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

February 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Preliminary and incomplete; please do not cite or circulate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 (by about 20 points). 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 the program raised the welfare of voucher recipients by nearly as much as its cost.

*We thank Chao Fu, Adam Kapor, Petra Todd, and seminar participants at Barcelona, Princeton, Toulouse, UPenn, 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 Data

2.1 The Andhra Pradesh School Choice Project

Our data are drawn from a randomized controlled trial conducted in 180 villages in the Indian state of Andhra Pradesh (AP) beginning in 2008. The AP School Choice project was designed to study the impact of school vouchers on student learning outcomes at both the student and market level. Students randomized into treatment were offered a voucher covering the costs of tuition and associated fees or expenses (e.g. books and uniforms) at private schools in their village for the duration of primary schooling (grades one through five). Expenses for transportation were not covered by the voucher. Villages selected for the project had to have at least one private school that agreed to participate in the project. Participation was voluntary at the school level, but schools were not allowed to screen or selectively admit voucher students.[4] The voucher value was set at around the 90th percentile of the fee distribution of private schools in the study sample, and was paid directly to schools' bank accounts.

Our research design uses the AP School Choice project's two-stage randomization. At baseline, parents of eligible students were invited to apply for the program with the knowledge that the voucher would be allocated by lottery and that applying would not guarantee receipt. The program was targeted to students likely to otherwise attend government schools. After eliciting interest, the project then randomized villages into treatment and control markets. Applicant households in treatment villages were then randomized into or out of the voucher treatment group. This double randomization design was important for estimating spillover effects on non-participants in the program (Muralidharan and Sundararaman, 2015). For this paper, the data collection for households and schools that were randomized out at the village-level provides a control group free of contamination from the experimental treatment. We use the data from the 90 control markets to fit empirical models of school choice, while we use the exogenous variation induced by the experiment to validate those models (and the estimated welfare impacts).

[4]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.

2.2 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.[5] 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 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.

2.2.1 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

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

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).[6] 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.[7] This leaves a sample of 574 individual households who were randomly offered a voucher in treatment villages. Of these, 489 (85%) intended to take-up the voucher offer.[8] The next section dives into these choice data for the experimental validation.

2.2.2 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

[6] 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.

[7] 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.

[8] *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%.

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.

3 Empirical Models

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

We compare the estimates and predictions for two classes of choice models. In the first, we explicitly model the the influence of an ability-to-pay constraint on choice. In relaxing this constraint, a private school voucher thereby potentially generates welfare benefits by expanding households' choice sets. We compare this model class, which 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.[9] 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:

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

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.[10]

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 characteristics. $Closest_{ij}$ allows that the closest school, if a government school, is especially salient.[11] 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.[12] Also contained in X_j is school j's value-added in math, which we estimate from the panel of student test scores.[13] ξ_j represents an index of commonly-valued amenities of school j unobserved

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

[10] 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.

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

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

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

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.[14] 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":

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

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.[15]

Our baseline specification uses two instruments. First, we use a summary measure of each

[14] 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).

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

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.[16] 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, Y_i, is inherently not contained in the data. This introduces unobserved heterogeneity across households in choice sets. Mis-specifying households' choice of school as unconstrained is liable to bias estimates of willingness-to-pay and underestimate the gains of a voucher.

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

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

[16] 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.

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:

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

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). $\Phi(\frac{\ln p_1 - I_i'\lambda}{\sigma})$ is the probability of not being able to choose *any* private school in their village.[17] 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.[18]

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).[19] In this class of models, the underlying choice problem is unconstrained—households are able to choose from any primary school in their village:

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

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

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

[18] 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).

[19] 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.

