Estimating the Effects of Private School Vouchers in Multi-District Economies*

Maria Marta Ferreyra
Tepper School of Business
Carnegie Mellon University

July 16, 2005

Abstract

This paper estimates a general equilibrium model of school quality and household residential and school choice for economies with multiple public school districts and private (religious and non-sectarian) schools. The estimates, obtained through full-solution methods, are used to simulate two large-scale private school voucher programs in the Chicago metropolitan area: universal vouchers and vouchers restricted to non-sectarian schools. In the simulations, both programs increase private school enrollment and affect household residential choice. However, under non-sectarian vouchers private school enrollment expands less than under universal vouchers and religious school enrollment declines for large vouchers. Fewer households benefit from non-sectarian vouchers. (JEL I22, H73, H42, C51).

---

* This paper is based on my PhD dissertation at the University of Wisconsin-Madison. I wish to thank my advisors Derek Neal, Steve Durlauf and Phil Haile for their guidance and encouragement. Comments from the editor and the referees substantially improved this paper. I am also grateful to Peter Arcidiacono, Angela Dills, Dennis Epple, Arthur Goldberger, Ken Judd, John Kennan, Yuichi Kitamura, Grigory Kosenok, George Levi-Gayle, Bob Miller, Tom Nechyba, Uday Rajan, Richard Romano, Ananth Seshadri, Holger Sieg, Fallaw Sowell, Steve Spear, and M. Kathleen Thomas for many helpful ideas and suggestions. For their help with and insights on the technical aspects of this project, I am thankful to Javier Pena, R. Ravi, and Jeff Reminga. I benefited from comments and suggestions by seminar participants at Brown, Carnegie Mellon, Clemson, U. of Colorado-Boulder, Columbia Business School, Duke, U. of Georgia, U. of Oklahoma, Rutgers, U. of Illinois at Urbana-Champaign, U. of Virginia, and U. of Wisconsin, and participants at the 2002 CIRANO conference on education, Camp Resources X, the Stanford Conference on Empirical Boundaries for IO, Public Economics and Environmental Economics, the LSE conference on Economics of Cities, the 2003 Public Economic Theory conference, the 2003 Stanford Institute for Theoretical Economics, the 2003 European Meeting of the Econometric Society, the 2004 Southern Economic Association meetings, and the 2004 conference on Education and Economic Development at the Federal Reserve Bank of Cleveland. I owe special thanks to the Condor Teams at the UW-Madison Computer Science Department and the Pittsburgh Supercomputing Center, whose help has been essential to carry out this project. Financial support from the UW Department of Economics to purchase data, and from the UW Robock Award in Empirical Economics are gratefully acknowledged. Surendra Bagde provided excellent research assistance. All errors are mine.

1. Introduction

Private school vouchers play an important role in the debate about education reform in the United States. Vouchers, it is argued, give households the opportunity to enroll their children in private schools and access their preferred type of education. Whereas some households prefer private schools over the public schools in their metropolitan areas, they face budget constraints that restrict them to public schools. Although public schools have no explicit tuition, in metropolitan areas where public schools have residence requirements households must choose public schools and residences as *bundles,* whose costs are determined by housing prices and property taxes. Therefore, to gain access to their preferred public schools households might choose to live in places they would not have selected in the absence of bundling. Vouchers may break this bundling by allowing households to choose private schools, which have no residence requirements.

Thus, vouchers may not only give households more school choices, but also alter household residential decisions. As a result, public school districts may experience changes in their property values, school funding, and the composition of their student populations. To gain insight into the potential impact of large-scale private school voucher programs, in this paper I examine these general equilibrium effects by estimating a general equilibrium model that jointly determines school quality and household residential and school choices in an economy with multiple public school districts and private schools. I then use the parameter estimates to simulate two different voucher programs.

Since the voucher programs enacted to date in the United States have included a small number of voucher recipients and have often restricted the set of eligible private schools,[2] researchers have had insufficient data to evaluate the general equilibrium effects of potential large-scale voucher programs and

[2] Publicly funded voucher programs currently exist in Florida; Cleveland, Ohio; Utah (for students with special needs); Milwaukee, Wisconsin; and the District of Columbia. Voucher-like laws in Maine and Vermont provide school choice in towns without public schools. See www.ij.org and www.heritage.org.

have turned to simulation to investigate them.[3] Thus, I build upon Nechyba's (1999) theoretical work to develop a framework that is rich enough for empirical implementation and counterfactual analysis. My approach differs from Nechyba's in a number of important ways. First, I incorporate household idiosyncratic tastes for location and school choices. This allows for the plausible heterogeneity that creates, for instance, a strong attachment to particular suburban public schools or urban Catholic schools on the part of otherwise similar households. Moreover, this addition gives rise to an equilibrium that mixes households with heterogeneous income and school choices even in districts with no housing quality variation, a significant departure from Nechyba's model.

Second, I include household religious preferences and two types of private schools, Catholic and non-Catholic. Religious schools comprised 85 percent of the 1990 private school enrollment in high school grades (U.S. Department of Education (1992)). Furthermore, the considerable variation in private school markets among metropolitan areas is related to geographic differences in the distribution of religious affiliations. Religious preferences are thus relevant to the understanding of these markets although private school modeling thus far has not considered their role.[4]

The inclusion of religion in the model has important implications for voucher analysis, since it answers such questions as: what type of private schools would expand under vouchers and at what rate, how religious preferences would affect the distributional effects of vouchers, how vouchers would affect schools' religious compositions, where Catholics would choose to reside under a voucher regime, and the minimum voucher level necessary to make a household better off given its wealth and religious preferences. Furthermore, the use of publicly funded vouchers at religious schools is a contentious issue. Although the Supreme Court upheld voucher use at religious schools in Zelman v. Simmons-Harris (06.27.02), states can still choose whether or not to include religious schools when designing voucher

---

[3] See, for instance, Bearse, Glomm and Ravikumar (2001), Caucutt (2002), Cohen-Zada and Justman (2005), Epple and Romano (1998, 2003a), Fernandez and Rogerson (2003), Manski (1992), Nechyba (1999, 2000, 2003).

[4] In simultaneous work, Cohen-Zada and Justman (2005) have considered religious preferences in their voucher simulations by calibrating a single-district economy without household idiosyncratic preferences. A number of their qualitative results agree with mine.

programs.[5] My model is thus able to analyze the effects of prohibiting voucher use in religious schools. My particular focus is on Catholic schools, which comprise the largest and most homogeneous group of schools within the private school market,[6] and attract most of the religious school enrollment in the states that have debated the participation of religious schools in voucher programs. Moreover, Catholic schools have historically had a strong presence among inner-city low-income students, who have often received subsidized tuition and who are targeted by most current voucher proposals.

Third, I include non-residential property in the model for a better reflection of the environment that determines educational expenditure for public schools. Since an important policy issue concerning vouchers is their ability to improve educational outcomes for low-income students in the central city of metropolitan areas, the model must represent the sources and magnitude of public school spending in those districts, often endowed with large amounts of non-residential property.

Furthermore, in contrast with all other researchers who have examined the general equilibrium effects of vouchers by relying on calibrated numerical examples, I estimate my model. [7] Estimating the model allows me to investigate its empirical properties, most importantly whether it captures the relevant aspects of the reality presumably affected by the policies of interest. The estimation also reveals which dimensions of the data are well fitted and which are not, and how specific features of the model affect the parameter estimates and their accuracy as well as the fit of the data. Such information is, in turn, invaluable guidance for building models equipped to answer the desired policy questions.

---

[5] Many states have constitutional provisions ("Blaine amendments") with more prohibitive criteria for the separation of church and state than those found in the First Amendment. While state voucher programs were upheld by the Wisconsin and Ohio Supreme Courts when challenged on the basis of Blaine amendments, a similar case before the Florida Supreme Court is still pending. See Bolick (2003), Viteritti (1999), and www.ij.org for further reference.

[6] According to the 1989 Private School Survey, 58 percent of private school enrollment in grades 9 through 12 was in Catholic schools, 27 percent in other religious schools, and the remaining 15 percent in non-sectarian schools in the United States in 1989. Moreover, Catholic schools captured 75 percent of the 1989 total private school enrollment. Neal (2002) reviews the literature that documents the positive effects of Catholic school attendance.

[7] The only other paper that simulates large-scale voucher experiments based on econometric estimates is Altonji, Huang and Taber (2004), which focuses on the impacts of vouchers on the peer group of the students who remain in public schools, holding locational decisions and the political economy equilibrium constant.

Although models of household sorting across jurisdictions originated with Tiebout's (1956) work,[8] only recently have researchers estimated them. Since household residential choices interact with housing prices, community compositions, and the level of local public goods such as education, researchers have developed estimation approaches consistent with the structure of the equilibrium. As the complexity of these models often precludes closed-form solutions, the challenge arises of ensuring that in the estimation *all* the conditions that characterize an equilibrium hold, as discussed below.

The inclusion of household idiosyncratic preferences in my model resembles the use of random utility models in demand models for differentiated products used in industrial organization (Berry (1994), Berry, Levinsohn and Pakes (1995)). Applying this framework, Bayer, McMillan and Rueben (2004) have estimated an equilibrium model of household sorting using restricted-access Census micro data for the San Francisco Bay Area. Although this model features a rich and flexible demand side for housing and location specific characteristics, it does not endogenize the provision of local public goods. Hence, this model does not account for the variation in spending, and possibly quality, across public schools in a given metropolitan area, which in turn relates to which households would be more likely to take up vouchers, and how vouchers would affect local public school quality. In addition, this model does not explore the role of privately provided alternatives to local public goods.

Furthermore, in this framework the demand for each jurisdiction is the aggregation of individual demands emerging from the random utility model and is thus a function of the mean utility level across households. This, in turn, is found by equating the predicted demand for each jurisdiction with the observed population –namely, by "inverting" the jurisdiction's population share. Since mean utility is assumed a linear function of the jurisdiction's observed and unobserved characteristics for a given set of parameters, and the observed characteristics – such as housing prices – are presumably correlated with the unobserved ones, the parameters of the model are estimated through an instrumental variable regression of the mean utility level on the observed characteristics. Hence, estimation relies on a two-step procedure

[8] See Epple and Nechyba (2004) for a recent survey of the literature on Tiebout models.

that first solves for the partial equilibrium of the housing market given other endogenous variables and then addresses this endogeneity through instrumental variables techniques.

This paper, in contrast, features a one-step, full-solution estimation method that solves for the general equilibrium of the model as part of the estimation procedure. This approach is particularly well suited to ensure that all equilibrium conditions hold at once because the very computation of an equilibrium is the search for an allocation that fulfills all those conditions. While clearly desirable, full-solution estimation is computationally more costly than a two-step procedure. Therefore, I have developed fast algorithms to compute the equilibrium that make my estimation approach computationally feasible. Furthermore, this equilibrium computation is the same one used for policy simulations, which yields internal consistency and makes policy outcomes completely transparent. Besides being the first attempt to estimate a multi-jurisdictional model by full-solution methods, this paper illustrates how a similar procedure might apply to estimate other types of equilibrium models.[9] Thus, this research lies at the frontier of computational analysis and estimation.

The first papers to estimate an equilibrium model of household sorting among local jurisdictions while accounting for the endogeneity of local tax and expenditure policies were Epple and Sieg (1999) and Epple, Romer and Sieg (2001). This model, however, does not include private schools. Furthermore, its local public good index aggregates elements with potentially dissimilar roles in a voucher environment.[10] While household preference heterogeneity plays a key role in household sorting across jurisdictions in this model, housing quality variation within and across districts is essential to sorting in my framework. The identification of Epple and Sieg's model, estimated through a two-step procedure by exploiting necessary conditions for the equilibrium, relies on variation across districts within a metropolitan area, whereas identification in my framework also relies on variation across metropolitan

[9] Full-solution estimation has also been developed by Calabrese, Epple, Romer and Sieg (2004) in work initiated after mine, for a model that extends Epple, Romer and Sieg (2001).

[10] For instance, even if all households have the same preferences for housing and school quality, some of them might be willing to live in districts with low public school quality yet relatively good housing and low property-tax inclusive housing prices, for the sake of sending their children to private schools.

areas. Moreover, policy simulations carried out with a model related to Epple and Sieg's have focused on exogenous public good changes (Sieg, Smith, Banzhaf and Walsh (2004)). In contrast, my paper recognizes that the provision of local public goods would adjust endogenously under vouchers both through household individual choices and voters' collective decisions. Thus, my model is uniquely suited to answer the policy questions of interest.

To estimate the model I match key features of the predicted equilibrium through a minimum distance estimator using 1990 data on a cross-section of metropolitan areas and school districts. The estimates successfully capture the pattern of income stratification and the distribution of housing values across districts within metropolitan areas, and replicate public school spending and private school enrollment rates reasonably well. Using the parameter estimates, I assessed the effects of two hypothetical policies for the Chicago metropolitan area: the introduction of universal vouchers and of vouchers restricted to non-sectarian private schools ("non-sectarian vouchers")[11]. Universal voucher analysis provides insight into the impact of an unrestricted voucher program, of which any other voucher program may be seen as a special case. In these simulations, a voucher is a set amount of money received by the household from the state for the exclusive purpose of paying private school tuition.

According to my simulation results, both programs increase private school enrollment and affect household residential choice. For instance, some voucher users migrate towards neighborhoods with lower tax-inclusive housing prices and send their children to private schools, thus weakening the residential stratification of the current public school system. In the two programs, most households gain school quality for vouchers of at least $3,000. While households with an approximate wealth of $35,000 experience the largest school quality gains in both programs, low-income households reap the largest welfare gains. Despite these similarities, universal and non-sectarian vouchers differ in important ways. Universal vouchers increase enrollment at both Catholic and non-Catholic private schools, yet when

[11] For simplicity, this paper considers only one type of religious school, namely Catholic schools. Thus, the non-sectarian voucher simulation is the simulation of vouchers restricted to non-Catholic private schools, although it is meant to capture the effects of restricting vouchers to non-sectarian schools.

vouchers are restricted to non-sectarian schools, overall private school enrollment expands less and Catholic school enrollment declines as the voucher rises. Further, fewer households benefit from non-sectarian vouchers, particularly in the low-income segment. Whereas households who prefer Catholic schools benefit the most from universal vouchers, they lose the most from non-sectarian vouchers.

The remainder of this paper is organized as follows: Section 2 provides descriptive statistics of the data employed; Section 3 presents the model; Section 4 discusses the computational version of the model used for estimation purposes; Section 5 describes the estimation procedure; Section 6 discusses the estimation results; Section 7 analyzes voucher effects in policy simulations, and Section 8 concludes.

2. Descriptive Statistics

My analysis focuses on the metropolitan areas of New York, Chicago, Philadelphia, Detroit, Boston, St. Louis, and Pittsburgh, and the secondary and unified school districts therein. As of 1990, these were among the twenty largest metropolitan areas in the United States. They also depended highly on local sources for public school funding and had populations that were at least 25 percent Catholic (see Table 1).

As Table 2 shows, the school districts in these metropolitan areas vary widely along the dimensions of interest, such as private school enrollment, average household income and rental value, and public school spending per student. Moreover, households with children in private schools tend to have higher incomes while living in higher-rental value houses than households with children in public schools. The central city district is the largest district in each of my sample's metropolitan areas and captures most of the private school enrollment and non-residential property.

In addition, the geographic variation in private school markets across metropolitan areas seems to be shaped, at least partly, by the geographic variation in the distribution of adherents to different religions. Among the twenty largest metropolitan areas in the United States, those with higher private school enrollment rates have higher Catholic school enrollment rates and proportionally more Catholics. Moreover, the correlation between the fraction of students enrolled in Catholic schools and the fraction of Catholics equals 0.79, which squares with the fact that in 1990 about 85 percent of Catholic high school students in the United States were Catholic (National Catholic Educational Association (1990b)).

