

IDEAL VOUCHERS

CAROLINE M. HOXBY*

Abstract

School access for disabled, poor, minority, or other groups of students are essentially externality problems. Such problems generically plague schools (not just choice programs), but school choice discussions often run aground on such issues. This paper argues that it is inherently *easier* to address such issues with vouchers than with the traditional public school system because vouchers are prices that can be set so as to internalize externalities, thereby inducing families and schools to make self-interested decisions that remedy externality-based problems. I describe a market-market mechanism for obtaining "externality-internalizing" vouchers and show how students with such vouchers can be given maximum choice among schools. The mechanism rewards schools that, given the same resources and access to students, do a better job of attracting parents. I consider passive parents who do not make choices for their children, optimal school composition, rewards for schools that follow state guidelines, and allowing vouchers to be topped up.

* Department of Economics, Harvard University and the National Bureau of Economic Research. The author gratefully acknowledges research support from the Carnegie Scholars program of the Carnegie Foundation and from the Bastiat Institute. For their helpful comments, the author gratefully thanks Martin Feldstein, William Gentry, Claudia Goldin, Jonathan Gruber, Lawrence Katz, Bridget Long, Michael Murray, Jeffrey Myron, Albert Saiz, and numerous seminar participants.

I. Introduction

When vouchers or other school choice reforms (such as charter schools) are discussed, certain worries arise. Will children with disabilities or limited English proficiency be barred from exercising choice because every school will find them too expensive to educate? Will choice increase the isolation of poor or minority children because it enables other children to choose schools that are inaccessible to them by virtue of high tuition or discrimination in admissions? Will there be rewards for good schools and will they be able to expand? Will there be safeguards against bizarre curricula? What will happen to children whose parents are so disengaged that they do not make an active school choice, even when it is fairly obvious that their current school is a failure?

These problems are not specific to school choice, they are generic problems of elementary and secondary schooling. In traditional public school systems, attempts to solve these problems are embodied in states' school finance laws, complex systems of mandates, quotas, and grants layered on top of local finance (locally controlled public schools that derive revenue from local taxes). Despite years of legislation and court action aimed at improving school finance law, its solutions are not widely regarded as successful. For instance, the isolation and inadequate education of poor children continue to be the subjects of perennial court battles.¹ Also, although desegregation plans have been in effect for 25 to 30 years in most big cities, big cities' black students are no more integrated with white students than they were in the early 1960s.² In short, the right question is not whether these issues arise in school choice. The right question is whether school choice will make it harder or easier to address these issues. To put it another way, let us assume that the issues addressed by current school finance law are a set of issues

¹ See Ladd, Chalk, and Hansen [1999] for an overall review or Hoxby [forthcoming] for an economic treatment.

² See Orfield [1996] and Reber [2001].

that any school system should address. Can vouchers address these issues successfully?

This paper argues that it is inherently easier to solve the problems addressed by current school finance law with vouchers than with the traditional system. The reason is simple: the problems are generated largely by externalities. Vouchers offer greater opportunities to resolve the problems because vouchers are *prices* that can be set so as to internalize externalities, thereby inducing families and schools to make decisions (in their own self-interest) that remedy what are essentially social problems. Prices can be enormously powerful and flexible instruments. I show that one can obtain "externality-intenalizing" vouchers with a *market-like mechanism*, so that vouchers address the aforementioned problems *and* maintain the flexibility and responsiveness (to families' preferences and schools' supply conditions) that we ordinarily associate with market-based prices.

The plan of the paper is as follows. I show that the aforementioned problems are essentially externality problems. I propose a market-like mechanism for obtaining vouchers under which the externalities are actually internalized. (This is the key part of the paper.) I then show that there is a mechanism that maximizes families' school choice, once they are furnished with the externality-internalizing vouchers. This mechanism also rewards schools that do a better job of attracting parents, given the same resources and access to students. In the final sections, I take up less central problems: passive parents who do not make choices for their children, optimal school composition, rewards for schools that follow state guidelines, and allowing vouchers to be topped up.