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

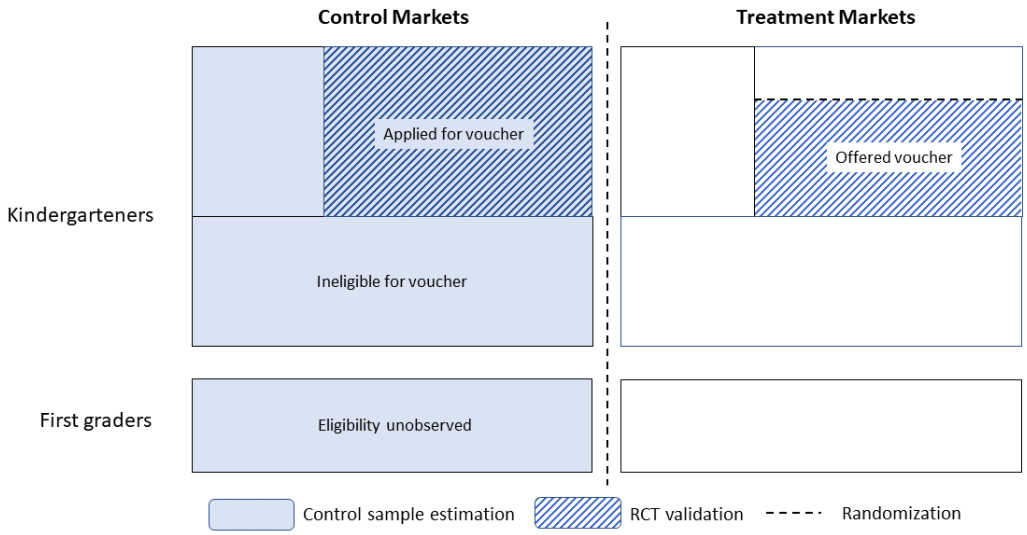


Figure 1: Control Market Sample and Treatment Market Validation

4 Validation

4.1 Control Market Estimation and Model Predictions

Per our research design, we estimate the empirical models using only data from the control markets. The estimation details and results are summarized here, with greater elaboration provided in Part 1. We use the models to generate predictions for school choices under the voucher in the treatment markets. The next section summarizes the results of the experimental validation.

Figure 1 presents a visual representation of our research design. As mentioned earlier, while the experimental validation focuses on the choices of kindergarten students in treatment villages who applied and were randomly offered the voucher, we use several other subgroups of students in estimation on the control data. These subgroups include kindergarten students who were eligible (by virtue of attending an Anganwadi at baseline) but who did not apply for the program as well as ineligible students. Though the model specifications allow for preferences (and ability-to-pay, in the constrained model case) to depend on eligibility, there is no parameter heterogeneity between eligible students with application status. Finally, note that we also use the retrospective school choices of older cohort students ("first graders") in the estimation. A complication this raises is that eligibility is unobserved for first graders. We therefore model latent eligibility for 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).

4.2 Treatment Market Validation

In this section, we compare predictions generated by the empirical models (estimated on data from the control villages) with experimental outcomes from the treatment markets. We first focus on out of sample validation by seeing how well our random coefficients and credit constrained models predict private school attendance and the tuition of the private schools attended for those in treatment villages who did not win a voucher, either because they were (i) ineligible, (ii) eligible but didn't apply, or (iii) applied but did not win a voucher. Both sets of models provide an excellent fit for those in groups (i) and (ii), (both for private school attendance and tuition conditional on attending). However, the control models significantly under-predict private school attendance of group (iii) (voucher losers).

Next we examine how well our models predict voucher take-up (that is, private school attendance) for voucher winners. Recall that both the random coefficients model predict substantial increases in private school attendance from the voucher, and that this was especially true for the credit-constrained model. Despite these high predicted take-up rates, we show that both models significantly under-predict private school attendance for voucher winners, especially the random coefficients model. Further, the under-prediction is especially prevalent at low tuition schools.

4.2.1 Out of Sample Validation

We begin by examining how well the model predicts behavior of those who do not have a voucher in treatment villages. This provides us with an out of sample fit for our model. To do this, we substitute in the data on treatment villages into the model and examine predicted private school attendance. We then compare these numbers to their actual counterparts. These comparisons are made for kindergartners in treatment villages for the following three groups: (i) ineligible, (ii) eligible but didn't apply, and (iii) apply but didn't win a voucher.