3. The Model

In the model, an economy is a set of public school districts with fixed boundaries that contain neighborhoods of different qualities. There are three types of schools: public, private Catholic, and private non-Catholic. Households that differ in endowment, religious preferences and idiosyncratic tastes for locations and school types maximize utility by choosing a location and a school for their children and by voting for property tax rates used to fund public schools. In equilibrium, no household wishes to move, switch to a different school, or vote differently.

Households and Districts

The economy is populated by a continuum of households, each one endowed with one house. The set of houses in the economy is partitioned into school districts. Every district d is in turn partitioned into neighborhoods, and there are H neighborhoods in total in the economy. Houses may differ across neighborhoods, but within a given neighborhood are homogenous and have the same housing quality and rental price. The size of the housing stock equals the measure of endowed houses and the housing stock cannot be varied in quantity or quality. Furthermore, each household has one child, who must attend a school, either public or private. One public school exists in each district[12] and the child may attend only the public school of the district where the household resides. If parents choose to send their child to a private school, Catholic or non-Catholic, they are not bound to any rule linking residence and school.

In addition to a house, households are endowed with a certain amount of income and there are I income levels. Besides endowment, households differ in their religious orientation, which is given by their valuation of Catholic schools relative to non-Catholic schools. Thus, a household has one of K possible religious types, where types $1,\ldots,L$ are Catholic, and the others non-Catholic. Not all the L

[12] This assumption rules out the existence of neighborhood schools, such as those in Epple and Romano (2003b).

Catholic types are necessarily identical, for they may differ in their relative valuation of Catholic schools, and the same is true for the non-Catholic types. Finally, households also differ in their idiosyncratic preference for each location and type of school.

Household preferences are described by the following Cobb-Douglas utility function:

$$U(\kappa, s, c, \varepsilon) = s^{\alpha} c^{\beta} \kappa^{1-\beta-\alpha} e^{\varepsilon}, \quad \kappa = k_{dh} \tag{1}$$

where $\alpha, \beta \in (0,1)$, k_{dh} is an exogenous parameter representing the inherent quality of neighborhood *h* in district *d* (i.e., housing size and age, geographic amenities, etc.), *c* is household consumption, *s* is the parental valuation of the quality of the child's school, which depends on the household's religious preferences, and ε is the household's idiosyncratic preference for the location and school type attended by the child. Furthermore, ε is distributed according to a continuous distribution $G(\varepsilon)$, and is independently and identically distributed across locations and school types for a given household and across households.

Household *i* seeks to maximize utility (1) subject to the following budget constraint:

$$c + (1 + t_d) p_{dh} + T = (1 - t_y) y_n + p_n \tag{2}$$

where y_n is the household's income, t_y is a state tax rate, and p_n is the rental price of the household's endowment house. Given its per-period total income, represented by the right-hand side of (2), the household chooses to live in location (d, h) with housing price p_{dh} and local property tax rate t_d. It also chooses a school for its child with tuition *T*, and *T* = 0 for public schools. The remaining income is used for consumption *c*.

Production of school quality

All schools in the economy produce school quality $\tilde{s}$ according to the following production function:

$$\tilde{s} = q^{\rho} x^{1-\rho} \tag{3}$$

where $\rho \in [0,1]$, *q* stands for the school's average peer quality and *x* is spending per student at the school. Denote by *S* the set of households whose children attend the school and by $\bar{y}(S)$ the average income of these households. Then the school's average peer quality is defined as $q = \bar{y}(S)$.[13] If the school is private, the spending per student x equals tuition T and may be supplemented with non-tuition revenue, whereas it equals the spending per student in district d, x_d, if the school is public and run by the local government.

The parental valuation of school quality (see s in equation (1)), depends on each household's religious preferences. A household of religious type $k = 1, \ldots, K$ whose child attends a school with religious orientation $j = 1, 2$ (Catholic or non-Catholic respectively, with public schools being non-sectarian and therefore non-Catholic) and quality $\tilde{s}_j$ perceives the school's quality as follows:

$$s_{kj} = R_{kj}\tilde{s}_j \qquad (4)$$

where $R_{kj} > 0$ is a preference parameter.

Public Schools

The quality of the public school in district d is $\tilde{s}_d = q_d^{\rho} x_d^{1-\rho}$, where q_d is the average income of households in district d with children attending this school. The public spending in education is funded by local property taxes, possibly aided by the state. Thus, the spending per student in district d is given by $x_d = t_d\left(P_d + Q_d\right)/n_d + AID_d$, where n_d is the measure of households choosing public school in district

[13] In this specification, peer quality captures all parental influences not mediated through school budgets that are positively correlated with parental income, such as parental involvement and monitoring, which have been found to be positively associated with education and income (McMillan (2000)). For simplicity I have assumed perfect correlation between peer quality and household income, although Nechyba (2000, 2003) has explored imperfect correlation by adding student's ability to the peer quality measure and found his fundamental results unchanged. Epple and Romano (2003) model achievement as a function of own ability and school quality, which in turn depends on the student body's average ability. They point out that any household variable, such as parental involvement, that positively affects both the child's performance and her school conforms to their model. Through similar reasoning one concludes that if income measures ability, and ability affects achievement, none of this paper's findings are affected by the peer quality measure. See Epple and Romano (2003b) for references to the theoretical literature that considers peer effects in the production of education.

d, AID_d is the amount of state aid per student for district d, funded through state income tax, and P_d and Q_d are the values of residential and non-residential district property, respectively.[14]

Private Schools

Private schools are modeled as clubs formed by parents under an equal cost-sharing rule. Since the school production function in (3) creates incentives for a household to join a school with households of equal or higher endowment and the production of school quality features constant returns to scale, households of a given endowment may optimally segregate into a private school and reject lower endowment households. Therefore, a private school formed by households of income level y_n has peer quality $q = y_n$.

Households of a given endowment share costs equally at a private school. Thus, the tuition equals the households' optimal spending on education, holding their residential locations fixed. That is, after choosing a location (d, h) with quality k_{dh}, household n of religious type k with income y_n may choose to send its child to a private school with tuition T and religious orientation $j = 1, 2$ (Catholic or non-Catholic) that maximizes utility (1) subject to the budget constraint (2) and the perception of school quality $s = R_{kj} q^{\rho} x_j^{1-\rho}$, where $q = y_n$, and $x_j = T$. Notice that the optimal tuition T determined by solving this optimal choice problem does not depend upon R_{kj}. Furthermore, private schools may supplement their tuition revenue with other sources. In particular, a private school of religious orientation j may match its tuition at the rate z_j, so that $x_j = (1 + z_j)T$. Parents who decide to open a private school choose the school religious orientation (Catholic or non-Catholic) that yields the higher utility.

[14] For simplicity, I model non-residential property as owned by an absentee landlord who does not participate in the elections to set property tax rates. I further assume full capitalization of property taxes for non-residential property, so that the gross-of-tax rental price of non-residential property is fixed. Since the tax base includes the net-of-tax value of this property, my treatment captures the incentive faced by voters when taxing non-residential property.

Household Decision Problem

Households are utility-maximizing agents that choose locations (d,h) and schools simultaneously, while taking tax rates t_d, district public school qualities s_d, prices $p_{dh,}$, and the composition of the communities as given. Household n chooses among all locations (d,h) in the budget set determined by the constraint $(1+t_d)p_{dh} \le (1-t_y)y_n + p_n$. For each location the household compares its utility under public, Catholic, and non-Catholic private schools. Migrating among locations is costless in the model and the household may choose to live in a house other than its endowed house.

Absolute Majority Rule Voting

Households also vote on local property tax rates. At the polls, households vote for property taxes taking their location, their choice of public or private school, property values, and the choices of others as given when voting on local tax rates. Households that choose private schools vote for a tax rate of zero, whereas households that choose public schools vote for a nonnegative tax rate. Because voters choose the tax rate conditional on their school choice, taking everything else as given, their preferences over property tax rates are single peaked. Property tax rates are determined by majority voting as long as they at least support an exogenously specified spending floor $\bar{x}_d$; if they do not, the property tax rate is set to cover the spending floor, which reflects adequacy clauses in state constitutions that seek to guarantee the minimum spending required to provide adequate school quality.

The state cooperates in funding public education in district d by providing an exogenous aid amount per student AID_d. This aid, which operates as a flat grant, is in turn funded by a state income tax whose rate t_y is set to balance the state's budget constraint.

Equilibrium

An *equilibrium* in this model specifies a partition of the population into districts and neighborhoods, local property tax rates t_d, a state income tax t_y, house prices p_{dh}, and a partition of the population into subsets of households whose children attend each type of school, such that: (a) every house is occupied; (b) property tax rates t_d are consistent with majority voting by residents who choose public versus private school, taking their location, property values, and the choices of others as given when voting on local tax rates; (c) the budget balances for each district; (d) the state budget balances, and (e) at prices p_{dh}, households cannot gain utility by moving and/or changing schools.

Though the equilibrium is proved to exist with a finite number of household types (Nechyba (1999)), no proof has been developed for the case of an infinite number of household types. Nonetheless, I have established conditions sufficient for determining whether an allocation is an equilibrium along the lines of the previous paragraph, and have developed an algorithm applying these conditions to compute the equilibrium.[15]

4. The Computational Version of the Model

In the computational version of the model the concept of "an economy" corresponds to a metropolitan area, and households do not migrate across metropolitan areas. The estimation strategy involves computing the equilibrium for each metropolitan area at alternative parameter points to search for the point that minimizes a well-defined distance between the predicted equilibrium and the observed data. Since the equilibrium does not have an analytical solution, I solve for it through an iterative algorithm for

[15] With a finite number of household types, the allocation of households to locations and schools is unique if the variation in district average housing quality is sufficiently large (Nechyba 1999). This condition is likely to hold for an infinite number of household types as well. Hence, for the empirical model I constructed neighborhoods so as to maximize such variation. Simulations have shown that the equilibrium is robust to the selection of different initial prices and assignments of households to locations. Although for a given variation in housing quality multiple equilibria are more likely when households place a sufficiently high value on school quality (i.e., high α) or when peer quality is very important relative to spending (i.e., high ρ), estimates for this model do not satisfy these conditions. Finally, notice that the type of equilibrium that I compute, with higher income households living in higher quality districts, seems to have been an empirical regularity in U.S. metropolitan areas for a number of years, thus becoming a reasonable focal point for estimation and policy simulations.

a tractable representation of each metropolitan area. Thus, this section describes the setup of districts and neighborhoods in this representation, the construction of household types, the state financial regime applied in computing the equilibrium, and the algorithm employed.

Community Structure

I measure the actual size of neighborhoods, districts, and metropolitan areas by the number of housing units. For computational tractability I aggregate the actual districts of each metropolitan area into *pseudo-districts* in order to compute the equilibrium, such that the largest district is a pseudo-district in itself, while smaller, contiguous districts are pooled into larger units. The actual 671 districts thereby yield 58 pseudo-districts. Figure 1 depicts Census tracts, and school districts and pseudo-districts for the Chicago metropolitan area. Once the pseudo-districts (henceforth called districts) are constructed, I split them into neighborhoods of approximately the same size, such that some districts have only one neighborhood while others have several. Larger metropolitan areas have larger neighborhoods.

Neighborhood Quality Parameters

In the theoretical model each neighborhood is composed of a set of homogeneous houses, such that neighborhood h in district d has a neighborhood quality index equal to k_{dh}. Since standard datasets do not measure neighborhood quality, I construct an index that captures housing quality and neighborhood amenities, *excluding* public school quality. The Census geographical concept which best approximates a neighborhood is the Census tract. Hence, I first compute the neighborhood quality index for each Census tract by regressing the logarithm of tract average rental price on a set of neighborhood characteristics and school district fixed effects for each metropolitan area,[16] then making each tract’s neighborhood quality index equal to the tract's fitted rental value net of school district fixed effects. The motivation for this

[16] I use tract average housing characteristics from the Census, and a linear and quadratic term in tract distance to the metropolitan area center. See Ferreyra (2002) for more details on the computation of the neighborhood quality parameter and for the data sources used to compute rental values and neighborhood characteristics.

regression is that, broadly speaking, rental prices reflect housing characteristics, neighborhood amenities, and public school quality. Thus, the district fixed effect nets out the school quality component from the measure of neighborhood quality. After obtaining the neighborhood quality index for each Census tract, I construct neighborhoods of the desired size by pooling contiguous tracts whose value for the neighborhood quality index lies in the same range. Lastly, I assign each neighborhood the median quality index from the neighborhood's tracts.[17]

For an example of the final representation of a metropolitan area through pseudo-districts and neighborhoods, see Figure 2 for the Chicago metropolitan area. The central city of Chicago overlaps entirely with the central district. Unlike the suburban districts, which have one neighborhood each, the central district has seven neighborhoods which differ in housing quality. On average, the central district has the lowest housing quality in the metropolitan area, although some neighborhoods in the central district are of higher housing quality than others in the suburbs.

Households

I consider income levels equal to the 10th, 30th, 50th, 70th and 90th percentiles of the income distribution of households with children in public or private schools in grades 9 through 12 in each metropolitan area. For computational purposes, the joint distribution of housing and income endowment is as follows. At the beginning of the equilibrium computation, the distribution of income in each neighborhood is initially the same and equal to the metropolitan area's. Hence, income and housing endowments are independently

[17] The process of constructing pseudo-districts and neighborhoods involves the following steps: a) based on the total size of the metropolitan area, determine the number of equal-sized neighborhoods that yield tractable computations; b) find the neighborhood quality parameter for each Census tract; and c) pool contiguous tracts with similar values for their neighborhood quality parameters into one neighborhood, such that no actual district is split between neighborhoods, each neighborhood comes as close to the size determined in (a) as possible, and the central district remains a pseudo-district in itself. For estimation-related reasons, I organize the neighborhoods thus constructed into as many pseudo-districts as possible rather than having fewer pseudo-districts with many neighborhoods each.

distributed as are religious preferences and endowments.[18]

Recall that each household is characterized by two religious matches, one with respect to Catholic schools and another with respect to non-Catholic schools. If household *n* is Catholic, its religious preferences are described by its matches with respect to Catholic and non-Catholic schools, $R^n_{C,C}$ and $R^n_{C,NC}$ respectively; if household *n* is non-Catholic, its religious preferences are given by its matches with respect to Catholic and non-Catholic schools, $R^n_{NC,C}$ and $R^n_{NC,NC}$ respectively. Since there are two types of schools, I make $R^n_{C,NC} = R^n_{NC,NC} = 1$ and focus on the relative valuation of Catholic schools. Unlike income, whose distribution comes straight from the data, the distribution of religious matches $R^n_{C,C}$ and $R^n_{NC,C}$ needs to be estimated. Therefore, I construct a discrete distribution of religious matches by assuming an underlying continuous distribution: I assume that $R^n_{C,C}$ and $R^n_{NC,C}$ are distributed uniformly over the intervals $[(1+r)(1-\delta),(1+r)(1+\delta)]$ and $[(1-r)(1-\delta),(1-r)(1+\delta)]$ respectively, where $0 < r < 1$ and $0 < \delta < 1$. The parameter *r* is both the premium enjoyed by the average Catholic in a Catholic school, and the negative of the discount suffered by the average non-Catholic in a Catholic school, whereas the parameter δ is proportional to the coefficient of variation of these distributions.[19] Finally, since households in the model also differ in their idiosyncratic preferences for locations and

[18] For computational convenience I place a measure of households equal to ten in each (house endowment, income) combination. For instance, if the proportions of Catholics and non-Catholics in the metropolitan area are 28 and 72 percent respectively, then Catholic and non-Catholic households have initial measures equal to three and seven, respectively, in each (house endowment, income) combination, given the lack of empirical evidence against income and religion being independently distributed (Ferreyra (2002)). Data on the fraction of Catholics in a metropolitan area come from the 1990 Church and Church Membership in America survey.