A. A Practical Approach

I attempt to confine the paper to voucher procedures that are *feasible*. The title is "Ideal Vouchers," but the paper really focuses on vouchers that are as nearly ideal as is practical. Also, although the discussion is phrased in terms of "vouchers," everything could be written with charter schools in mind. Simply substitute "charter school fee" for "voucher." The voucher systems I discuss are universally

available to students within a district, not (say) given to 1000 lucky students who win a lottery.³

It will be useful for the reader to think of most schools in a state, including all existing public schools, being under the voucher system. That is, existing public schools are "voucherized" or "charterized," making their control indistinguishable from that of schools that were begun as voucher or charter schools. Existing private schools can also be voucherized or charterized, but only if they use the voucher system with all of their students. (That is, they cannot admit some students based on the school choice mechanism I describe and admit others under another system of their own. The final section, on topping-up, contains more discussion of private schools that want to deviate from the voucher program.) Of course, existing private schools could choose not to participate in the voucher system at all, leaving them in their current situation—in which their finances are entirely based on private sources.

A side-effect of constraining the paper to be practical is that I accept much of the structure of current school systems. I focus on the district as the unit that controls local finance and the state as the unit that controls categorical aid for the disabled, limited English proficient, *et cetera*.

Once one begins to consider the capacity of vouchers, it is tempting to propose vouchers that solve any number of social problems, including problems that are so subtle and seemingly intractable that school law experts have not even suggested that they be addressed. In this paper, I strongly resist such temptations, in part because there is no consensus on the nature of many social problems and in part because the key benefit of vouchers is improving the productivity of schools. (Both empirical and theoretical work suggest that improving school productivity would be the primary benefit of school choice.) For instance, there is no consensus on optimal ability grouping of students. This is not only because the problem is an ethical snarl and has a complex, general equilibrium nature. It is also because

³ Vouchers with enrollment constraints are not enacted because they have desirable qualities. They are enacted because they are cheap as policy experiments and because vested interests (such as teachers' unions) insist upon enrollment constraints in order to limit their own losses.

there is very little empirical evidence on how students are affected by the abilities of their schoolmates. We do not generally know even the *sign* of the effect of having one schoolmate substituted by another. We certainly have no idea about the effect of many classroom compositions that are possible but rare in practice. Given the fact that the productivity of American schools is rapidly worsening,⁴ it would be folly to entangle a choice program that has the potential to improve productivity in interminable debates about problems that are not contemplated by current school finance law.

In short, the issues that I consider are those, and only those, that one can identify in current school law. I consider racial issues because race-based mandates do appear in current school law, but if such laws are overturned (as seems possible), it may be less appropriate to consider racial issues.

B. Vouchers and Mandates Differ, as Means of Solving Social Problems

There is an essential difference between how externalities are handled by a voucher system and the traditional public school system. Vouchers can be set to price externalities, so that someone who has a negative externality imposed upon him receives compensation for it. The traditional public system imposes largely unpriced externalities on parents who are incumbents in some school. The hope is that these parents will find it so expensive to switch schools that they will allow themselves to be imposed upon and (thereby) permit the burden to be distributed. In fact, the reliance of the current system on unwary incumbent parents is obvious in the debate over school choice. Opponents of choice routinely suggest that there are parents who would, if they were able, escape the burden that they are bearing for society.

To make this point concrete, consider the education of disabled children, who impose an externality on their school fellows because they are unusually expensive to educate. The addition of disabled children tends to reduce the regular per-pupil spending experienced by non-disabled children in

⁴ See Hanushek [1994] or Hoxby [2001].

their district (or increase the taxes required to fund a given level of regular per-pupil spending).⁵ In the traditional public system, a district must accommodate all disabled children who reside in its attendance area by providing them with the services listed in their Individual Education Plans. Mandatory accommodation translates into forcible imposition of the costs on the district's population. There is no procedure for redress or compensation of parents who experience a negative shock to the inputs their children experience.

