay constrained less so—and reveal a data pattern that the constrained model better captured. The table shows that the voucher elasticity is highest for the low asset households and lowest for the high asset households. Both

[29]While these comparisons presented represent the baseline specification of the models, these gaps are robust across the different specifications estimated on the control markets data.

models match this, but the difference in the elasticity between the low and high asset households is matched more closely by the ability-to-pay constrained model. Finally, Table A5 compares effects of the voucher offer on characteristics of households' chosen schools in terms of treatment-control differences in levels and elasticities. As expected, the offer raised tuition at chosen schools (by about 1000 Rs. on average) and increase attendance at an English medium school by 13 points. It also increased attendance at a school offering Hindi by 33 points. Both models underpredict the effect on Hindi, but produce similar ITT effects as the experiment on English and tuition. This is despite underpredicting private school attendance significantly and thus suggest the models overvalue English and overpredict use of the voucher at high tuition schools.

5.2.2 Hypothesis Tests

This section examines model fit and mis-specification by estimating auxiliary models on the choices of the treatment group (applicants in treatment markets randomly offered a voucher). Specifically, we begin by estimating models of the following form:

$$u_{ij}^m = \hat{u}_{ij}^m + \hat{\alpha}_i^m p_j + \pi_V^m Private_j \tag{9}$$

for each empirical model m estimated on the control sample. $\hat{u}_{ij}^m + \hat{\alpha}_i^m p_j$ is treated household i's predicted indirect utility from choice j, according to the estimates from model m. Like before, if model m accurately captures treated students' take-up of the voucher offer (i.e. their preferences over private schooling), we expect that $\pi_V^m = 0$.

Table 6: Experimental Validation: Hypothesis Tests Comparing Random Coefficient and Ability-to-pay Constrained Models

		RC			CC	
	(1)	(2)	(3)	(4)	(5)	(6)
Private voucher school		4.72	7.49		2.60	5.28
		(0.30)	(0.46)		(0.22)	(0.40)
Tuition and fees (@ voucher school)			-1.32			-1.32
			(0.17)			(0.16)
$\hat{\text{Use}}$	0.56	0.84	0.84	0.65	0.84	0.84
AIC	1496	1198	1135	1400	1235	1164

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

Columns (5) and (9) of Table 6 simultaneously estimate an intercept for voucher-eligible private schools and a "slope" on tuition at those schools:

$$u_{ij}^{m} = \hat{u}_{ij}^{m} + \hat{\alpha}_{i}^{m} p_j + \pi_V^m Private_j + \tau_V^m p_j$$

The result is surprising: while both models under-predict voucher use, they *over*-predict usage at higher tuition private schools. In other words, offered students use the voucher at lower tuition schools than expected. Further, the coefficient on tuition is remarkably similar between models and columns (5) and (10) show this pattern holds across levels of household wealth. This finding is key for understanding the sources of mis-specification in the control market estimates. In particular, it suggests that conventional unobserved school characteristics (i.e. insufficiently addressing tuition endogeneity in the control markets) are not the issue. This is because offered students do not have to pay the tuition, so the presence of school unobservables unaccounted for by the control models would instead predict a positive slope on tuition. Rather, if there is an unobservable school "quality" at play, it is negatively correlated with tuition.

We use the hypothesis testing framework to examine several other kinds of mis-specification pre-committed to in Arcidiacono et al. (2021) (Table 17). These tests are restricted to the ability-to-pay constrained model. Column (7) of Table A6 adds interactions between students' baseline math scores and school characteristics—a private school intercept, whether English medium, estimated math valued-added, and whether offers Hindi instruction—to the model to test for ability sorting. The control models did not include ability heterogeneity. The results suggest that some mis-specification may come from greater take-up among higher ability students, but higher ability students actually "prefer" lower value-added schools (and vice versa). Column (8) the allows for the possibility that offered students value voucher-eligible private schools' attributes differently than implied by the control models. These results indicate higher disutility of travel to voucher schools, much weaker preferences for English medium instruction and for value-added, and greater preferences for Hindi classes. While interesting, the inclusion of these covariates does little to explain the overall under-prediction of voucher take-up nor does their inclusion meaningfully modify the negative coefficient on voucher school tuition.