[19] To exemplify the determination of household religious matches, assume that 28% of all households are Catholic, r=0.2 and δ=0.1, which implies that $R^n_{C,C}$ and $R^n_{NC,C}$ are uniformly distributed between 1.08 and 1.32, and between 0.71 and 0.88 respectively. Hence, the first, second, and third Catholic household types have matches equal to 1.08, 1.20 and 1.32, and the first through seventh non-Catholic household types have matches equal to 0.72, 0.75, 0.77, 0.80, 0.83, 0.85, and 0.88.

schools, I assume that ε follows a type I extreme value distribution with scale parameter 1/*b*, where *b*>0. Thus, $F(\varepsilon) = \exp(-\exp(-\varepsilon / b))$, and the variance of ε equals $(1/6)\pi^2 b^2$.

State Aid and Non-Residential Property

Since the metropolitan areas included in my analysis on average fund more than two thirds of public school spending through local sources, it is clear that the efforts from those states to equalize spending across districts are quite limited. Furthermore, these metropolitan areas differ in the actual (and extremely complex) state formulas for the allocation of funds among districts (Hoxby (2001)), and the allocation process involves subtleties often unobservable to the researcher (Nechyba (2003)). Hence, I simplify by using the same mechanism across metropolitan areas – a local funding system with a state flat grant per student which may differ across districts and is equal in value to the state aid reported by the 1990 School District Data Book.[20] As for the non-residential property, its gross-of-tax value for a district is such that in the absence of property taxes, the ratio between this value and the observed value of the residential property tax base equals the district observed ratio of assessed non-residential to residential property.[21]

The Algorithm

In the model, the parameter vector is $\theta = (\alpha, \beta, \rho, \mathrm{r}, \delta, b)$. Computing the equilibrium for each parameter point and metropolitan area is an iterative process in which households choose locations and schools and vote for property taxes until no household gains utility by choosing differently. The input for the algorithm consists of the community structure, initial distribution of household types, initial housing

[20] Whereas the actual matching grant mechanisms in these metropolitan areas create incentives for higher property tax rates, the assumed flat grants have the opposite effect. However, by creating an incentive for higher property tax rates, non-residential property mimics the effect of matching grants. Furthermore, districts that receive large matching grants tend to have large stocks of non-residential property as well.

[21] I constructed the observed ratio of assessed non-residential to residential property using data from the Departments of Revenue of Illinois, Pennsylvania and New York, the Massachussetts Taxpayers Foundation, the Citizens Research Council of Michigan, and the 1987 U.S. Census of Governments.

prices, state aid, non-residential property and spending floor for each district. The output is the computed equilibrium from which I extract the variables whose predicted and observed values I match in the estimation (see the Appendix for further details).

5. Estimation

I estimate the model using a minimum distance estimator. I match the observed and simulated values of the following district-level variables, which I construct based on the 1990 School District Data Book: y_1 = average household income, y_2 = average housing rental value, y_3 = average spending per student in public schools, and y_4 = fraction of households with children in public schools. In addition, I match y_5, the fraction of households with children in Catholic schools at the metropolitan area level, calculated from the 1989 Private School Survey.[22] These variables, which are scaled to have unit variance in the sample, are of interest because they provide the basic characterization of household sorting across districts and schools, and the resulting spending in public schools.

Let D denote the total number of districts in the sample (D=58), M the number of metropolitan areas (M=7), and N_j the number of observations available for variable y_j, $j=1, \ldots 5$, so that $N_1 = N_2 = N_3 = N_4 = D$, and $N_5 = M$. Assume that district i is located in metropolitan area m. Then, denote by X_i the set of exogenous variables for district i, such that $X_i = x_i \cup x_m \cup x_{-i}$. Here, x_i is district i's own exogenous data (state aid, non-residential property and spending floor, number of neighborhoods, neighborhood quality in each neighborhood), x_m is exogenous data pertaining to metropolitan area m (10th, 30th, 50th, 70th and 90th income percentiles, and the fraction of Catholic households in the metropolitan area), and x_{-i} is the "own" data from the other districts in metropolitan area m. In addition, the set of independent variables for y_{5m} is X_m, which is the union of all the X_i sets corresponding to the districts that

[22] The fraction of households who reside in a district and send their children to Catholic schools is not available, since no data source links households' residences with different types of private schools. However, it seems reasonable to assume that households with children enrolled in Catholic schools located in a given metropolitan area reside there, which allows me to match Catholic school enrollment at the metropolitan area level.

belong to metropolitan area *m*. Finally, let n_i denote the number of housing units sampled in district *i*, and let n_m denote the number of housing units sampled in metropolitan area *m*.

I assume the following:

$$E\left(y_{ji} \mid X_i\right) = h_j(X_i, \theta) \qquad j = 1,...4;\ \mathrm{i} = 1,...\mathrm{N_j} \tag{5}$$

$$E\left(y_{5m} \mid X_m\right) = h_5(X_m, \theta) \qquad \mathrm{m} = 1,...\mathrm{M} \tag{6}$$

where the *h*'s are implicit nonlinear functions that express the equilibrium value of each endogenous variable I match as a function of the exogenous data and the parameter vector θ. Since the y_{ji}'s are (district-level) sample means, $C(y_{ji}, y_{ki'} \mid X_i, X_{i'}) = \sigma_{jk} / n_i$ if $i = i'$ and 0 otherwise, with $V(y_{ji} \mid X_i) = \sigma_{jj} / n_i = \sigma_j^2 / n_i = \sigma_{ji}^2$, where σ_{jk} and σ_j^2 denote population covariances and variances, respectively. Similarly, given that the y_{5m}'s are also sample means, $C(y_{ji}, y_{5m} \mid X_i, X_m) = \sigma_{j5} / n_i$ if district *i* is located in metropolitan area *m*, and $C(y_{ji}, y_{5m} \mid X_i, X_m) = 0$ otherwise. Also, $V(y_{5m} \mid X_m) = \sigma_{55} / n_m = \sigma_5^2 / n_m = \sigma_{5m}^2$.

Estimation Strategy

Because the number of observations is rather small, I estimate the model using Feasible Weighted Least Squares to account for heteroskedasticity across observations and then use the cross-equation covariances to obtain correct standard errors. The first stage of Feasible Weighted Least Squares determines the value for θ that minimizes the following loss function:

$$L(\theta) = \sum_{j=1}^{4} \sum_{i=1}^{N_j} \left(y_{ij} - \hat{y}_{ij}(\theta)\right)^2 + \sum_{m=1}^{M} \left(y_{5m} - \hat{y}_{5m}(\theta)\right)^2 \tag{7}$$

and the residuals from this regression are used to compute $\hat{\sigma}_j^2$ and $\hat{\sigma}_5^2$. The second stage runs Nonlinear Least Squares on variables transformed to account for heteroskedasticity, and seeks to minimize the following loss function in the transformed variables:

$$\tilde{L}(\theta) = \sum_{j=1}^{4}\sum_{i=1}^{N_j}\left(y_{ji}^{*} - \hat{y}_{ji}^{*}(\theta)\right)^2 + \sum_{m=1}^{M}\left(y_{5m}^{*} - \hat{y}_{5m}^{*}(\theta)\right)^2 \qquad (8)$$

where * denotes division by $\hat{\sigma}_{ji}$ or $\hat{\sigma}_{5m}$. The value of θ that minimizes this function, $\hat{\theta}$, is my estimate for the parameter vector. In addition to the model in Section 3, I estimate three simplified models to highlight the empirical richness of my theoretical framework. In particular, Model 1 excludes household idiosyncratic preferences (i.e., b=0) while the others include them.

Computational Considerations

Since Model 1 has a finite number of household types, it exhibits a coarseness that poses challenges for the equilibrium computation, the estimation and the fit of the data. To estimate this model I use a refined grid search, which allows for the objective function to be evaluated at each parameter point independently of others and lends itself to the type of parallel computing that I exploit in the estimation. However, the disadvantage of a grid search is that the grid size grows exponentially with the number of parameters. Using Condor to estimate this model,[23] I evaluate the objective function at about 250 parameter points simultaneously using a separate processor for each point. A function evaluation takes approximately ten minutes on a 1 Ghz Intel processor, and the full two-stage procedure takes about a week.

In contrast, the presence of an infinite number of household types in Models 2, 3 and 4 facilitates the computation of the equilibrium and the estimation of the model, for which I employ a cyclical coordinate descent algorithm (Bertsekas (1995)). A function evaluation takes about forty seconds. Furthermore, the full estimation takes between one and two days in a 3 Ghz Intel processor and can be run on a desktop, which is a clear simplification over Model 1.[24]

---

[23] A project of the Computer Science Department at the University of Wisconsin-Madison, Condor is a software system harnessing the power of a cluster of UNIX workstations on a network (http://www.cs.wisc.edu/condor).

[24] In Model 1 the loss function is discontinuous because of the discreteness of household types and the presence of a median voter in each district. However, a sufficiently large number of household types and school districts would yield a smooth objective function. Although majority voting still generates some minor discontinuity in the objective

Identification

The model is identified if no two distinct parameter points generate the same equilibrium for each metropolitan area. A sufficient condition for local identification is that the matrix of first derivatives of the predicted variables with respect to the parameter vector has full column rank when evaluated at the true parameter points, a condition which requires sufficient variation in the exogenous variables across districts and metropolitan areas. Evaluated at my parameter estimates, the matrices of first derivatives of the estimated models have full column rank in my sample.

Although a change in one parameter produces changes in several endogenous variables given the nature of the model, one can still identify the first-order effects from varying each parameter. A higher coefficient on school quality in the utility function (α) implies higher educational spending and lower housing prices, and a higher coefficient on consumption (β) implies higher household consumption and lower housing prices. A higher elasticity of school quality with respect to peer quality (ρ) lowers the importance of spending in the production of school quality, hence lowering spending. Furthermore, a higher ρ makes households more willing to segregate themselves by forming private schools.

An increase in the Catholic school premium (r) raises the relative valuation of Catholic schools among Catholics yet lowers it among non-Catholics. If r=0, then Catholics have the same preferences as non-Catholics and households sort themselves across private school types randomly. The greater the value of r, the higher the fraction of Catholic school enrollment accounted for by Catholics. An increase in δ raises the variation around the mean religious match among Catholics and non-Catholics. As δ rises, more Catholics come to prefer non-Catholic over Catholic schools and the reverse happens to non-Catholics. In particular, the identification of r and δ is largely driven by the variation in religious affiliation and Catholic school enrollment across metropolitan areas. Finally, an increase in the value of the variance of

function for Models 2, 3 and 4, a sufficiently large number of districts would yield a completely smooth function. For the sake of computing standard errors, I proceed as if I had good approximations to the continuous functions and rely on numerical derivatives.

idiosyncratic preferences (*b*) strengthens the role of idiosyncratic preferences in household location and school choice. When *b*=0, households' choices are only determined by their wealth and their religious preferences; when *b* is sufficiently large, household choices are only determined by their idiosyncratic preferences, which results in a random sorting of households across locations and schools.

6. Estimation Results

Table 3 presents the parameter estimates for the estimated models, each of which is discussed in turn below. I present three simplified models before turning to the most general formulation, which corresponds to the model in Section 3. Model 1 excludes household idiosyncratic preferences for locations and schools, non-residential property and spending floor in public schools, and assumes a zero subsidy rate for Catholic schools. Catholic school enrollment in this model is mostly driven by the heterogeneity of preferences for Catholic schools parameterized through r and δ. In this model, when faced with the choice between schools that are identical in everything except religious orientation, at least some Catholics would choose non-Catholic schools if $\delta > r/(1+r)$, whereas at least some non-Catholics would choose Catholic schools if $\delta > r/(1-r)$. The parameter estimates for this model generate sufficient overlap that Catholics may enroll in non-Catholic schools and vice versa. These estimates also lead to the rejection of the hypothesis that Catholics' and non-Catholics' preferences for Catholic schools follow the same distribution,[25] hence predicting that most of the equilibrium Catholic school enrollment proceeds from Catholic households. For example, in the non-voucher equilibrium for Chicago, 84 percent of the Catholic school enrollment is accounted for by Catholic households, a prediction that squares very well with the observed religious composition of Catholic schools (see Section 2).

Although the parameter estimates are highly significant, largely as a result of fitting sample means from Census data based on thousands of observations, the fit of the data displays some

[25] Interestingly, when Model 1 is estimated without including household religious preferences or religious schools, the point estimate for ρ is higher as it captures all factors different from spending which lead to the formation of private schools, including preferences for religious education. See Ferreyra (2002).

shortcomings. First, only districts with variation in housing quality (i.e., with more than one neighborhood) exhibit any variation in predicted household income and school choices. Hence, this model is not capable of replicating the observed private school enrollment in suburban districts, which have only one neighborhood in my representation. Second, Model 1 faces problems at matching the observed public school spending. In particular, in the central districts the predicted average household income and property values are sufficiently low that it is optimal for residents to vote for a property tax rate of zero and have public schools funded exclusively through state aid.

In order to explore the role of idiosyncratic preferences, Model 2 builds on Model 1 by adding these preferences for location and school type. Thus, in Model 2 households sort themselves across locations and schools based not only on their income and religious preferences, but also on their idiosyncratic tastes. Hence, in equilibrium all districts attract households of varying incomes who make heterogeneous school choices. This, in turn, facilitates the fit of private school enrollment in suburban districts and of other variables as well. However, perhaps not surprisingly, the addition of a second type of household preference heterogeneity results in reduced precision for the estimates of parameters characterizing private school enrollment, ρ, r and δ. While Model 2 fits private school and Catholic school enrollment better than Model 1, it does so almost solely on the basis of idiosyncratic preferences. Furthermore, a consequence of r not being significantly different from zero is that Catholics attend Catholic schools at the same rate as non-Catholics. This result counters the empirical evidence that 85 percent of Catholic school students come from Catholic households (see Section 2) and reveals Model 2's failure to capture an essential feature of private school markets – namely, who chooses which type of private school. Hence, in Model 4 I exploit additional information that permits more precise estimation of the parameters associated with private school enrollment.

Model 3 generalizes Model 2 to better reflect the environment in which educational expenditures are determined. First, it incorporates the non-residential property tax base, which is particularly important in central cities. Second, it reflects the recognition that state constitutional requirements for provision of education place an effective minimum on expenditure per student. Since constitutions do not state a minimum in explicit dollar terms, I use the empirical approach of finding a minimum expenditure level based on fit to the data. This leads to a choice of a spending floor equal to 60 percent of a district's observed spending. Third, it incorporates tuition subsidies received by Catholic school students. In the central district, the best-fitting subsidy rate is 100 percent, though the results are not substantially affected by the choice of a lower subsidy rate.[26] The resulting Model 3 fits the data better than Model 2, particularly with regards to spending.

Turning to Model 4, I improve the precision of the estimates of the parameters characterizing preferences for Catholic schools by requiring that the weighted average percent of Catholic in Catholic schools be equal to the observed national average of 0.85 as data on the religious composition of Catholic schools are not available at the metropolitan area level. I do this through the following iterative procedure: (1) fixing *r*, find the remaining parameters that minimize the objective function; (2) fixing the remaining parameters, update *r* to replicate the observed national average of .85; repeat steps (1) and (2) until all the parameter values converge. Since Model 4 contains the most general formulation of the theoretical model, I use its parameter estimates hereafter.

Analyzing the fit of the individual variables

Figures 3a through 3e depict the predicted and observed values for each variable. Overall, the model fits the data reasonably well, particularly for the central districts and the largest metropolitan areas. Furthermore, rank-order correlation analyses reveal the model's ability to replicate the observed district rankings within metropolitan areas. This relatively good fit is an encouraging result given the parsimonious parameterization of the model, the aggregation into pseudo-districts and the coarse

[26] I explored a number of generalizations to address the misfit of spending besides the one featured here, such as district variation in the number of public school children per household based on Census data, and different levels for the spending floor. For the chosen floor, the predicted spending in the city of Chicago is approximately equal to eighty percent of the observed one. Although the national average subsidy rate in Catholic high schools is about 40 percent (National Catholic Educational Association (1990a)), subsidies are heavily based on need, leading to the expectation that a substantially higher one prevails in the city.

discretization of the distributions of income and religious preferences.