There are two problems with the traditional system of imposing externalities. First, it is *ex post* inequitable. Households are arbitrarily selected to bear a burden that should be spread more optimally across society. In fact, household selection is not just arbitrary: it is regressive. It is regressive because savvy, well-off people are least likely to let themselves be imposed upon and are least constrained by their economic circumstances to bear unforeseen burdens. Second, while the traditional system of imposing externalities works in the short-run, it works poorly in the long-run. As soon as people register the fact that residents of a certain school's or district's attendance area are being imposed upon, people take the burden into account when they consider moving into the area. Property prices fall, monetizing the loss of wealth imposed on the incumbents and making it impossible to fund regular per-pupil spending and the services in the Individual Education Plans at the previous level. The system is out of equilibrium and someone will move, plunging the system into another round of trying to impose the externalities.⁶ (Note take Summers' [1989] argument that mandated benefits can be efficient cannot be extended to apply to this situation.⁷)

⁵ For an introduction to the finances of education for the disabled, see Cullen [1999].

⁶ The last few sentences contain straightforward implications of the Tiebout [1956] model. A good formal treatment may be found in Epple, Filimon, and Romer [1984].

⁷ Summers' argument is an extension of the labor economics literature on equalizing wage differentials. He focuses on benefits that an employer is mandated to offer workers. His argument does

II. Externality-Based Issues that can be Identified in Current School Law

In this section, I describe some problems that are typically addressed by states' school finance laws and show that each problem is essentially an externality problem. It is useful to introduce some common terminology. Students are "classified" into "categories," where one category is regular education and the remaining categories are "protected categories"—that is, categories for which school law imposes mandates.

A. Disability

The disabled are the most obvious example of negative externality-imposing students. This is because they are entitled to services, spelled out in their Individual Education Plans (IEPs), that are generally more costly than those to which non-disabled students are entitled. (An IEP is a legal agreement about the services that a school must provide to a specific disabled student. It is written with the advice of school administrators, counselors, psychologists, parents, and attorneys.) The externality imposed is a simple function of the extra cost (above and beyond regular per-pupil spending) of fulfilling the IEP..

In most states, disabled students are classified according to the type and intensity of their disability. These classifications tend to be reasonably good indicators of how much a disabled student will cost, although the details of IEPs make costs vary within classifications. From now on, I describe disabled students as though their classifications corresponded closely to the cost of educating them. To the extent that current classifications do not fulfil this condition, I assume that the classifications could be

not apply for a combination of two reasons. First, school mandates impose externalities on schoolfellow. Second, it is illegal for school districts to have individual-specific tax rates or allow individual students to experience a level of regular education spending that differs from that which other students experience. Thus, some people who do not benefit from the mandate pay its costs, while others who equally do not benefit pay no costs. Also, there is no force pushing the beneficiaries of the mandate towards the districts that can fulfil the mandate most cost effectively.

altered. In my examples, I describe learning disabled students and physically disabled students as though each group had one cost associated with it. I make this obvious simplification purely to ease my illustrations.

B. Limited English Proficiency

Limited English proficiency is, essentially, a special case of disability. Like disabled children, children with limited English proficiency are entitled to a variety of services that exceed those to which regular students are entitled. The level of services is related to a student's knowledge of English. For instance, students with partial English proficiency are often assigned to an English-as-a-Second-Language, while students with no English proficiency are often assigned to Bilingual Education. One can price the extra cost of providing the required language services. The externality imposed is a simple function of this extra cost.

One thing that makes English proficiency differ from disability is that the services to which a child is entitled often depend on how many other children belong to his language group in his district. For the purposes of this paper, one should think of an English proficient student as being classified by the services that he is entitled to receive, not just by his home language and knowledge of English.

C. Race

If one says that white parents discriminate against minority students as school fellows for their children, this is no more than saying that white parents *act* as though minority children impose negative externalities. Minority students may impose only imagined externalities (purely psychological discomfort) or they may impose real externalities because their racial background makes them systemically more expensive to teach. The discussion of race in schools suggests that real negative externalities, to the extent that they exist, may be indirect—that is, not emanating the students themselves. If school staff, say, have bad attitudes towards minority students and consistently under-perform in schools where

minority students are enrolled, then externalities associated with minority students emanate from staff. If bureaucrats in charge of allocating school inputs routinely discriminate against minority students, then externalities associated with minority students emanate from bureaucrats. (Desegregation cases often cite evidence that suggests that black students are assigned worse buildings, less qualified teachers, and so on.) Any form of discrimination against minority students that spills over onto their school fellows is a negative externality. It may be useful to think of the externality as the cost of bribing people to provide a minority student with the same inputs as a non-minority student would receive.