6 Unified Model

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

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

6.1 What Was Missing?

6.1.1 Search

That voucher losers in treatment villages—in contrast to eligible who did not apply and ineligibles—enrolled in private schools at much higher rates than what the control models predicted suggests the intervention impacted their choice in some way. While numerous possibilities exist, our proposed explanation is that voucher losers expected they would get a voucher too and, hence, searched for private school options.[30] For many, search revealed a high match quality, so they were willing-to-pay anyhow when it was later revealed they had to pay tuition.

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

Our principal evidence for this explanation comes from a handful of treatment market villages where, in the end, no household was able to actually use a voucher to go to private school.[31] Call these villages "flagged." We then estimate a linear probability model where the dependent variable is private school attendance. We control for whether the student applied for the voucher interacted with treatment village and flagged villages, whether they were offered a voucher interacted with flagged villages, where they were eligible for the voucher.

Table 7: Private School Attendance in Non-compliant Villages

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

The results are presented in Table 7. The interaction between applying for a voucher and treatment village is large and positive (and matches the evidence shown earlier of a 15 point discrepancy), as is the effect of winning a voucher. But what is interesting is what happens in flagged treatment villages: Namely, we no longer see an effect of winning a voucher nor do we see that flagged villages have private school attendance that is any different from other villages for non-applicants. Yet, both voucher winners and voucher losers attend private schools at similar rates to voucher losers in other treatment villages and correspondingly attend at higher rates than applicants in control villages.

[31] Likely because the government-recognized private school(s) backed out of the program.

6.1.2 Kickbacks

Why do voucher winners appear to prefer low tuition private schools? We propose pass through of the voucher surplus an the explanation. Such kickbacks make rational sense given the program's design: the voucher's yearly value was set at 2,600 Rs., which is about 44% more than the annual tuition and fees charged by the average private school.[32] This implies that, for those private schools with a sticker price below the voucher value, replacing a non-voucher student with a voucher student would be worth up to at least 2.6-p_j. A profit-maximizing school would thus try to attract voucher students by sharing the surplus generated and, importantly, this incentive should be stronger for lower tuition private schools.

Table 8: Voucher Pass-through: Survey Responses for Focal Child and their Siblings

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

While rational, a challenge for this explanation is the question how the surplus could feasibly be shared with voucher students. We present evidence this is achieved by offering scholarships to voucher students' siblings. Specifically, we examine post-intervention survey responses of households in control and treatment markets regarding private school attendance and their expenditure on tuition and fees. Importantly, the survey includes responses pertaining to the focal child, who did (treatment) or would have received a voucher (control), as well as for their siblings in the household. The top panel of Table 8 shows, as expected, that randomly offered households report

[32] Moreover, two years' worth was paid up front.

54 point greater private school attendance for the main child (column 1) and report spending about 600 Rs. less on the main child's tuition and fees (column 3). The bottom panel of the table reports analogous intent-to-treat estimates for school-aged siblings of the main child. The key finding is: the offer raises the probability their sibling attends private school by 15 points (column 1) without changing the household's spending on tuition and fees for the sibling child (column 3).[33]

6.2 Unified Model

In this subsection, we detail an empirical model that we then take to the entire dataset that has two new features:

1. Kickbacks—participating private schools in treatment villages share a fraction of the program's surplus with voucher recipient households;
2. Search—households must pay a cost to reveal their match qualities at private schools and all voucher applicants in treatment villages anticipate receiving a voucher.

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

$$u_{ij}^{V} = u_{ij} + \alpha p_j + \theta Surplus_j \tag{10}$$

where u_{ij} is the "control" utility previously given by equation (2). αp_j is added to this because these households anticipate recieving a voucher (recall that the coefficient on tuition in u_{ij} was α). $Surplus_j$ is then the per-recipient surplus generated by the voucher program, with θ—a new parameter to be estimated using the combined dataset—representing how it is shared with households. If $\theta = \alpha$, then the household captures all the surplus; if $\theta = 0$ then the private school captures all the surplus.