The presence of idiosyncratic preferences, non-residential property, spending floor, and Catholic school subsidies lead to reasonably good predictions for private school enrollment (Fig. 3a), although the truncation of my five-point income distribution at the 90th percentile of income prevents greater predicted enrollment rates in the wealthiest districts. The model fits district average household income and rental value well (Figs. 3b and 3c, respectively), thus replicating sorting patterns across jurisdictions. Nonetheless, the model tends to under predict rental value. While this might be partly due to limitations of the neighborhood quality parameters, which do not include a number of actual physical neighborhood amenities, it also points to the possibility that housing prices may reflect neighborhood demographic composition above and beyond public school peer quality. The generalizations made on the spending side have helped fit public school spending reasonably well (Fig. 3d). The fact that the observed spending in the central cities of Pittsburgh, Boston and St. Louis rank almost at the top of their respective spending distributions is not replicated by the model. It suggests that additional factors such as details of the state aid allocation and interjurisdictional productivity differences in public schools may be important. Finally, Catholic school enrollment is particularly well fitted for New York, Chicago and Detroit (Fig. 3e).

Table 4 shows the correlations between the matched variables, both for the observed and fitted values. The correlations for fitted values resemble the actual correlations reasonably well. Hence, while acknowledging the limitations in the fit of the data and considering them informative for future extensions, I view the evidence presented here as indicative that the model successfully captures the patterns observed in the data.

7. Simulating Private School Vouchers

I simulate two types of voucher programs for the Chicago metropolitan area. The first type is a universal voucher program in which every household is eligible for a voucher that may be used for any type of private school. The second program (“non-sectarian vouchers”) differs in that the voucher may only be used for non-Catholic schools. In either program, households may supplement the voucher with additional payments towards tuition but cannot retain the difference when the tuition is lower than the voucher level. Consequently, the tuition is never set below the voucher level.

Furthermore, the voucher level, v, is set exogenously by the state in these simulations. Since vouchers are funded through a state income tax, the state income tax has to fund both the flat grants for public school students and the vouchers for private school students. Moreover, during the policy simulations, household n's budget constraint differs from the one given in (2) as follows:

$$c + (1+t_d)p_{dh} + \max(T - v, 0) = (1 - t_y)y_n + p_n \tag{9}$$

Notice that when forming private schools, households choose their optimal tuition taking into consideration voucher availability and dollar amount. Other things equal, vouchers lead to a higher tuition and school quality level while reducing the share of tuition paid by parents. Since it is not clear how donors to Catholic schools might respond to vouchers, I assume that the *total* tuition subsidy for urban Catholic schools provided in the non-voucher equilibrium remains constant throughout the simulations, which means that the subsidy per child may rise or decline depending on enrollment. In addition, an important issue concerns the benchmark equilibrium self-selection of households into locations and schools. For instance, a household with a high idiosyncratic preference for Catholic urban schools who chooses to live in the city and send their children to Catholic schools may respond differently to vouchers than an otherwise identical household with different idiosyncratic preferences who chooses suburban public schools. The Appendix provides details on the treatment of self-selection.

Before discussing the outcomes of the simulations, the benchmark equilibrium, which is the equilibrium simulated using the parameter estimates for a non-voucher regime, must be discussed. The first column of Tables 5a and 5b report the data, whereas the second column presents the benchmark equilibrium. In addition, Figure 4 depicts the geographic distribution of income, rental value, private school enrollment and public school spending in the benchmark equilibrium, which reasonably mirror the data. In particular, the simulated benchmark equilibrium correctly predicts that urban public schools have the lowest spending and peer quality in the metropolitan area. Furthermore, the benchmark equilibrium captures the fact that private school enrollment rates are higher in the city than in the suburbs. It also predicts that urban private school attendees reside in the central district's best neighborhoods, whose low tax-inclusive housing prices reflect the district's low public school quality. On average, private school households are wealthier and have a stronger preference for Catholic schools than public school households. Most private school attendees are enrolled in Catholic schools, whose students are primarily Catholic, whereas public and private non-Catholic school students are mainly non-Catholic.

Universal Vouchers

Below I analyze the effects of universal vouchers on school choices, residential decisions, and school quality in both public and private schools. In addition, I discuss the welfare implications of universal voucher policies and provide some perspective on my findings.

Household Sorting across Schools and Jurisdictions under Universal Vouchers

Table 5a presents some results from the simulation of universal vouchers for $1,000, $3,000, $5,000 and $7,000.[27] Private school enrollment grows with the voucher amount, and reaches a maximum of 0.74 of the entire population for a $7,000-voucher. Voucher availability gives rise to both new Catholic and private non-Catholic schools yet the private school market share for Catholic schools decreases with the voucher level. Furthermore, universal vouchers enable more households to attend the type of school that best suits their preferences. Thus, Catholic schools attract an increasing number of non-Catholics with a preference for Catholic schools, while non-Catholic schools capture fewer Catholics.

Figure 5 depicts household sorting across public, Catholic and private non-Catholic schools for selected universal and non-sectarian voucher levels. Since the majority of households enjoy a spending per student above $1,000 in the absence of vouchers, only households who can supplement a low voucher

[27] Notice that voucher amounts are up to $4,300, $2,700, $5,900 and $7,500 for the Florida, Cleveland, Milwaukee and D.C. programs respectively. Meanwhile, per pupil spending in those places is approximately equal to $7,500, $11,000, $11,000, and $12,000, respectively. See www.ij.org and www.heritage.org.

take it, and most of them already attend private schools in the benchmark equilibrium. Higher voucher levels, however, appeal both to low-income households in urban public schools, and to middle- and high-income households in suburban public schools.

The majority of voucher users are urban residents before vouchers, yet between 10 and 15 percent are suburbanites who move into the city. These migrants are middle- and high-income households with children in suburban public schools who move to the central district because of its relatively low tax-inclusive property values and choose private schools. Hence, vouchers attenuate the residential stratification generated by the residence-based public school system as they break the bundling of residence and public schools. Since tax-inclusive property values continue to be relatively high in the suburbs given their spending floors, the effect persists as the voucher grows. Interestingly, voucher availability reduces the housing premium in the best school districts but raises it in the locations favored by voucher users, thus causing capital losses or gains, respectively, to those homeowners.[28]

Vouchers also induce migration from the city to the suburbs on the part of households who seek private schools and better housing after reaping capital gains in the city. Although a quarter of voucher users remain in the suburbs, an increasing fraction relocates across them. Furthermore, private schools progressively spread to the suburbs, although the presence of non-residential property and high spending floors keep property tax rates relatively high and maintain high-quality public schools. Not surprisingly, private schools appear last in the best school districts.

While vouchers bring higher income households to the city, suburban districts with good housing yet low public school quality also attract wealthier voucher users. In addition, the greater affordability of districts with the best public schools appeals to wealthy households who strongly prefer public schools. Furthermore, the largest capital gains accrue to homeowners in the best urban neighborhoods *and* in the suburban districts chosen by voucher users. Although suburban locations lose property value with low vouchers, some of them gain for sufficiently high vouchers when majorities choose private schools, thus

[28] See Brunner, Sonstelie and Thayer (2001) for empirical evidence that homeowners in good public school districts are less likely to vote in favor of universal vouchers.

lowering the property tax burden and further attracting households who bid up rental values.

Whether or not vouchers have the ability to improve school quality for the low-income segment of the population is an important policy issue. Most low-income households, which are taken to include those with an income less than or equal to $20,000, need a relatively high voucher amount to compensate for the good public school peers they would lose in a private school and the spending per student they might forego. Yet half of the low-income segment takes up a $3,000-voucher and vouchers above $1,000 are used at a higher rate among these households than in the rest of the population. Encouraged by tuition subsidies, they also attend Catholic schools at higher rates.

With the expansion in the use of vouchers, funded by the state, fiscal burdens progressively shift from district property taxes onto the state income tax. While the average property tax rate falls with the increasing number of voters who favor zero property taxes, the income tax rate rises as the net outcome of lower state aid expense for public schools and higher voucher expense for private schools.

Universal Voucher Effects on School Quality

By affecting household residential and school choices, vouchers affect the quality of public and private schools as the evolution of school quality indicators shows in Table 5b. Most households gain school quality for vouchers of at least $3,000. Average school quality declines slightly for low voucher amounts but rises for larger ones; under a $7000-voucher the average school quality is 11 percent higher than in the benchmark equilibrium, a gain of 2.1 percent of the average household endowment. At the same time, the variation in spending and school quality rises for vouchers up to $5,000 and then falls for larger vouchers as private schools converge to a tuition equal to the voucher. The variation in peer quality experienced by students, however, grows as the educational system becomes increasingly private.

Behind this aggregate pattern lie remarkable differences across public and private schools. The public schools that remain open are the best ones, which accounts for the rising public school indicators. Average private school indicators, on the other hand, first drop as households with an income lower than the original private school population take up the voucher and then rise as wealthier households do so.

Furthermore, vouchers affect public schools through multiple channels. For instance, property tax rates fall as a greater number of voters choose a zero-property tax rate, and the value of the residential property tax base falls in several districts with the decline in demand for their public schools yet rises in other locations favored by voucher users. Moreover, higher spending per student may result even in the presence of a lower property tax base because fewer children need support in public schools.

These various forces, in turn, play out differently across school districts. For urban public schools, spending initially falls with the substantially lower property tax rate induced by vouchers but then rises when the effect of the lower public school enrollment prevails and leads to a higher per-student property tax revenue, particularly from non-residential sources. In suburban schools, the declining public school enrollment raises per-student property tax revenue until private schools comprise the majority of district voters, at which point the effect of the lower tax rate prevails. Moreover, voucher impact on public school peer quality also varies widely across districts. For instance, districts with the lowest benchmark equilibrium peer quality further lose good peers, whereas the reverse takes place in districts with the highest peer quality, whose declining housing prices entice wealthier households.

To evaluate how school quality gains and losses are distributed in the population, Figure 6 depicts the average school quality for the 10^{th}, 30^{th}, 50^{th}, 70^{th} and 90^{th} percentile endowments in the benchmark and universal voucher equilibria. Not surprisingly, school quality is increasing in endowment in all scenarios, yet converges across endowments as the voucher grows to cover most of the tuition payment. Vouchers below $5,000 mainly favor –albeit modestly- middle and high-income households with the ability to supplement the voucher. In contrast, higher vouchers favor households below the 75^{th} percentile endowment though they slightly damage those above, who are affected by the growing income tax burden and often attend declining public schools. Yet the greatest gains for sufficiently large vouchers accrue to households with an approximate endowment of $35,000. These households attend urban public schools in the benchmark equilibrium and mix with lower-wealth peers but then access higher quality schools, either by moving to suburban districts with better public schools or by switching into private schools.

Low-income voucher users gain spending with sufficiently high vouchers yet always lose peer quality. In this segment, school quality losses peak for the $3,000-voucher, which causes the loss of a significant number of good peers to those who remain in public schools yet is not high enough to match the pre-voucher quality for those who switch into private schools. Hence, only a voucher larger than the benchmark equilibrium spending per student in urban schools ($4,300 in these simulations) yields school quality gains for *all* low-income voucher users. These gains are sizeable (12 and 44 percent on average for $5,000 and $7,000-vouchers, respectively), and are even larger for Catholic school students.

Welfare Implications of Universal Vouchers

Among the most relevant issues concerning vouchers is who wins or loses when they are introduced. As Table 5b shows, the majority of the population benefits, slightly, from $1,000-vouchers, whereas just below half gains from larger voucher amounts (compensating variation measures welfare gains). Furthermore, the average welfare gain reaches a maximum at $237 for $3,000-vouchers and a minimum at -$1,040 for $7,000-vouchers. These outcomes, in absolute value equal to 0.5 and 2 percent of the average household endowment, respectively, show that the average welfare gains are relatively small, although the distributional effects are large. For instance, while the average winner may reap gains of up to 3 percent, the average loser may suffer losses of up to 5 percent. In these simulations the average winner is less wealthy yet more strongly prefers Catholic education than the average loser. Wealthy households, who already enjoy high school quality before vouchers, tend to lose under high voucher levels *regardless of their school choices* due to their high income tax burden and capital losses. Moreover, households who remain in public schools make up the largest fraction of welfare-losing households for vouchers below $7,000 despite the school quality gains attained by many of them.

While more households reap school quality gains as the voucher grows, fewer experience welfare gains. Furthermore, winners at low voucher levels are *less* likely to gain school quality than losers, a fact which is reversed at high voucher levels**.** Although seemingly counterintuitive, these findings simply highlight the multiplicity of channels that give rise to welfare changes: low vouchers mostly lead to savings in school or tax-inclusive housing spending that allow for greater consumption, which in turn brings forth welfare gains, whereas high vouchers yield school quality gains as well.

The top row of Figure 7 provides additional insights into the distributional effects of universal vouchers. On average, welfare gains are decreasing in endowment, yet households with the strongest preference for Catholic schools experience the largest gains at each wealth level. In addition, the large fiscal cost of the $7,000-voucher is evidenced by the fact that all households gain less for $7,000 than for $5,000-vouchers. A salient outcome is that the greatest relative welfare gains accrue to low-income households, who reap average gains between 2 and 3 percent of their wealth. Furthermore, low-income households benefit from vouchers at a higher rate than the rest of the population, and virtually all voucher users in this segment experience welfare gains. With small vouchers, low-income households reap consumption gains from lower property taxes in the central district although their urban public schools lose quality; with large vouchers they experience both consumption and school quality gains.

Further Issues

While an important aspect of vouchers is their ability to expand household residential and school choices, thus unleashing a variety of equilibrium effects, one might worry that the absence of moving costs in the model over predicts relocations and their associated effects. To provide some perspective on this issue, I simulated universal and non-sectarian voucher programs without allowing household relocation.[29] Since most voucher users relocate under full mobility, private school formation is slower now and heavily concentrated in the central district regardless of the voucher amount. Due to the lower voucher user rate, urban public schools keep their good peers yet lose spending at higher rates by foregoing the increase in property values induced by immigrants under full mobility. Moreover, the reduction in property tax rates leads to lower public school spending yet is not large enough to yield welfare gains through greater

[29] Detailed results for these simulations are available from the author upon request.

consumption. Overall, lack of mobility, which may be viewed as a short-run constraint, leads to fewer households experiencing school quality and welfare gain.

One might also wonder how much of the voucher effects described for the full mobility case are associated with the greater residential and school choice set afforded by vouchers, and how much with the mere increase in public spending for education. To investigate the normative effects of vouchers per se, I simulated an increase in per pupil state aid of equal dollar amount for each district, whose equilibrium total state expense equals that under vouchers. The greater state aid crowds out property tax effort on the part of suburban districts, which end up with almost the same levels of spending and quality. However, it significantly increases spending in the central district, where property taxes grow in response to a wealthier electorate and a greater public school enrollment associated with the improved public schools. Nonetheless, the higher property tax burden overwhelms the school quality gains for urban households, in contrast with the welfare-enhancing reduction in property tax burden under vouchers. The increase in state aid, which effectively reduces the variation in spending across districts by benefiting the central district proportionally the most, motivates less relocation than the voucher as it eliminates much of the property tax incentive to migrate across locations. By favoring public schools, this policy also leads to fewer households attending their optimal type of school and to current and former private school attendees losing the most welfare of all households. The combination of higher property tax burden for urban households and fewer school choices leads to only 12 percent of the population benefiting from this policy, as opposed to 43 percent that gains from the $5,000-voucher.

Finally, I investigated the robustness of the Chicago findings by simulating both voucher programs for the New York metropolitan area. While the qualitative results are similar, the main differences are centered on the fact that New York City accounts for 70 percent of the metropolitan area whereas Chicago comprises about 50 percent. Since the central district is larger in New York, public school attendees come from more varied income levels. Hence, proportionally more households benefit from the opportunity to segregate into private schools and the largest school quality gains accrue to households at the median rather than the 30th percentile of the endowment distribution. The larger central district also creates more opportunities for the establishment of private schools in the city's best neighborhoods and also generates property tax savings to more households. The combination of greater opportunities for school quality improvement and property tax savings yields welfare gains to the majority of the population at every voucher amount in New York. This contrasts with Chicago and warns against generalized conclusions regarding the political support for vouchers.