Although it is uncertain how the negative externalities associated with minority students arise (and even whether they are real or just imagined), it is easy to see that current school law treats minority students as though they imposed externalities. Minority students are a protected category in most states' school laws. I take this situation as given.

D. Poverty

Poor students impose negative externalities because their parents cannot afford to give them necessary services and inputs that better-off children receive from their parents. For instance, in order to concentrate in class, poor students may need to be provided with meals at school. Schools may need to buy books, supplies, and even computers for poor children that would be provided by well-off parents. Poor students are more likely to use after-school programs and health services provided by a school. Because their parents often lack stable housing and employment, poor students tend to move often among schools. Such movement imposes administrative and teaching costs on schools.

For these reasons and more, poor students impose a negative externality, the size of which depends on the cost of making them equivalent to a student whose well-off parents provide a stable home that is more rich in inputs.

III. “Simple” Externalities

At this point, it is useful to define “simple” and “complex” school-based externalities. A student is defined as imposing a simple externality if his presence imposes a cost that does not vary with the enrollment composition of the school. For instance, it might cost a school \$10,000 to fulfil a disabled student’s Individual Education Plan. The externality is simple if the \$10,000 amount does not vary with the school’s number of students, percentage of students who are black, and so on.

In truth, many externalities are probably not simple, but there are good reasons for focusing on simple externalities initially, as I do in the following sections. (I take up complex externalities in the final section.) First, many externalities are approximately simple. Although a school might be able to fulfil the IEP for a severe disability at much lower cost if it had a critical mass of children with the disability, children with serious disabilities are such a small share of the population that such critical masses are unrealistic unless disabled children are given little or no school choice. For common disabilities (say, minor learning disabilities), economies of scale are reached at such a low number of students that the externality is approximately simple. Second, states apparently believe it would be impractical to incorporate anything but simple externalities into school law. To the extent that current legal mandates address externality-based problems, they treat all externalities as though they were simple. For instance, a school’s compensation for a disabled child depends on that disabled child’s presence, not on interactions between that child’s presence and the presence of other children. Focusing on simple externalities is, thus, in keeping with my goal of describing vouchers that are practical and that deal with externalities at least as well as the current law does. Third, as I describe below, designing vouchers for complex externalities is a task that requires a great deal of (external) information. It is a task that is most likely to be worth pursuing for externalities that are both highly complex and large in scale—say, blind students.

IV. A Market-Like Mechanism that Generates Externality-Internalizing Vouchers

In this section, I describe a market-like mechanism that generates vouchers that internalize simple externalities. That is, the vouchers generated are such that schools are indifferent to a student's category. I also show a simulation of how the mechanism works.

A. The Notion of Basic Vouchers and Voucher Add-Ons

The goal of this exercise is to obtain a set of vouchers "add-ons" that apply to the categories of externality-imposing students. Ideally, each educational market in a state would have its own set of add-ons. The functional translation of an educational market tends to be a metropolitan area or a statewide non-metropolitan remainder. However, given current practice, it is likely that most states would create a statewide set of add-ons. (The state is the level at which most categorical aid is currently funded and disbursed). Therefore, I will refer to add-ons as though they were state-wide, from now on.

A set of voucher add-ons implies a "basic" voucher for each district: one simply takes total expenditure of a district, subtracts what must be spent on the add-ons after examining the district's student population, and divides the remaining spending by the total enrollment of the district. In this way, one obtains the basic voucher that a regular education student (who fits into no protected category) would carry. A student in a protected category would carry a voucher equal to the basic voucher plus his category's add-on.