The voucher program paid two years' worth of 90th percentile of annual tuition and fees up front: 5,200 Rs. To impute $Surplus_j$, we thus consider the opportunity cost to each private school of accepting a voucher student and receiving this payment, which we take as instead using that

[33] The middle column of Table 8 (column 2) shows that, conditional on the focal child attending a private school, offered households report spending essentially zero on the main child's tuition and fees and report spending about 60% less than control households on their siblings' tuition and fees.

seat on a tuition-paying student. This is given by:

$$Surplus_j = (1+\kappa)(V - p_j) \times \mathbf{1}[V > p_j]$$

where $V = \frac{5.2}{1+\kappa}$ and κ represents the implied discount on the next year, which we set to the product of an annual discount factor (0.9) and the probability a student persists in attending (calculated as 0.79). This expression thus reflects that the surplus is both a function of the large amount paid and the fact that is was paid in advance and is zero for sufficiently high annual tuition private schools. It is intuitive to see from this equation that kickbacks to households will be larger at low-tuition private schools, potentially reconciling why more voucher winners than expected by the control models attend such schools.

The second extension is to introduce search. In particular, households have full information about government schools, but must pay a cost to reveal their match (represented by the preference shocks, the ϵs) with private schools. Denote by u_s and u_{ns} the expected utility of searching and not searching for information on private schools, respectively. Letting c_i represent the cost and G_i denote the set of government schools in i's village, applicants in treatment villages search when:

$$\begin{aligned} c_i &< u_s - u_{ns} \\ &< \ln\left(\sum_{j \in \mathcal{V}_i} \exp u_{ij}^V\right) - \ln\left(\sum_{j \in G_i} \exp u_{ij}\right) \\ &< -\ln(P_{iG|S}^V) \end{aligned} \tag{11}$$

where the equality in the second line follows from McFadden (1978) and where $P_{iG|S}^V$ is the probability i chooses a government school conditional on searching and receiving a voucher. In contrast, "control" households will search for private schools when

$$\begin{aligned} c_i &< \ln\left(\sum_{j \in \mathcal{V}_i} \mathbf{1}[p_j \leq \omega_i] \exp u_{ij}\right) - \ln\left(\sum_{j \in G_i} \exp u_{ij}\right) \\ &< -\ln(P_{iG|S}) \end{aligned} \tag{12}$$

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

It is this search channel that also provides a mechanism to explain higher private school attendance by applicants in treatment villages who did not receive a voucher. Specifically, we treat these households, consistent with the patterns presented earlier, as *expecting* to get the voucher, as in equation (11), resulting in increased search. Then, at the stage where they must make a decision as to which school to attend, they receive no kickbacks and must pay full price at participating private schools (i.e. their ex-post utility is given by u_{ij}, not u_{ij}^V). In the case of ineligible students, we assume they paid the search cost earlier; the reason they were ineligible for the voucher program is that they were attending a private school pre-kindergarten.

6.3 Unified Model Results and Fit

Table 9: Estimates: Selected Parameters—Control Ability-to-Pay and Unified Models

	Control	Unified
Tuition and fees (1000s of Rs.)	-1.28	-1.53
	(0.58)	(0.19)
First stage residual	1.77	1.62
	(0.63)	(0.22)
Private random effect σ	2.66	1.90
	(0.27)	(0.50)
Kickback (1000s of Rs.)		1.13
		(0.15)
Search		
Location		-0.19
		(0.09)
Scale		0.35
		(0.04)
Ability-to-pay constraint		
Intercept	2.96	3.43
	(0.55)	(0.68)
Eligible for AP voucher	-1.29	-0.92
	(0.41)	(0.37)
Asset factor	1.09	1.29
	(0.23)	(0.30)
σ	1.34	1.51
	(0.28)	(0.32)

Table 9 shows the estimation results incorporating these two features. The estimated parameters are similar to estimates of the control village alone (see Table XX) with the exception of the

Table 10: Estimates: Ability-to-pay Constraint and Search Probability

	Share unable to pay for...					
	Any private		Priciest private		Search privates	
	Control	Unified	Control	Unified	Control	Unified
First graders						
Overall	0.09	0.06	0.18	0.11	1.00	0.52
Lower caste	0.13	0.08	0.25	0.16	1.00	0.41
Parents completed primary	0.11	0.06	0.23	0.13	1.00	0.34
Asset level < 3	0.24	0.17	0.44	0.31	1.00	0.41
Asset level = 3	0.09	0.05	0.20	0.11	1.00	0.51
Asset level = 4	0.03	0.01	0.08	0.04	1.00	0.56
Asset level > 4	0.01	0.01	0.03	0.02	1.00	0.60
Voucher program applicants						
Control markets	0.13	0.08	0.25	0.15	1.00	0.46
Voucher losers	0.12	0.08	0.27	0.17	1.00	0.82
Voucher winners	0.17	0.09	0.33	0.20	1.00	0.81