Non-sectarian vouchers

Non-sectarian vouchers raise the price of Catholic schools relative to non-Catholic private schools, thus inducing some households, depending on their religious preferences and budget constraints, to substitute non-Catholic private schools for Catholic schools. Tables 5a and 5b compare the results of universal and non-sectarian vouchers. Non-sectarian vouchers induce *less* private school enrollment than universal vouchers precisely because many households that would use a universal voucher would choose Catholic schools. Furthermore, under non-sectarian vouchers fewer households *than in the benchmark equilibrium* choose their optimal school type. Whereas enrollment grows in all private schools under universal vouchers, it now rises at private non-Catholic schools but falls at Catholic schools, since only households with a high taste for Catholic schools and the ability to pay for them remain (see the third row of Figure 5). Nonetheless, the enrollment losses in Catholic schools are tempered by the existence of tuition subsidies in urban Catholic schools, and Catholic schools' share falls from 11 percent of the total enrollment in the benchmark equilibrium to 6 percent under the $7,000-voucher, with most of this decline occurring for high vouchers.

Since both universal and non-sectarian vouchers subsidize private school attendance, they produce some qualitatively similar effects. They induce residential changes and generate comparable effects on property values. In addition, they have a similar impact on school quality and public schools. Both programs progressively shift the fiscal burden towards the state income tax, thus redistributing income from the wealthy to the poor. The combination of a higher fiscal burden, capital losses, and relatively small gains in school quality makes the average loser wealthier than the average winner in both programs, and households who remain in public schools after vouchers comprise the largest fraction of losing households. Moreover, households at the 30th percentile wealth reap the largest school quality gains in the two programs while low-income households enjoy the largest welfare gains.

Despite these similarities, universal and non-sectarian vouchers differ in other important ways. Fewer households attain school quality gains under non-sectarian vouchers, and urban schools experience greater losses because public school enrollment does not fall enough to offset the declining property tax rate while property values do not grow enough to raise spending. More low-income households use a universal than a non-sectarian voucher given the financial incentive to attend Catholic schools. Furthermore, a greater number of low-income households gain school quality through universal vouchers – this is true for both voucher users and non-users. Hence, a non-sectarian program seeking to match the universal vouchers' success at increasing school quality for the low-income segment needs to provide a more generous voucher.

The two voucher programs also differ in private school location patterns. Since fewer households compete for urban housing under non-sectarian vouchers, a greater fraction of non-sectarian private schools locate in the city and no new Catholic school opens in the suburbs. At low voucher levels, fewer users originally come from the city, given that urban Catholic school attendees cannot use the voucher and low-income households cannot supplement it.

Moreover, universal and non-sectarian vouchers have different welfare implications. For instance, the average welfare gain is higher for universal vouchers below $7,000, and households benefit from $1,000- and $5,000-universal vouchers at a higher rate.[30] While at each endowment level Catholic households gain the most, they lose the most under non-Catholic vouchers of at least $5,000 (see the

---

[30] The average welfare gain and the percent of winners from a non-sectarian $7,000-voucher are higher than for a $7,000-universal voucher because the lower adoption rate of the non-sectarian voucher imposes a lower fiscal burden. Although more households gain from a non-sectarian than a universal $3000-voucher, the difference in welfare gains is negligible for the households who win with non-sectarian yet lose with universal vouchers.

bottom row of Figure 7). Although most households either gain or lose in both programs, about 15 percent of all households gain from high universal yet lose from non-sectarian vouchers. These households, who either attend Catholic schools in the benchmark equilibrium or would choose them through universal vouchers, now turn to public or private non-Catholic schools for sufficiently high vouchers. With welfare outcomes that differ by 5 percent of their wealth across voucher regimes, these households are, not surprisingly, those whose welfare is most affected by the choice of voucher program.[31]

8. Concluding Remarks

Few policies are as controversial in the United States as private school vouchers. Although no large-scale voucher program has been implemented to date, one can learn about their potential effects through policy simulation within a general equilibrium framework. Thus, in this paper I estimate a general equilibrium model with multiple public school districts and private schools and use the parameter estimates to simulate voucher programs. An important contribution in this paper is the inclusion of religious schools and household religious and idiosyncratic preferences, which has enabled me to compare the effects of universal vouchers with vouchers restricted to non-sectarian schools in the Chicago metropolitan area. Whereas the two programs give rise to some similar effects, the evolution of the private school market differs in each case. In addition, fewer people, particularly from the lower-income segment of the population, benefit from non-sectarian vouchers. While those with the strongest preference for Catholic education gain the most under universal vouchers, they lose the most under non-sectarian vouchers.

The fact that households who care the most about religious education are the ones who lose the most in a non-sectarian program may seem an obvious result, and a skeptical reader may question its

[31] The Catholic Church has publicly stated its full support of unrestricted parental choice programs and specifically advocates against the exclusion of religious schools from voucher programs. See, for instance, http://www.usccb.org/bishops, http://www.ncea.org/publicpolicy/policystatements, and http://www.flacathconf.org.

usefulness. However, it is important to bear in mind the tradition in the United States federal jurisprudence that upholds parents' right to choose the type of education they want for their children – including, of course, religious education (see Viteritti (1999) and the references therein). Moreover, the US Supreme Court upheld the Cleveland voucher program as "entirely neutral with respect to religion" and as a "program of true private choice" (Zelman v Simmons-Harris (2002), p. 2473). An outcome of this decision was the enactment of the first federally funded voucher program in the District of Columbia, which started in the 2004/05 academic year, and the proposed Choice Incentive Fund for the 2005/06 fiscal year to enable other cities to develop parental choice programs. At the state level, however, the role of Blaine amendments in court battles continues to raise questions on the consequences of excluding religious schools from voucher programs. Besides providing tools for answering these questions, this paper's framework is appropriate to analyze other relevant issues such as voucher targeting and child-centered funding, the expansion of current voucher programs, and the potential effect of voucher proposals originating in the last five years but not yet implemented.[32]

While powerful, this framework certainly leaves room for important extensions such as the refinement of neighborhood quality measures, the inclusion of preferences for neighborhood demographic composition and of neighborhood schools, the distinction between renters and owners, and the introduction of households without school-age children. As more data become available on private schools, exploring private schools' pricing schemes and their productivity relative to the public sector should also prove worthwhile. Although no simulation exercise will be able to replace the actual enactment of a large-scale voucher program, developing and estimating general equilibrium models of local jurisdictions that incorporate private school markets can still shed much light on the potential outcomes of school choice programs.

---

[32] Between 2000 and 2005, thirty-nine states including California, New York and Texas have considered enacting voucher programs, and the legislatures of Wisconsin and Ohio are currently studying expansions to their programs in Milwaukee and Cleveland, respectively. While this paper analyzes vouchers for K-12 education, vouchers are also being used for pre-kindergarten in Louisiana and Florida. See www.heritage.org.

References

Altonji, Joseph; Huang, Ching-I and Taber, Christopher. "Estimating the Cream Skimming Effect of Private School Vouchers on Public School Students". Working Paper, Northwestern University, October 2004.

Bayer, Patrick; McMillan, Robert and Rueben, Kim. "A Equilibrium Model of Sorting in an Urban Housing Market". Working Paper, Yale University, June 2004.

Bearse, Peter; Glomm, Gerhard and Ravikumar, B. "Majority Voting and Means-Tested Vouchers". Working Paper, Michigan State University, 2001.

Berry, Steven. "Estimating Discrete-Choice Models of Product Differentiation". *RAND Journal of Economics*, Summer 1994, 25(2), pp. 242-262.

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

Bertsekas, Dimitri. *Nonlinear Programming.* Athena Scientific, Belmont, MA, 1995.

Bolick, Clint. *Voucher Wars: Waging the Legal Battle over School Choice*. Washington, DC: Cato Institute, 2003.

Brunner, Eric; Sonstelie, Jon and Thayer, Mark. "Capitalization and the Voucher: An Analysis of Precinct Returns from California's Proposition 174". *Journal of Urban Economics*, November 2001, 50(3), pp. 517-536.

Calabrese, Stephen; Epple, Dennis; Romer, Thomas and Sieg, Holger. "Local Public Good Provision, Myopic Voting, and Mobility". Working Paper, Carnegie Mellon University, 2004.

Caucutt, Elizabeth. "Educational Policy When There Are Peer Group Effects – Size Matters". *International Economic Review*, February 2002, v. 43, pp. 195-222.

Cohen-Zada, Danny and Justman, Moshe. "The Religious Factor in Private Education". *Journal of Urban Economics*, May 2005, 57(3), pp. 391-418.

Epple, Dennis and Nechyba, Thomas. "Fiscal Decentralization," in V. Henderson and J. F. Thisse, eds., *Handbook of Regional and Urban Economics*, vol. 4. North Holland: Elsevier, 2004, pp. 2423-24.

Epple, Dennis and Romano, Richard. "Competition Between Private and Public Schools, Vouchers, and Peer-Group Effects". *American Economic Review*, March 1998, 88(1), pp. 33-62.

Epple, Dennis and Romano, Richard. "Educational Vouchers and Cream Skimming," Working Paper, University of Florida, 2003a.

Epple, Dennis and Romano, Richard. "Neighborhood Schools, Choice, and the Distribution of Educational Benefits", in Caroline Hoxby, ed., *The Economics of School Choice*. Chicago, IL: U. of Chicago Press, 2003b.

Epple, Dennis; Romer, Thomas and Sieg, Holger. "Interjurisdictional Sorting and Majority Rule: An Empirical Analysis". *Econometrica*, November 2001, 69(6), pp. 1437-1465.

Epple, Dennis and Sieg, Holger. "Estimating Equilibrium Models of Local Jurisdictions". *Journal of Political Economy*, August 1999, 107(4), pp. 645-681.

Fernandez, Raquel and Rogerson, Richard. "School Vouchers as a Redistributive Device: An Analysis of Three

Alternative Systems", in Caroline Hoxby, ed., *The Economics of School Choice*. Chicago, IL: University of Chicago Press, 2003.

Ferreyra, Maria. "Estimating the Effects of Private School Vouchers in Multi-District Economies". PhD Dissertation, University of Wisconsin-Madison, 2002.

Hoxby, Caroline. "All School Finance Equalizations are Not Created Equal". *Quarterly Journal of Economics*, November 2001, 116(4), pp. 1189-1231.

Manski, Charles. "Educational Choice (Vouchers) and Social Mobility". *Economics of Education Review*, December 1992, 11 (4), pp. 351-69.

McMillan, Robert. "Parental Involvement and Competition: An Empirical Analysis of Public School Quality". Working paper, University of Toronto, 2000.

National Catholic Educational Association. *Catholic High Schools and Their Finances - 1990*, by Michael Guerra and Michael Donahue. Washington, DC, 1990a.

National Catholic Educational Association. *United States Catholic Elementary and Secondary Schools, 1989-90*, by Frederick Brigham, Jr. – Washington, DC, 1990b.

Neal, Derek. "How Vouchers Could Change the Market for Education". *Journal of Economic Perspectives*, Fall 2002, 16(4), pp. 25-44.

Nechyba, Thomas. "School Financed Induced Migration Patterns: The Case of Private School Vouchers". *Journal of Public Economic Theory*, January 1999, 1(1), pp. 5-50.

Nechyba, Thomas. "Mobility, Targeting and Private School Vouchers". *American Economic Review*, March 2000, 90(1), pp. 130-146.

Nechyba, Thomas. "Introducing School Choice into Multi-District Public School Systems", in Caroline Hoxby, ed., *The Economic Analysis of School Choice*. Chicago, IL: University of Chicago Press, 2003.

Sieg, Holger; Smith, Kerry; Banzhaf, Spencer and Walsh, Randall. "Estimating the General Equilibrium Benefits of Large Changes in Spatially Delineated Public Goods". *International Economic Review*, 2004, 45(4), pp. 1047-1077.

Tiebout, Charles (1956). "A Pure Theory of Local Expenditures". *Journal of Political Economy*, October 1956, vol. 64, 416-424.

US Department of Education. *Private School Universe Survey, 1989-90*. Washington DC: National Center for Education Statistics, 1992.

US Supreme Court of Justice. *Zelman v. Simmons-Harris*. 122 S.Ct. Washington, DC, 2002.

Viteritti, Joseph. Choosing Equality. School Choice, the Constitution, and Civil Society. Washington, DC: Brookings Institution, 1999.

TABLE 1
Selected Metropolitan Areas

Metropolitan Area	1990 Census Population (in thousands)	No. of School Districts	Largest District's Relative Size	Fraction of Catholic Population	Share of Local Sources
Boston, MA	2,871	87	.158	.49	.72
Chicago, IL	6,070	50	.472	.41	.68
Detroit, MI	4,382	110	.261	.35	.78
New York-Long Island, NY	11,156	167	.644	.43	.76
Philadelphia, PA—NJ	4,857	106	.315	.34	.57
Pittsburgh, PA	2,057	80	.141	.47	.63
St. Louis, MO-IL	2,444	71	.146	.26	.65

Secondary and Unified School Districts included. District Relative Size = number of housing units in district / number of housing units in metropolitan area. Share of local sources for public school funding is the district average share in each metropolitan area.
Source: 1990 Census and School District Data Book (SDDB), and 1990 Churches and Church Membership in America.

TABLE 2
School Districts in Selected Metropolitan Areas: Summary Statistics

	Mean	Std. Dev.	1^{st} Percentile	99^{th} Percentile
Fall Enrollment	1,988	10,741	115	10,698
No. Households	1,970	11,829	37	10,619
Fraction of hhs. w/ children in private schools	.112	.082	.000	.357
In Central District	.194	.051	.097	.268
Avg. Household Income ($)–All Households	63,589	27,985	24,066	169,457
Hhs. w/ Children in Public Schools	60,496	25,363	23,690	154,423
Hhs. w/ Children in Private Schools	84,736	54,219	19,664	297,177
Avg. Housing Rental Value ($)–All Households	13,525	8,007	3,499	40,349
Hhs. w/ Children in Public Schools	13,156	7,785	3,453	39,867
Hhs. w/ Children in Private Schools	16,159	9,775	2,905	46,656
Avg. Spending per Student in Public Schools ($)	7,674	3,987	3,221	22,500
District Size Relative to Metro Area	.010	.036	.000	.141
Share of Local (District) Revenues for Pub. Sch.	.697	.193	.186	.989
Share of State Revenues	.276	.177	.007	.697
Non-Residential Property Value / Resid. Prop.Value	.457	.518	.000	2.651
In Central District	1.181	0.535	.678	2.146

No. observations: 671 school districts - Household data and Fall enrollment are for grades 9 through 12.
District size = number of housing units in the district / number of housing units in the metropolitan area.
Source: 1990 SDDB and 1989 Common Core of Data. Non-residential property sources: see Section 4.

TABLE 3
Parameter Estimates

Parameter	Model 1 Estimates	Model 2 Estimates	Model 3 Estimates	Model 4 Estimates
α	0.12	0.086	0.068	0.077
	(0.001)	(0.002)	(0.001)	(0.003)
β	0.72	0.661	0.680	0.678
	(0.001)	(0.006)	(0.009)	(0.007)
ρ	0.24	0.010	0.121	0.221
	(0.001)	(0.070)	(0.029)	(0.008)
r	0.11	0.010	0.010	0.451
	(0.009)	(0.063)	(0.049)	
δ	0.25	0.270	0.713	0.275
	(0.001)	(0.723)	(0.278)	(0.037)
b		0.031	0.030	0.029
		(0.001)	(0.001)	(0.0003)
Sum of Squared Residuals	871.564	451.572	298.463	316.187
Unweighted Sum of Squared Residuals	714.460	255.581	225.953	244.215

Standard Errors in parentheses. Number of observations: see section 5.
Sum of Squared Residuals uses second-stage weights from Model 3. Unweighted Sum of Squared Residuals uses no weights. Ranking of models by sum of squared residuals is robust to the use of weights from any model.