For instance, if a district's only externality-imposing students were physically disabled students and b_2 was their add-on, then one would solve for the district's basic voucher by solving for x_j in the equation:

$$(1) \quad [x_j (1 - \%SpEdl_j)] + [(x_j + b_2) \%SpEdl_j] = pps_j .$$

where pps_j is the school's statutory per-pupil spending and $\%SpEdl_j$ is the share of students who are physically disabled. More generally, one solves for the basic voucher by solving for x_j in the equation:

$$(2) \quad [x_j (1 - \%SpEd1_j - \%SpEd2_j - \dots - \%Biling_j)] + [(x_j + b_2) \%SpEd1_j] + [(x_j + b_3) \%SpEd2_j] + \dots + [(x_j + b_{k-1}) \%Biling_j] = pps_j .$$

In equation (2), $\%SpEd2_j$ is the share of students who are learning disabled, $\%Biling_j$ is the percentage of students in need in of bilingual education, and the ellipses represent all of the other externalities discussed in Section II.

Nothing in this formulation precludes the possibility that a given student belongs to more than one externality-imposing group. For instance, a student could get a voucher equal to $x_j + b_2 + b_{k-1}$ if he were physically disabled and in need of bilingual education. Moreover, “interaction groups,” such as students who are simultaneously physically disabled and in need of bilingual education, can be categories of their own. In other words, we need not impose the restriction that a student who belongs to multiple categories gets the sum of the add-ons; he could get a smaller or larger add-on so long as we make “multiple categories statuses” into categories of their own.

It is very important to realize that the basic voucher incorporates all factors that affect school finance, apart from those based on categories of students. The basic voucher will reflect local demand for school spending. The basic voucher will also reflect non-categorical revenues that a district receives—for instance, school finance equalization revenues. The basic voucher will be high in states with high per-pupil spending. In states where districts are largely financially autonomous and per-pupil spending varies a lot from district to district, the basic voucher will vary a lot. (Below, I discuss how parents can take basic vouchers across district lines.) In short, in an effort to be practical, I am describing a voucher that allows states to retain as much of their current spending determination as possible. The level and equality with which states ought to fund educational spending are logically separate issues from vouchers, and they require separate analysis.

B. The Market-Like Mechanism

For now, let us take for granted that the state begins the process with a set of initial vouchers (a set of add-ons and the basic vouchers implied for each district). In the next section, I show how the state can use empirical evidence to obtain a set of initial vouchers for the first year of the program. For now, it is enough to know that the initial vouchers are an *attempt* to price the externalities associated with student categories. In the first year of a voucher program, the initial vouchers are likely to be in the ballpark of what the market-like mechanism will generate, but this is all they are likely to be. In subsequent years, the initial vouchers will be the vouchers generated by the market-like mechanism in the previous year, so they are likely to be close to what the mechanism will generate.

It is worth noting in advance that parents are wholly *uninvolved* in the market-like mechanism and do not need to know what the initial vouchers are. Indeed, as will be seen later, parents *never* need to know what add-ons the market-like mechanism generates. Parents only need to know the basic vouchers. This is akin to current practice, in which parents often know what local districts' regular spending is, but almost never know the formulae in the categorical programs of their states' school finance laws.

Require each school to rank the students in its district (and students who are transfers from other local districts—see below) in order of preference. Such ranking will necessarily be somewhat crude and is *not* analogous to the ranking performed by highly selective colleges or graduate programs. Elementary and secondary students have very short *curriculum vitae*, so the ranking will necessarily focus on data in students' administrative files (especially their categories, if any) and knowledge that the school gathers about students during open houses or similar events (for instance, how much the family likes the school's mission). All schools will be provided with the same basic information, so their rankings may be quite similar, except for differences due to location and mission (a school may have a classical curriculum or an

arts focus, say).

If the initial vouchers happen to be exactly right, which is unlikely, the schools' rankings will show that—while schools prefer specific students within each category—they do not prefer whole categories of students. That is, schools' preference orderings will show them to be generally indifferent between enrolling a physically disabled child with the larger voucher, say, and a regular education child with basic voucher. More precisely, if one tests each school's preference ordering against the multinomial distribution (in which the category shares are the empirical category shares among district students), each school's ordering will not reject the distribution *regardless of the part of the ordering that one tests*. For instance, suppose that one divides a school's ordering into, say, quarters. Then, none of the four quarters will reject the multinomial distribution. More generally, one needs a statistical test for the school having sampled from its multinomial distribution in a non-random order. There are non-parametric statistical tests for this problem that are based purely on rank. See Randles and Wolfe [1991]. In the spirit of practicality, however, I suggest the following test, which is easy to implement using multinomial logit estimation.