Table XX shows how well our two models match the private school attendance rates for these three groups. Both the random coefficients model and the credit constrained model match the private school attendance rates for ineligibles and eligible who didn't apply. However, both models

significantly under-predict private school attendance for voucher losers by over 14 percentage points.

To further examine the fit of our models, we begin by fixing u_{ij} for each option $\{i, j\}$ in t treatment villages to that predicted from the control model estimation. Label this utility as u_{ij}^C. We then estimate a model for kindergartners in treatment villages who do not win a voucher where the utility is specified as:

$$u_{ij}^T = u_{ij}^C + \sum_{k=1}^{3} g_{ik}\phi_{0k} + \phi_1 p_j \tag{9}$$

where g_{ik} are indicators for whether i belongs to each of the three groups (ineligible, eligible didn't apply, voucher loser). This specification allows us to see whether the overall private dummy is different for these groups (ϕ_{0k}) as well as whether the price coefficient is different in treatment villages.

Results are presented in Table XX. Consistent with Table XX, the estimates of the private dummy are not different from zero for ineligibles and for eligibles who do not apply, but is significantly positive for voucher losers. The coefficient on price is small and insignificant, suggesting the control estimates are providing good estimates of the tuition gradient.

Overall, our estimates suggest that the control model provides a good fit for the treatment data for those who did not apply for a voucher. The results also suggest, however, some effect of the voucher program on voucher losers, a feature we explore in later sections of the paper.

4.2.2 Experimental Validation

To generate the predictions, we set tuition and fees at private schools for applicants to zero in the models and simulate their choice of school under that counterfactual. 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.

Table 1 presents actual and model-predicted take-up of the voucher offer. The first column reports attendance at private schools by applicant kindergarteners in control markets. Overall, 27% of applicants in control markets choose to attend a government-recognized 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. Though qualitatively similar, these predictions do not exactly match those reported in Arcidiacono et al. (2021). This is because Table

1 re-computes the predictions on the cleaned sample of 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 treatment-control differences in observables.[20]

Table 1: Validation: Take-up of Voucher Offer

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

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 1 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 A2 reports, 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 1 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 pre-

[20]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.

diction, 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 A1 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.[21]

Table A2 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 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 A3 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.

4.2.3 Hypothesis Tests

This section takes a different perspective on model fit to the data under the voucher offer. For those kindergartners randomly offered a voucher in treatment markets, we estimate choice models of the following form:

$$U_{ij}^m = \hat{U}_{ij}^m + \delta^m Private_j \tag{10}$$

[21] 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.

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

$$\hat{U}_{ij}^m = -\hat{\alpha}_i^m p_j^T + X_j'\hat{\beta}_i^m + \hat{\gamma}_i^m \ln D_{ij} + \hat{\delta} Closest_{ij} + \hat{\xi}_j^m + \epsilon_{ij}$$

where p_j^T is set to zero at all government-recognized (and participating) private schools. This setup allows us to frame validation in terms of hypotheses tests. For example, if model m accurately captures treated students' take-up of the voucher offer (i.e. their preferences over private schooling). we expect that $\delta^m = 0$. It also allows us to understand sources of model mis-specification.

Table 2: Validation: Hypothesis Tests Comparing Random Coefficient and Ability-to-pay Constrained Models

	Random coefficient (RC)					Ability-to-pay constrained (CC)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Private voucher school		4.72		7.49	7.39		2.60		5.28	5.12
		(0.30)		(0.46)	(0.47)		(0.22)		(0.40)	(0.40)
Private voucher school × Asset factor					-0.96					-0.88
					(0.67)					(0.57)
Tuition and fees (@ voucher school)			0.86	-1.32	-1.29			0.52	-1.32	-1.25
			(0.10)	(0.17)	(0.17)			(0.08)	(0.16)	(0.16)
Tuition and fees × Asset factor					0.41					0.43
					(0.24)					(0.24)
$\hat{\text{Use}}$	0.56	0.84	0.69	0.84	0.84	0.65	0.84	0.74	0.84	0.84
AIC	1496	1198	1417	1135	1136	1400	1235	1360	1164	1165
Count R^2	0.46	0.80	0.63	0.81	0.80	0.56	0.77	0.67	0.77	0.77

Table 2 reports results that compare fit and mis-specification between the random coefficient and ability-to-pay constrained models. Columns (1) and (6) report measures of goodness-of-fit to offered students' choices under the voucher; the constrained model achieves a lower AIC and a higher "count" R^2.[22] 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 (3) and (8) instead insert tuition and fees at voucher schools (and remove the private school intercept). The results show there is greater switching to higher tuition schools (unconditionally) than either model predicts, though, again, the constrained model fits relatively better.