TABLE 4
Goodness of Fit : Some Correlations

a. Observed Data

	Average Hh. Income	Average Rental Value	Spending per Student	Fraction Public
Average Hh. Income	1			
Average Rental Value	0.98	1		
Spending per Student	0.52	0.61	1	
Fraction Public	0.25	0.21	-0.13	1

b. Fitted Data

	Average Hh. Income	Avg. Rental Value	Spending per Student	Fraction Public
Average Hh. Income	1			
Average Rental Value	0.88	1		
Spending per Student	0.74	0.57	1	
Fraction Public	0.46	0.31	0.73	1

Number of observations: 58 districts. Weighted correlations - weight: district measure of households.

TABLE 5a

Universal and Non-Catholic Vouchers in Chicago: School Choice, Demographic, and Fiscal Effects

	Data	B.E. (1)	Universal Voucher Amount				Non-Catholic Voucher Amount			
			$1,000	$3,000	$5,000	$7,000	$1,000	$3,000	$5,000	$7,000
Private School Enrollment										
Fraction Households in Private Schools	0.16	0.16	0.22	0.43	0.60	0.74	0.19	0.29	0.48	0.62
Fraction Hhs. in Catholic Schools w.r.t. Private Schools	0.84	0.72	0.69	0.61	0.55	0.50	0.60	0.34	0.17	0.11
Fraction Private Schools Hhs. in Central District	0.61	0.72	0.72	0.75	0.64	0.55	0.71	0.71	0.68	0.60
School Choice before and after Vouchers										
Fraction Hhs. choosing Public-Public (2)			0.78	0.57	0.41	0.26	0.81	0.71	0.52	0.37
Fraction Hhs. choosing Public-Private			0.06	0.27	0.43	0.58	0.03	0.13	0.33	0.47
Fraction Hhs. choosing Private-Public			0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Fraction Hhs. choosing Private-Private			0.16	0.16	0.16	0.16	0.16	0.16	0.15	0.15
Fraction Hhs. choosing Catholic-Catholic			0.11	0.11	0.11	0 .11	0.11	0.09	0.08	0.06
Fraction Hhs. choosing Catholic-Non-Catholic			0.00	0.00	0.00	0.00	0.00	0.01	0.03	0.05
Fraction Hhs. in Optimal School Type (3)		0.68	0.70	0.75	0.76	0.77	0.68	0.67	0.66	0.65
Religious Composition of Public and Private Schools										
Catholic Schools: fraction of Catholic students		0.86	0.83	0.85	0.75	0.73	0.84	0.85	0.87	0.87
Private Non-Catholic Schools: fraction non-Catholic students		0.67	0.69	0.74	0.78	0.80	0.66	0.65	0.64	0.63
Public Schools: fraction of Catholic students		0.34	0.32	0.35	0.23	0.20	0.35	0.35	0.36	0.36
Demographics										
Average Household Income Ratio (4)	2.42	2.71	2.67	2.72	2.62	2.48	2.68	2.68	2.63	2.60
Average Housing Rental Value Ratio (5)	3.37	2.64	2.56	2.54	2.53	2.73	2.61	2.54	2.62	2.57
Fraction of Hhs. that move			0.07	0.27	0.43	0.56	0.04	0.15	0.35	0.50
Fiscal Effects										
Income Tax Rate		0.03	0.03	0.04	0.08	0.12	0.03	0.04	0.06	0.10
Avg. Property Tax Rate		0.24	0.23	0.20	0.15	0.08	0.23	0.23	0.17	0.13
Avg. Tax Burden (property tax + income tax)		$3,500	$3,500	$3,900	$4,800	$6,200	$3,400	$3,700	$4,300	$5,500

(1) "B.E." denotes "Benchmark Equilibrium"- (2) "Public-Public" is short for "public schools before vouchers, and public schools after vouchers" – (3) "Optimal School Type" is the type of school for which the household has the highest religious match - (4) Average Household Income Ratio = avg. hh. income in highest housing quality district / avg. hh. income in lowest housing quality district- (5) Average Housing Rental Value Ratio = id. Avg. Hh. Income Ratio, but for housing rental value.

TABLE 5b

Universal and Non-Catholic Vouchers in Chicago: Effects on School Quality and Household Welfare

			Universal Voucher Amt				Non-Catholic Voucher Amt			
All Schools	Data	B.E.	$1,000	$3,000	$5,000	$7,000	$1,000	$3,000	$5,000	$7,000
Avg. Quality		$10,200	$10,100	$10,000	$10,400	$11,300	$10,100	$11,000	$10,000	$10,200
Avg. Spending		$6,800	$6,700	$6,600	$7,100	$8,100	$6,700	$6,700	$6,700	$7,530
Avg. Peer Quality	$45,000	$45,000	$45,000	$45,000	$45,000	$45,000	$45,000	$45,000	$45,000	$45,000
Fraction of Hhs. w/higher School Quality			0.50	0.51	0.57	0.65	0.44	0.53	0.53	0.63
Fraction Low Income Hhs. w/higher Sch. Qual.			0.05	0.06	0.41	0.79	0.06	0.06	0.05	0.69
Public Schools										
Avg. Quality		$10,200	$10,400	$11,600	$12,300	$11,700	$10,100	$10,700	$11,400	$11,900
Central District		$6,100	$6,000	$5,000	$5,100	$5,700	$5,900	$6,100	$4,900	$5,000
Suburbs (average)		$13,000	$13,200	$13,800	$13,000	$10,800	$13,100	$13,400	$13,300	$12,300
Avg. Spending	$6,700	$6,800	$7,000	$7,700	$8,300	$7,900	$6,700	$7,200	$7,600	$8,000
Central District	$5,500	$4,300	$4,200	$3,300	$3,500	$4,000	$4,100	$4,300	$3,300	$3,400
Suburbs (average)	$7,800	$8,500	$8,700	$9,200	$8,600	$7,000	$8,600	$8,900	$8,800	$8,100
Avg. Peer Quality	$49,200	$43,000	$43,800	$48,800	$50,900	$49,100	$43,100	$44,500	$48,900	$49,500
Central District	$26,600	$21,400	$21,200	$20,400	$20,400	$19,900	$21,000	$21,000	$20,100	$19,200
Suburbs (average)	$63,900	$58,500	$58,300	$58,300	$57,100	$55,400	$58,400	$58,300	$57,500	$56,400
Private Schools										
Avg. Quality		$10,700	$9,200	$7,800	$9,100	$11,200	$9,900	$8,700	$8,600	$10,300
Avg. Spending		$6,800	$5,800	$5,100	$6,400	$8,200	$6,290	$5,500	$5,800	$7,300
Avg. Peer Quality	$67,700	$55,800	$49,500	$40,100	$41,000	$43,700	$53,600	$46,600	$40,900	$42,300
Welfare Implications										
Fraction of Hhs that Win with Vouchers			0.71	0.47	0.46	0.42	0.69	0.60	0.35	0.33
Fraction of Low Inc. Hhs. that Win w/Vouchers			1.00	0.81	0.69	0.65	0.97	1.00	0.59	0.56
Avg. Welfare Change			$170	$237	-$106	-$1,040	$97	$215	-$135	-$978
Avg. Welfare Change / Avg. Hh. Wealth			0.00	0.00	0.00	-0.02	0.00	0.00	0.00	-0.02
Winners										
Avg. Wealth			$47,700	$42,700	$43,900	$42,700	$49,100	$44,900	$41,900	$42,400
Avg. Taste for Catholic Schools			0.95	1.05	1.05	1.06	0.93	0.94	0.88	0.88
Losers										
Avg. Wealth			$65,700	$61,900	$60,500	$60,300	$61,500	$65,000	$58,700	$58,000
Avg. Taste for Catholic Schools			0.82	0.79	0.79	0.80	0.86	0.87	0.93	0.93

FIGURE 1
Chicago: Census Tracts, School Districts, and School Quasi-Districts

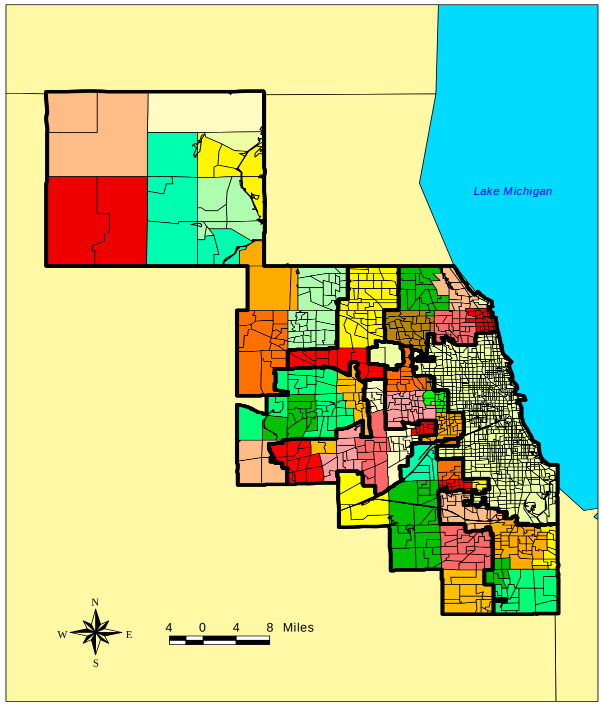

Note: The fine lines are the boundaries of Census tracts, and different shades identify Census tracts located in different school districts. The thick black lines are the boundaries of the pseudo-districts.

FIGURE 2
Chicago: Housing Quality by Neighborhood

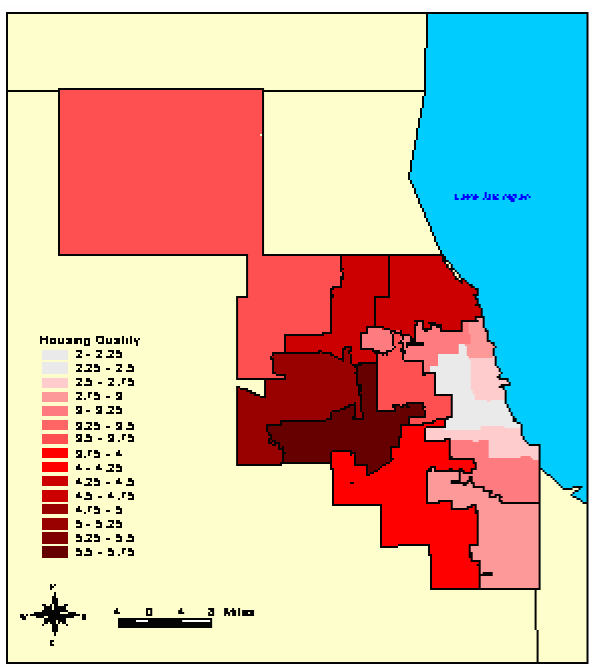

FIGURE 3 - *Fitted vs. Observed Values*

Figure 3a - Fraction of Households w/ children in Public Schools

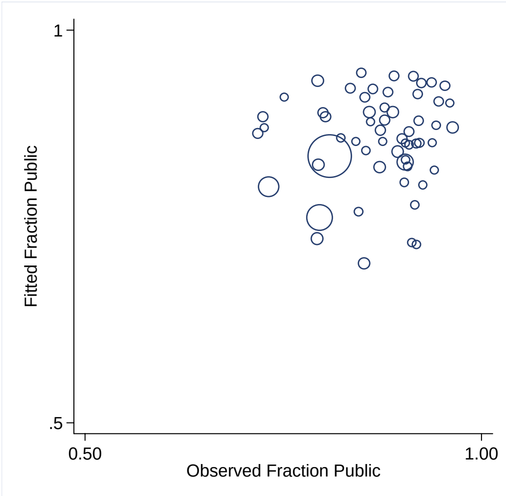

Figure 3b - Average Household Income (in $10,000)

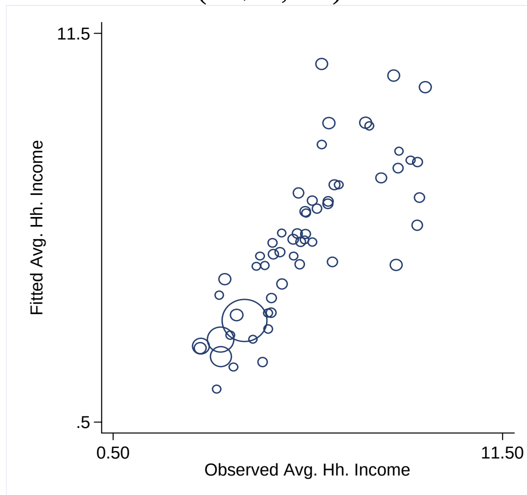

Figure 3c - Average Rental Value (in $10,000)

Figure 3d – Spending/Student in Public Schools (in $10,000)

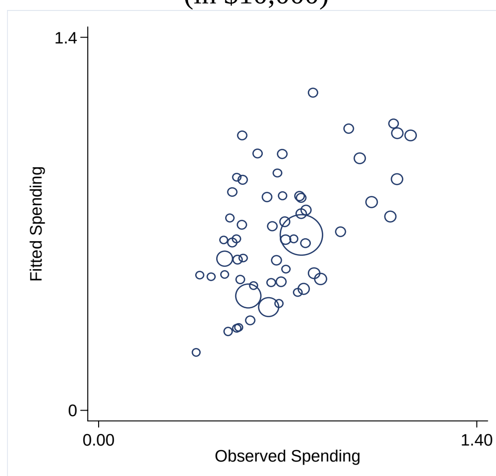

Figure 3e – Fraction of Households with Children in Catholic Schools

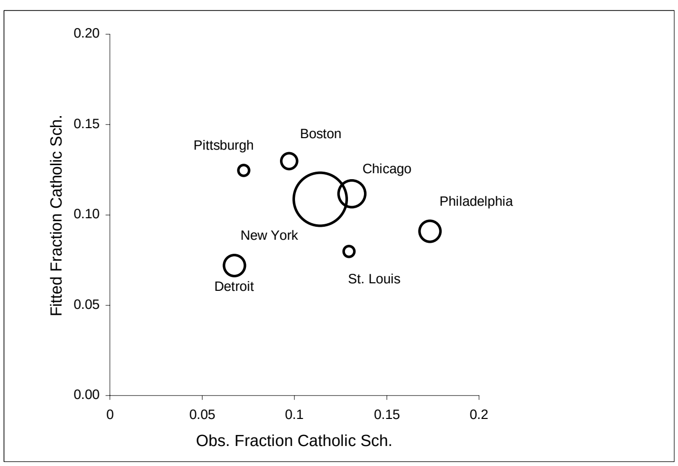

Note: observed values on the horizontal axis; fitted values on the vertical axis. Circle size is proportional to the observation's total measure of households.