Table 11: Unified Model Goodness-of-Fit

	Attend Private		Tuition|Private	
	Data	Unified	Data	Unified
First graders				
Overall	0.57	0.59	1.71	1.70
Lower caste	0.34	0.37	1.65	1.62
Parents completed primary	0.27	0.29	1.48	1.60
Asset level < 3	0.28	0.34	1.45	1.58
Asset level = 3	0.52	0.55	1.72	1.71
Asset level = 4	0.68	0.66	1.84	1.77
Asset level > 4	0.78	0.77	1.67	1.69
Voucher program applicants				
Control markets	0.34	0.34	1.88	1.66
Voucher losers	0.48	0.45	2.13	1.92
Voucher winners	0.81	0.79	2.09	2.07

coefficient on private school. This coefficient increases substantially due to the incorporation of search costs. With those who do not search ruled out from attending private school, the model rationalizes the observed private attendance by increasing this coefficient.

The coefficient on the kickback term is large and positive.[34] It is almost 74% of the magnitude of the coefficient on price, implying that voucher students are capturing almost 74% of the surplus from the voucher being set well above tuition at most schools.

The share of each group that paid the search costs as well as how binding the credit constraint is are shown in Table 10. Voucher winners are substantially more likely to search than applicants in control villages, who are in turn more likely to search than eligible non-applicants in either control or treatment villages. Voucher losers in treatment villages are also more likely to search as we treat them as expecting to receive a voucher at the search stage in order to reconcile their high rates of private school attendance.

Compared to our control model, the unified model has a credit constraint is less binding; the control model forces any search effects to operate through the credit constraint. Eight percent of applicants in both the treatment and control villages cannot afford in any private school and more than twenty-five percent cannot afford the most expensive private school.

Table 11 shows that these two additional features–kickbacks and search costs–significantly improve the fit of the model, both with regard to the rate at which different groups attend private school but also providing a better match with the posted tuition of the schools voucher winners attend. The first set of columns shows actual private school attendance for different groups of students, the predicted rates using the control model, and the predicted rates using the unified models. The predicted rates of private school attendance–both for voucher winners and voucher losers in treatment villages–now are within three percentage points of what is observed in the data. The second set of columns repeats the exercise but now focuses on tuition conditional on attendance. Kickbacks are important here, with expected posted tuition now much more in line with what is observed for voucher winners who attend private school.

[34] Recall that this coefficient is identified from the types of schools voucher winners attend relative to what we would expect based on the behavior of those in control villages.

6.4 Implications for Welfare from the Unified Model

In this subsection, we present estimates of welfare impacts per the unified model for the AP voucher program and for counterfactual programs of interest. Two major components for the total social welfare generated by a voucher program include: 1) the gain in consumer surplus to recipients of vouchers (in present value terms); and 2) the expected cost of financing the program, which enters social welfare negatively. Consumer surplus change is given by the added inverse of the estimated compensating variation—the amount of income that each household would need to be compensated to keep their utility level with the voucher the same as without it.[35] Additionally, we estimate the fiscal gain that would arise from re-allocating students out of government schooling to compute a net welfare change. Our calculations 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.

Table 12: Welfare Effects of AP Voucher

	Control	Unified
Gain in Consumer Surplus (1000s of Rs.)	3.46	5.81
Cost of Program	5.39	6.91
Fiscal Externality	4.73	7.20
Net Welfare Change	2.81	6.11

Table 12 presents estimates of the welfare impacts of the AP voucher program on recipients.[36] We compare estimates from the control models with those produced by the unified model. The first row shows that the present value of the voucher offer to the average household offered it is about 5,800 Rs. This is about 60% more than estimated by the ability-to-pay constrained model estimated on only the control data. Further, the unified model estimates that each dollar of program cost is worth about 84 cents to the average voucher recipient; the corresponding figure from the constrained control model is 60 cents. Since the control models underpredict takeup, the unified model also produces a considerably larger estimate of the fiscal externality from the program. The total surplus estimate from the unified model is over twice as large as that obtained from the control models.

[35] Appendix B details the calculation.