[22]This is the "count" R^2—the fraction whom use the voucher to attend private school and the model predicts a greater than 50% chance they attend private school under the voucher.

Columns (5) and (9) of Table 2 simultaneously estimate an intercept for voucher-eligible private schools and a "slope" on tuition at those schools. These estimates thus ask how the models do in predicting *which* private school an offered student chooses to attend (as a function of tuition). The result here 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 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.

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 A4 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.

In summary, the results in this section point to several important findings. First, the ability-to-pay constrained model, consistent with the predictions shown earlier, achieves relatively better fit to the experimental patterns. Nonetheless, a large gap between predicted take-up and experimental take-up persists; this gap is larger for students with a sibling in government school; and, although

the hypothesis tests identify several other sources of mis-specification, collectively they do not account for the major portion of the residual. A key question for welfare analysis is thus what this gap represents. In addition, the other major pattern that emerges from these results is that, conditional on take-up, the control models over-predict using the voucher at high-tuition schools. The next subsection proposes, provides supporting evidence for, and tests an explanation for this unexpected pattern.

4.2.4 Endogenous Response: Voucher Pass-through

There are several reasons to believe that pass through of the voucher is a major explanation for under-prediction of take-up. Recall that the voucher amount was paid directly to schools (not the households). Its annual value was also set at $V = 2,600$ Rs., which is about 44% more than the annual tuition and fees charged by the average private school. 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 $V - p_j$. 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. This raises the possibility that pass through can rationalize the negative slope on tuition and explain a portion of the private school intercept.

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 3 shows, as expected, that randomly offered households report 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).[23]

[23]The middle column of Table 3 (column 2) shows that, conditional on the focal child attending a private school,

Table 3: 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

4.2.5 Reconciling Results for Voucher Winners and Losers

The results in Section 4.2.1 showed that voucher losers in treatment villages–in contrast to eligibles who did not apply and ineligibles–enrolled in private schools at much higher rates than what the control models predicted. The results in Section 4.2.2 showed that the control models also underpredicted voucher take-up, especially at low-tuition schools. In this section we examine the behavior of voucher winners in villages that, in the end, did not have private schools that affected the voucher which we refer to as flagged villages.

In particular, we estimate a linear probability model where the dependent variable is private school attendance. We control for whether the student applied for the voucher interacted with treatment village and flagged villages, whether they were offered a voucher interacted with flagged villages, where they were eligible for the voucher.

Results are presented in Table 4. The interaction between applying for a voucher and treatment village is large and positive, as is the effect of winning a voucher.

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.

But what is interesting is what happens in flagged 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. This suggests the very possibility of a voucher changed behavior even if, in the end, the individual was not able to use the voucher. As we describe in Section 5.1, this is consistent with a model where individuals must first choose to search for information about private schools, with the possibility of having a voucher changing search behavior.

Table 4: 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)

N =2,960 kindergartners

5 Unified Model

The previous section showed that our model estimated on control villages (i) missed the supply-side incentives faced by schools and (ii) missed the increased enrollment in private schools for voucher losers in treatment villages. In this section, we develop and estimate an extended version of our model to capture these two features.

We then show that this extended model matches the moments we missed from using the control villages alone. Finally, we show how the unified model affects our welfare calculations.

5.1 Kickbacks and Search Costs

We extend the credit-constrained model in section 3.1 in two ways.[24] First, we treat supply-side incentives as 'kickbacks' whereby voucher winners enrolling in private schools receive some of the difference between the voucher amount and the tuition of the schools they attend. This difference is especially large for low tuition schools and will be able to reconcile why voucher winners attend such schools.