FIGURE 4 - *Chicago: Non-Voucher Equilibrium*

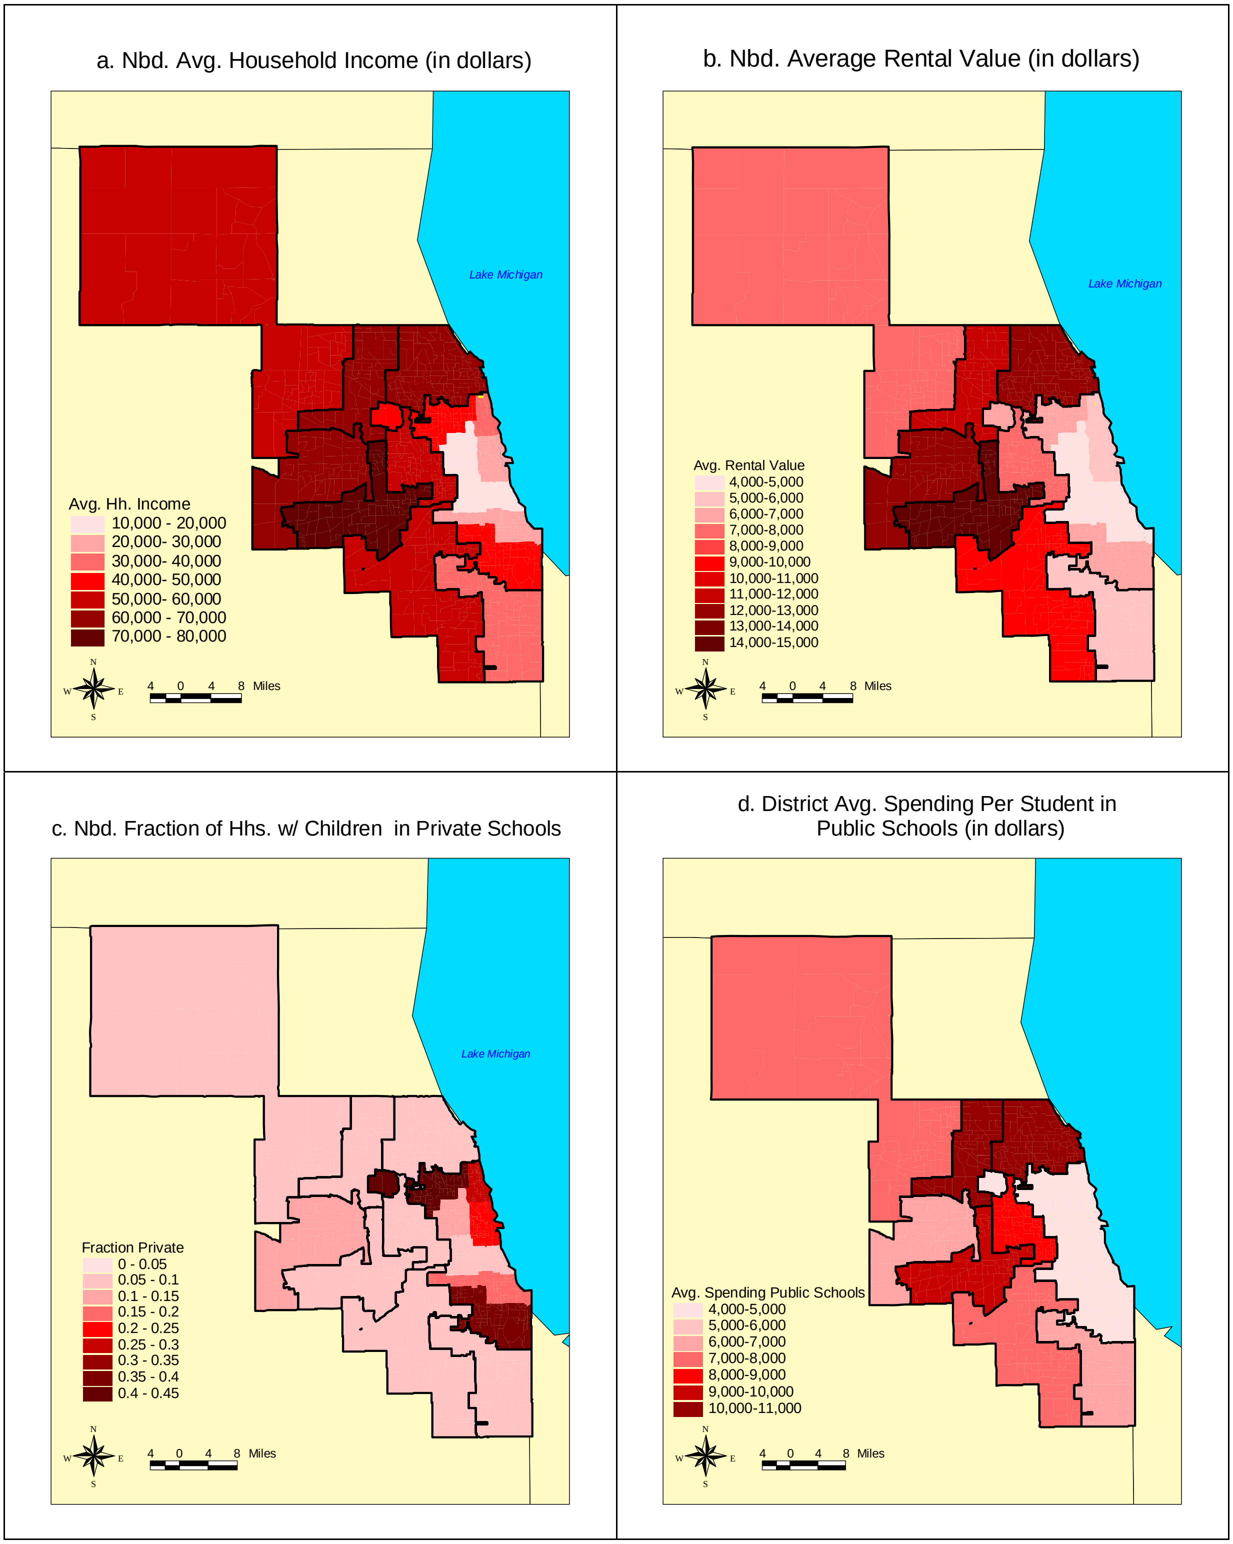

FIGURE 5 - ***Chicago: Predicted Household Sorting Across Schools for Benchmark Equilibrium, and Universal and Non-Sectarian Vouchers***

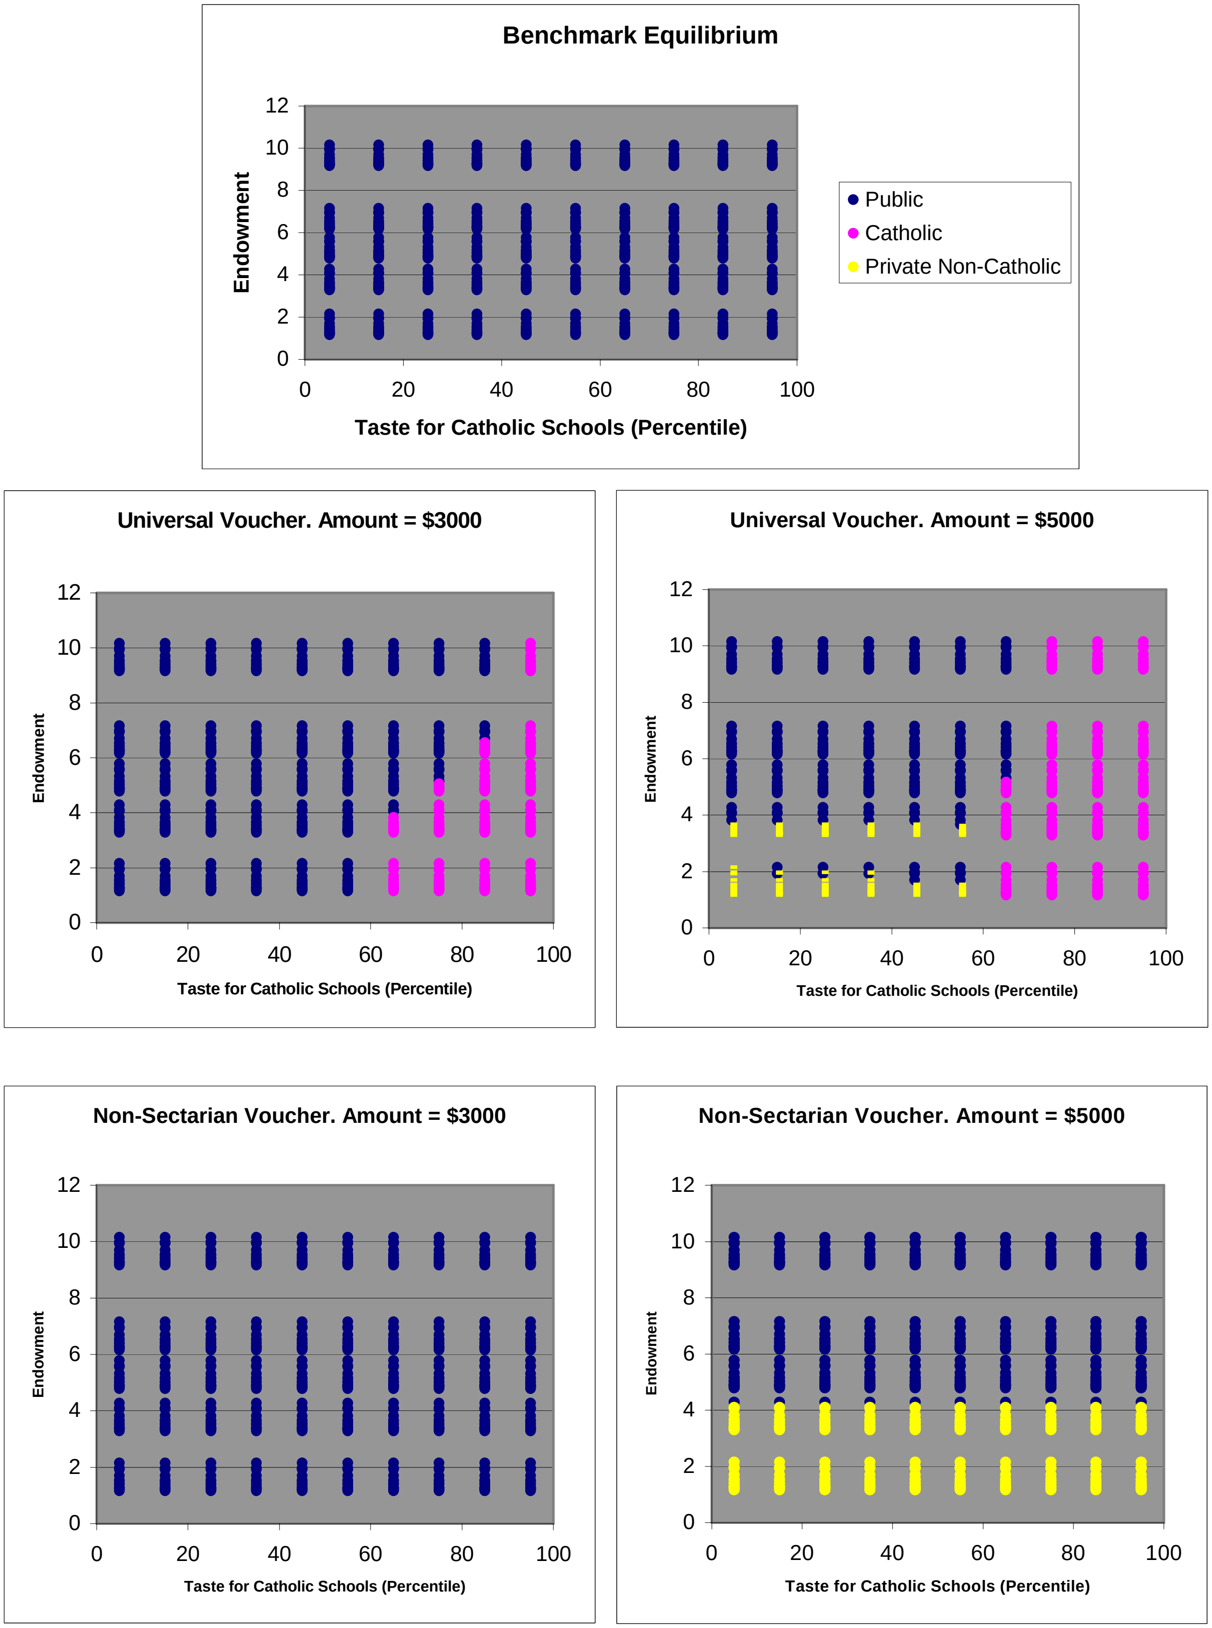

Note: Endowment expressed in $10,000. Each graph depicts the most popular choice made by each group of households with a given endowment and taste for Catholic schools.

FIGURE 6 – ***Chicago: Predicted Effects of Universal Vouchers on School Quality***

Average school quality for the 10th, 30th, 50th, 70th and 90th percentile endowment

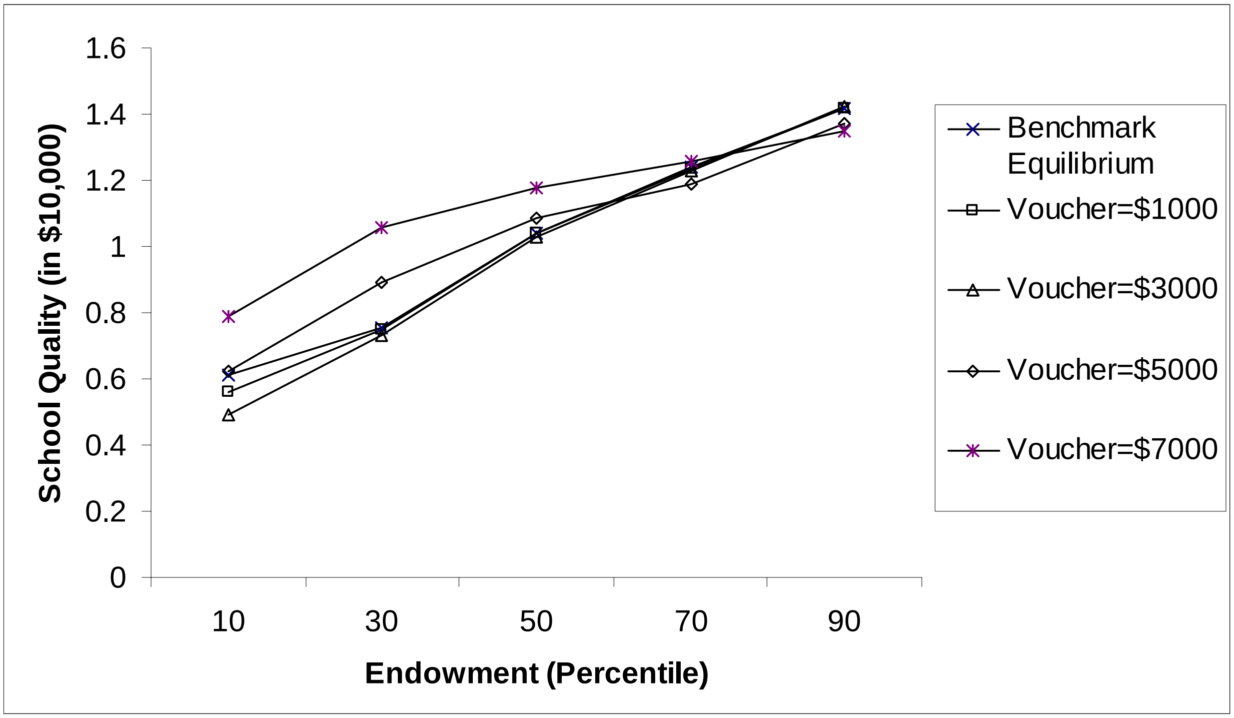

FIGURE 7 – ***Chicago: Predicted Welfare Gains by Household Endowment and Religious Preferences***