[36] In focusing on recipients, we exclude voucher losers—who searched for private schools under false pretenses—from the calculations. This corresponds to the idea of scaling the program to all those who applied for it.

Table 13: Welfare Effects of AP Voucher by Treatment Subgroup

	Always takers		Compliers	
	Control	Unified	Control	Unified
Share of Applicants	0.32	0.34	0.30	0.46
Gain in Consumer Surplus [without kickback]	5.48	9.09 [4.18]	5.60	5.90 [1.10]
Cost of Program [tuition expense]	8.61 [5.63]	8.61 [4.91]	8.61 [6.88]	8.61 [5.13]
Fiscal Externality	0	0	15.67	15.68

In Table 13, we decompose the welfare impacts by treatment subgroup: always takers, who would've chosen a private school even in the absence of the AP voucher program, and compliers, who are induced to switch to a private school by the voucher offer. One thing highlighted by the table is that the kickbacks raise the cost of the program (present value 8,600 Rs. per user) well above the amount of tuition actually paid for. A big fraction of the AP program's overall value to households is due to the kickbacks. While it stands to reason that always takers will value the tuition reduction at close to the tuition displaced, the table shows the kickback distorts compliers' choices such that the average household induced into private schooling by the offer values the non-kicbback aspects of their choice at almost one fifth of the associated tuition expense.

The findings in Table 13 motivate welfare evaluation of counterfactual, idealized voucher programs that instead do not generate or otherwise allow for kickbacks. Removing kickbacks has several likely effects on aggregate welfare: First, it predictably reduces take-up of the offer, mechanically lowering the size of the fiscal externality and, notwithstanding correction of its distortionary effects, also surplus per applicant. Both effects are reflected in the upper panel of Table 14, which keeps the targeting of the AP program in place: the complier share of applicants goes from 46% to 28%. Meanwhile, the consumer surplus gain to the average applicant shrinks by about 50%. Second, removing kickbacks stops distorting the choices of voucher users. This can be seen most clearly in the consumer surplus for the average complier with the targeted no kickbacks program, which is nearly four times greater than the surplus absent kickbacks to the average complier with the actual program.

The lower panel of Table 14 also implements a no-kickback voucher program, but instead expands eligibility to the entire population (not just those who applied to the AP program). Intuitively, the first order effect of this is just to make the program more "cash like" since the share of inframarginal always takers swells to nearly 60%. At the same time, another 20% of the population

Table 14: Welfare Effects of Counterfactual No-Kickbacks Programs (Unified Model)

	Overall	Always takers	Compliers
Targeted to AP voucher applicants			
Share of Applicants		0.34	0.28
Gain in Consumer Surplus	2.76	4.67	4.26
Tuition Expense	3.84	5.59	6.89
Fiscal Externality	4.35	0	15.68
Net Welfare Change	3.28	-0.92	13.05
Universal voucher			
Share of Population		0.58	0.20
Gain in Consumer Surplus	4.04	5.06	5.27
Tuition Expense	5.05	5.98	7.57
Fiscal Externality	3.09	0	15.68
Net Welfare Change	2.08	-0.92	13.38

would use this universal voucher to attend a private school instead of a government school. The interesting finding here is that the average household in this group, reflecting meaningful constraints on their school choice otherwise, would actually be willing to pay a little more than the average always taker would for the voucher program. The magnitude of the surplus gain to the average complier (5,200 Rs. in present value) is roughly equivalent to 6% of median annual household consumption.

7 Conclusion

References

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

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

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

Banerjee, Abhijit V., and Esther Duflo. 2014. "Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program." *The Review of Economic Studies*, 81(2): 572–607.

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

Carneiro, Pedro, Jishnu Das, and Hugo Reis. 2016. "The Value of Private Schools: Evidence from Pakistan."

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

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

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

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

McFadden, Daniel. 1978. "Modelling the Choice of Residential Location." In *Spatial Interaction and Planning Models.* , ed. F. Snickars and J. Weibull, 823–848. North-Holland, Amsterdam.

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

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

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

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

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

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

Rosenzweig, Mark R., and Kenneth I. Wolpin. 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.

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

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica*, 62(3): 539–591.