The utility of school j net of the preference shock (given in Equation (2)) is now modified to u_{ij}^V:

$$u_{ij}^V = u_{ij} + \alpha p_j \underbrace{\theta(V - p_j) \times \mathbf{1}(V > p_j)}_{\text{"kickback"}} \tag{11}$$

where $V - p_j$ is the surplus to be divided between the individual and the school from the voucher being higher than the tuition of the school. Recall that the coefficient on tuition in u_{ij} was α. If $\theta = \alpha$, then the individual would capture all the surplus; if $\theta = 0$ then the school captures all the surplus. By allowing for kickbacks we provide a mechanism to explain why individuals both use the voucher more than expected but at lower tuition schools.

The second extension is to introduce search costs. Namely, individuals only have expectations regarding their match (as measured by the preference shocks, the ϵ's) to particular schools. By paying the search cost, these preference shocks are revealed.

Absent the voucher, poorer households will be less likely to pay the search cost because many of the schools will be outside of their price range, limiting the benefits search. Having a voucher then has an effect on private school attendance both through searching (by increasing the expected gains from search) and through making private school more attractive conditional on searching.

It is this search channel that gives us a mechanism for higher private school attendance by voucher losers. Namely, we treat voucher losers as expecting to get the voucher, resulting in increased search. Then, at the stage where they must make a decision as to which school to attend, they face the actual tuition rather than the voucher-subsidized tuition.

Both voucher winners and voucher losers in treatment schools then expect–at the search stage– that they will have a voucher. We denote u_s and u_{ns} as the utility of searching and not searching for information on private schools, respectively. After the search stage, ϵ's for all schools are revealed

[24] We focus on the credit-constrained model as its fit was substantially better than the random coefficients model.

to those who searched; ϵ's for all government schools are revealed to those who did not. Denoting c_i as the cost of search and G as the set of government schools, individuals who expect a voucher search when:

$$\begin{aligned} c_i &< u_s - u_{ns} \\ &< \ln\left(\sum_j \exp(u_{ij}^V)\right) - \ln\left(\sum_{j\in G} \exp(u_{ij})\right) \\ &< -\ln(P_{iG|S}^V) \end{aligned} \tag{12}$$

where the equality in the second line follows from **?** and where $P_{iG|S}^V$ is the probability of attending a government school conditional on searching and having a voucher.

The search decision for those who are not expecting a voucher is similar except for two changes. First, individuals expect to pay tuition so u_{ij}^V is replaced by u_{ij}. Second, the first sum in Equation (12) is only over schools that the individual can afford. Finally, we assume that ineligible students have already paid the search costs at an earlier point; the reason they were ineligible is that they were attending a private school as kindergartners.

We specify search costs as following an XX distribution with scale parameter β. This implies that there are two parameters to be estimated, with the estimate of β coming from the coefficient on expected utility being normalized to one. In total, there are then three additional parameters in the unified model: one parameter for the kickback and two parameters for the search cost.

5.2 Unified Model Results and Fit

Panel A of Table XX shows the estimation results incorporating these two features (DO WE WANT TO ALSO ESTIMATE WITH RANDOM COEFFICIENTS AND HAVE ONE COLUMN FOR EACH, SHOWING AGAIN THAT THE MODEL FIT IS BETTER WITH THE SEARCH MODEL?). The estimated parameters are similar to estimates of the control village alone (see Table XX) with the exception of the 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.[25] It is XX% of the magnitude of

[25]Recall that this coefficient is identified from the types of schools voucher winners attend relative to what we

the coefficient on price, implying that voucher students are capturing XX% of the surplus from the voucher being set well above tuition at most schools.

The share of each group that paid the search costs is shown in Panel B of Table X. 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.

Table XX 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 X 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.

5.3 Implications for Welfare from the Unified Model

6 Conclusion

would expect based on the behavior of those in control villages.

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.

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.

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.

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 Tables and Figures

Table A1: 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
≥ 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 A2: 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 A3: 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 A4: 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