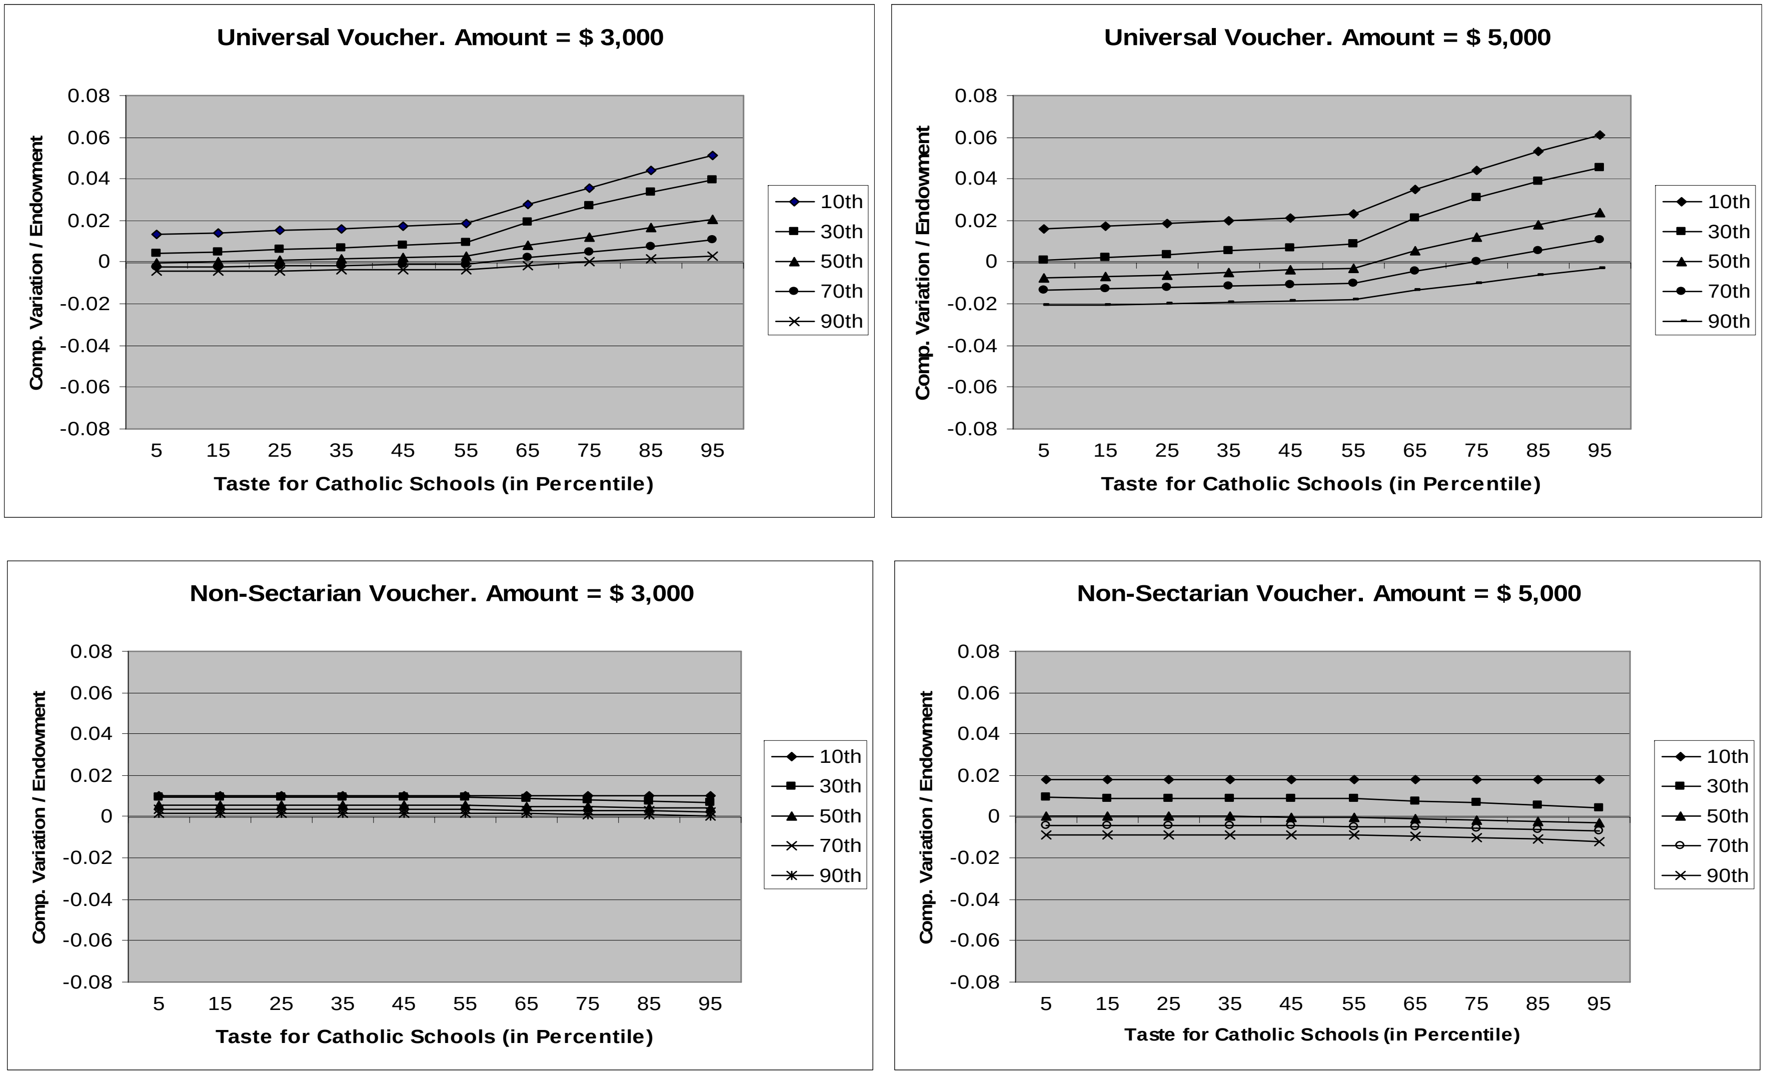

Note: each line represents a household endowment; households in the 10th, 30th, 50th, 70th, and 90th percentile of endowment are plotted.

Appendix

Computation of the Equilibrium

In this appendix, I first explain the algorithm used to compute the benchmark and voucher equilibria, and then discuss specific aspects related to the treatment of household location and school choice. This Appendix refers to Models 2, 3 and 4, which include household idiosyncratic tastes. The algorithm for Model 1 is analogous to the one presented here, although with specific provisions to deal with the discreteness of household types. See Ferreyra (2002) for further details on the algorithm for Model 1.

The computation of the benchmark equilibrium for the Chicago metropolitan area takes between five and ten seconds in a 3Ghz processor, and the computation of the voucher equilibrium for vouchers ranging from $1,000 to $7,000 takes between one and seven hours depending on the voucher program and amount.

The Algorithm: Overview

Chart I depicts the algorithm that calculates an equilibrium. The computation of an equilibrium consists of two nested loops: an outer loop of *major iterations* for voting and adjustment of community compositions and school quality, and an inner loop of *minor iterations* for the choice of location and school type. The sequence of major iterations concludes when agents cannot gain any utility by moving or switching to a different type of school, and all endogenous variables have converged.

The algorithm, which is coded in ANSI C++, follows the steps explained below. Notice that steps 3 through 5 comprise the inner loop, in which property tax rates, community compositions, public school quality, spending per student, number of people in public schools, and property tax base are held constant and taken as given by households when making choices.

1. For the benchmark equilibrium, set up the community structure for the computational version of the model, and define initial prices for houses and non-residential property in all locations. For the voucher equilibrium, take the benchmark equilibrium as the starting point.
2. Households vote for tax rates and spending in public schools. In the first iteration while computing the benchmark equilibrium, households vote for property taxes in the district where their endowed houses are located, as though they all attended public schools. Otherwise, households vote for taxes in the districts where they would prefer to live. A district's new tax rate is the one chosen by the district's median voter. Given the new tax rates, districts' spending per student and public school quality are updated.

3. Households choose their optimal location and type of school given the election outcomes; when considering a location, households evaluate the public and private schools available there.
4. Once all households have made their choice, the algorithm computes the excess of demand for houses in each location and adjusts the price of houses in all neighborhoods proportionally to each neighborhood's excess of demand according to the following price adjustment rule: $P_1(d,h) = P_0(d,h) + cED_1(d,h)$, where the subscript *1* stands for the current iteration, and *0* for the previous one. *ED(d,h)* denotes the excess of demand for housing in district *d* and neighborhood *h*. The parameter *c* is the adjustment factor. It is set to 0.01 at the beginning of every round of price adjustments, and it is adjusted using a bisection rule as departures are detected from the convergence path on the way to the equilibrium.
5. Check whether there is a nonzero excess demand in some location. If there is, go back to (3), which will in turn lead to a new adjustment of prices in (4). The inner loop of location choice and price adjustments continues until supply equals demand in all locations. In each of these minor iterations, the value of a household's endowment is updated as the price of its endowed house changes during price adjustments. Households keep their endowment house throughout the inner loop.
6. Repeat steps (3), (4) and (5) until prices clear the housing market in each location.
7. Adjust community compositions, non-residential property tax base, number of public school households, and public school quality for each district.
8. If households have experienced utility gains in the current major iteration by moving and/or changing schools, start a new major iteration. The outer loop continues until households cannot gain any utility by moving or switching to a different school, and all endogenous variables have converged.

Further Details on the Algorithm

a. Benchmark Equilibrium

To simplify the explanation of the computational treatment of household location and school choice, some additional notation is in order. Each household *i* in the economy belongs to a non-idiosyncratic household type *j*, which represents a combination of income endowment y_n, house endowment valued at p_n , and religious type *k*. This non-idiosyncratic type exists with measure μ_j , and the total number of non-idiosyncratic types is $J = I \times H \times K$. In addition, location (*d*,*h*) has a measure of houses equal to μ_{dh} , so that district *d* has a total measure of houses equal to $\mu_d = \sum_h \mu_{dh}$. Let *m* denotes school type; *m*=1, 2, 3 represents public, private Catholic and private non-Catholic school respectively. Furthermore, the triplet (*d*,*h*,*m*) denotes the joint choice of location (*d*,*h*) and school type *m*.

When choosing (*d*,*h*,*m*), household *i* of non-idiosyncratic type *j* obtains utility $U_{idhm} = s_{jdhm}^{\alpha} c_{jdhm}^{\beta} k_{dh}^{1-\beta-\alpha} e^{\varepsilon_{idhm}}$. Since all households of non-idiosyncratic type *j* choosing (*d*,*h*,*m*) experience the same consumption and parental valuation of school quality, we can write this utility as $U_{idhm} = U_{jdhm}^{*} e^{\varepsilon_{idhm}}$. Define $\widetilde{U} \equiv \log(U^{*})$. Hence, the share of non-idiosyncratic type *j* households who choose (d,h,m) equals $P\left(\widetilde{U}_{jdhm} + \varepsilon_{idhm} > \widetilde{U}_{jd'h'm'} + \varepsilon_{id'h'm'}\right)$ for all (d',h',m') not equal to (*d*,*h*,*m*). Since ε is distributed i.i.d. type I with scale parameter (1/*b*), this share equals $P_{jdhm} = \frac{\exp\left(\widetilde{U}_{jdhm}/b\right)}{\sum_d \sum_h \sum_m \exp\left(\widetilde{U}_{jdhm}/b\right)} = \frac{\left(U_{jdhm}^{*}\right)^{1/b}}{\sum_d \sum_h \sum_m \left(U_{jdhm}^{*}\right)^{1/b}}$. Thus, the total demand for any given location (*d*,*h*) is $\sum_j \sum_m P_{jdhm}\mu_j$, and the prices p_{dh} that equalize the demand for each location to the supply μ_{dh} clear the housing market.

With regards to the voting equilibrium, households of non-idiosyncratic type *j* who choose (*d*,*h*,*m*) vote for their optimal property tax rate t_{jdhm} in district *d*'s polls. The measure of such households is $P_{jdhm}\mu_j$. The median voter for this district is the household who votes for a tax rate equal to the median of the distribution of the selected tax rates.

b. Voucher equilibrium

Since households of non-idiosyncratic type *j* vary in their idiosyncratic preferences, they potentially make different benchmark equilibrium choices of location and school. Furthermore, their different location choices imply different amounts of capital gains or losses in the voucher equilibrium, which gives rise to different budget constraints. Hence, for computational reasons I re-define a non-idiosyncratic type as the set of households with the same income and house endowment, religious type, *and* benchmark equilibrium location choice. In other words, non-idiosyncratic type *v* is the set of households of the *original* non-idiosyncratic type *j* that choose benchmark equilibrium location $(\hat{d},\hat{h})$. In this re-definition, the number of non-idiosyncratic types equals $V = J \times H$. Hereafter, an asterisk denotes the benchmark equilibrium. Non-idiosyncratic type *v* has measure $\mu_v = \left(\sum_m P^*_{j\hat{d}\hat{h}m}\right)\mu_j$; of course, $\sum_d \sum_h \sum_m P^*_{jdhm}\mu_j = \mu_j$. During the computation of the voucher equilibrium, the budget constraint for type v equals $c + (1+t_d)p_{dh} + \max(T - \upsilon) = (1-t_y)y_n + p^*_n + \left(p_{\hat{d}\hat{h}} - p^*_{\hat{d}\hat{h}}\right)$, where υ is the voucher amount, $p^*_{\hat{d}\hat{h}}$ is the benchmark equilibrium price of the house chosen by the original idiosyncratic type *j* in location $(\hat{d},\hat{h})$, $p_{\hat{d}\hat{h}}$ is the price for that house in the current iteration, and p^*_n is the proceeds from selling the endowed house in the benchmark equilibrium.

Denote by $V_{idhm} = V^*_{vdhm}e^{\varepsilon_{idhm}}$ the utility enjoyed by household *i* of non-idiosyncratic type *v* by choosing (*d*,*h*,*m*) under vouchers, and define $\tilde{V} = \log(V^*)$. Under vouchers, the share of households of non-idiosyncratic type *v* who choose (*d*, *h*, *m*) equals $P\left(\tilde{V}_{vdhm} + \varepsilon_{idhm} > \tilde{V}_{vd'h'm'} + \varepsilon_{id'h'm'}\right)$ for all (*d'*, *h'*, *m'*) not equal to (*d*, *h*, *m*). From the definition of the non-idiosyncratic type *v*, this share equals the share of households of the original non-idiosyncratic type *j* who choose (*d*, *h*, *m*) under vouchers conditional on having chosen $(\hat{d},\hat{h})$ in the benchmark equilibrium, or

$$P\left(\tilde{V}_{jdhm} + \varepsilon_{idhm} > \tilde{V}_{jd'h'm'} + \varepsilon_{id'h'm'} \mid \max_m\left(\tilde{U}_{j\hat{d}\hat{h}m} + \varepsilon_{i\hat{d}\hat{h}m}\right) > \max_m\left(\tilde{U}_{j\bar{d}\bar{h}m} + \varepsilon_{i\bar{d}\bar{h}m}\right)\right) \text{ for all } (d,h,m)$$

different from (d',h',m') and all $(\hat{d},\hat{h})$ different from $(\bar{d},\bar{h})$.

Since this share does not have a closed form solution, I compute it by simulation as follows. I randomly draw *R* independent vectors of idiosyncratic tastes, each one of dimension $H \times 3$, such that all vector elements come from a type I extreme value distribution with scale parameter (1/*b*). I use the same *R* vectors for each non-idiosyncratic type *v*. For each non-idiosyncratic type *v*, I keep the N_v vectors whose elements are consistent with the type having chosen its benchmark equilibrium location. Choosing a sufficiently large *R* ensures that the simulated shares approximate well the benchmark equilibrium closed-form shares. In any given iteration, I compute the share of households of type *v* making choice (*d*,*h*,*m*) as $\hat{P}_{vdhm} = N_{vdhm} / N_v$, where the numerator is the number of vectors of idiosyncratic tastes that render (*d*,*h*,*m*) as the optimal choice for households of non-idiosyncratic type *v*. The computation of the demand for a given location, and of the voting equilibrium, is analogous to that in the benchmark equilibrium.

For the analysis of the outcomes of policy simulations, I proceed as follows. Since voucher effects differ across households of even the same non-idiosyncratic type, I keep track of the benchmark and voucher equilibrium choices for each household involved in the voucher simulation. In the case of

Chicago, the computation of the benchmark equilibrium manipulates *J*=750 non-idiosyncratic types, whereas the computation of the voucher equilibrium handles *V*=750x15=11,250 non-idiosyncratic types. These, in turn, become 750,000 households in order to simulate the choice shares as described above. By storing pre- and pos- voucher information for each of these households, I am able to estimate the distribution of voucher effects for any desired group of households.

Chart I. Algorithm to Compute the Equilibrium