A Willingness-to-Pay and Compensating Variation Calculations

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

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

where m indexes the models (e.g. $m \in \{\text{RC,CC,unified}\}$), $\hat{\alpha}_i^m$ corresponds to estimated coefficients on tuition and fees (presented in Table 9), and δ is the effective annual discount factor. For δ, we use the product of 0.90 (a 10% annual discount rate) and 0.79 (the annual probability that a voucher recipient remains in private school). Note that we also use δ for calculating the costs and fiscal externalities of voucher program.

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.

B Additional Tables

Table A1: Summary Statistics: Household Characteristics by Subgroup with Balance Checks

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

Means refer to households in the 90 treated villages; differences are taken with respect to households in the 90 control villages.
** significant the 5% level; *** significant at the 1% level

Table A2: Summary Statistics: Characteristics of Primary Schools

	Government	Private
Tuition and fees (Rs.)	0.81	1,924
English medium	0.02	0.57
Unrecognized	0	0.23
Mid-day meals	0.99	0.03
Kitchen facility	0.26	0.01
Full pucca building	0.89	0.52
Library	0.94	0.77
Functional water tap	0.42	0.62
Functioning toilet	0.65	0.84
Separate toilet for girls	0.34	0.60
Staffroom for teachers	0.20	0.72
Playground	0.52	0.70
Has secondary school	0.00	0.27
Total school enrollment	74.28	286.18
Average teacher salary (Rs. / month)	16,959	2,127
Multi-class teaching	0.70	0.24
Pupil-teacher ratio	26.53	16.68
Share teachers absent	0.21	0.09
Share teachers with BA	0.78	0.54
Share teachers with formal certificate	0.90	0.16
Share teachers female	0.50	0.71
Share teachers lower caste	0.24	0.12
Share teachers Muslim	0.02	0.07
Share teachers from village	0.25	0.48
Offers Hindi instruction	0	0.44
Offers computer skills	0.01	0.13
School value-added	-0.04	0.04
N schools	686	570

Table A3: Validation: Take-up of Voucher Offer (no sibling in gov't school)

	RCT	RC	CC
Overall	0.89	0.66	0.74
Female	0.88	0.65	0.73
Muslim	0.97	0.86	0.90
Lower caste	0.80	0.57	0.67
Both parents completed primary school	0.91	0.79	0.83
$\geq$ 1 parent completed secondary	0.78	0.82	0.87
Both parents laborers	0.80	0.62	0.72
Asset level < 3	0.88	0.67	0.75
Asset level $= 3$	0.85	0.63	0.71
Asset level $= 4$	0.93	0.65	0.75
Asset level > 4	0.91	0.70	0.76

Table A4: Validation: Voucher Elasticity of Private Schooling

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

Table A5: Validation: Voucher Intent-to-Treat Effects and Elasticities on Characteristics of Chosen School

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

Table A6: Validation: Hypothesis Tests for Mis-specification of Ability-to-Pay Constrained Model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Private voucher school		2.60		5.28	4.53	4.70	4.58	3.98
		(0.22)		(0.40)	(0.42)	(0.50)	(0.43)	(0.48)
Private voucher school × Asset factor					0.03	-0.76		
					(0.31)	(0.63)		
Private voucher school × Older sibling in gov't school					1.63	0.89	1.72	1.74
					(0.42)	(0.81)	(0.43)	(0.45)
Tuition (@ voucher school)			0.52	-1.32	-1.34	-1.43	-1.37	-1.26
			(0.08)	(0.16)	(0.16)	(0.21)	(0.16)	(0.18)
Tuition × Asset factor						0.39		
						(0.26)		
Tuition × Older sibling in gov't school						0.36		
						(0.33)		
Private voucher school × Ability							0.47	0.27
							(0.26)	(0.29)
English × Ability							-0.37	-0.34
							(0.28)	(0.29)
VA × Ability							-1.08	-1.08
							(0.29)	(0.31)
Has Hindi × Ability							0.00	-0.12
							(0.32)	(0.34)
Distance × Private voucher								-0.70
								(0.19)
English × Private voucher								-0.66
								(0.32)
VA × Private voucher								-2.09
								(0.63)
Has Hindi × Private voucher								0.63
								(0.37)
Control function × Private voucher								-0.08
								(0.21)
$\hat{\text{Use}}$	0.65	0.84	0.74	0.84	0.84	0.84	0.84	0.84
AIC	1400	1235	1360	1164	1153	1153	1137	1105
Count R^2	0.56	0.77	0.67	0.77	0.79	0.79	0.80	0.79