Assign each category of student an arbitrary number $c=0, \dots, C$. The value of each number does not matter, since multinomial logit assumes that the categories are unordered. For ease of interpretation, though, basic students (students who do not fit into any externality-imposing category) should be assigned to category 0. The applicants of each school are its observations, and the category number assigned to each student is the data upon which the dependent variable for the multinomial logit regression is constructed. For instance, the category assignment vector for applicants $k=1, \dots, K$ to school j might look like:

$$\text{Category}_i = \begin{bmatrix} \text{Basic} \\ \text{SpEd1} \\ \text{Basic} \\ \text{SpEd2} \\ \text{Biling} \\ \text{SpEd1} \\ \text{Basic} \\ \text{Basic} \\ \dots \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 2 \\ C \\ 1 \\ 0 \\ 0 \\ \dots \end{bmatrix}$$

The above vector shows only the first eight students in the school's preference ordering; a real vector would obviously be longer. The preference rank assigned to each student is the sole independent variable in the regression, apart from the constant. The multinomial logit model lets the probability that a given student belongs to a group c depend on his rank, according to the following functional form:

$$\text{Prob}(\text{Category}_i = c) = \frac{e^{\lambda_0 + \lambda_1 \cdot \text{rank}_i}}{1 + \sum_{c=1}^C e^{\lambda_0 + \lambda_1 \cdot \text{rank}_i}} \quad \text{for } c = 1, 2, \dots, C,$$

$$\text{Prob}(\text{Category}_i = 0) = \frac{1}{1 + \sum_{c=1}^C e^{\lambda_0 + \lambda_1 \cdot \text{rank}_i}}.$$

Estimating the multinomial logit equation allows us to determine whether students' category assignments systemically affect their rank. Put another way, if rank has a statistically significant association with the probability that a student belongs to a certain category, then it must be true that the school is not indifferent among categories of students. In order to perform this test, we need only see whether the coefficient on the rank variable is statistically significantly different from zero for a category of students. This is because the probability that a student belongs to category c , as opposed to the basic category, is a

simple function of the coefficients:

$$\ln \left[\frac{\text{Prob}(\text{Category}_i = c)}{\text{Prob}(\text{Category}_i = 0)} \right] = \text{lamba}_0 + \lambda_1 \cdot \text{rank}_i .$$

Moreover, the sign and size of the coefficient will tell us whether and how much the category is preferred to basic students. A positive, statistically significant coefficient on rank tells us that the category is preferred to basic students; the larger the coefficient, the more the category is preferred. A negative, statistically significant coefficient on rank tells us the opposite. Naturally, one can be more or less stringent about deviations from category-indifference by varying the level at which one declares a coefficient to be statistically significant. For instance, the 0.25 level of statistical significance requires schools to show more category-indifference than the 0.10 level.

Suppose that the add-ons are insufficient to make schools indifferent between regular education students and a particular type of externality-imposing students—say, physically disabled students. Then, the coefficient for that category of students will generally be negative and statistically significant in schools’ multinomial logit results. The state should raise the add-on for this category and then adjust the basic voucher so that vouchers comply with the adding up constraint (equation 2). Then, the state should ask schools to re-order their applicants. This process should iterate until no school’s ordering rejects the null hypothesis of category indifference. (Below, I suggest that states may prefer to set a more lenient threshold, such as allowing 5 percent of schools to reject category indifference. Also, a state may prefer to exempt certain schools from exhibiting indifference about certain categories.)

In summary, the algorithm for adjusting the voucher is as follows. It requires negligible computer time, so the only substantial effort is made by the schools in creating preference orderings.

1. Set a significance level for the tests, such as 0.25. Set a threshold for the share of schools that

will be permitted to reject the null hypothesis of category indifference in their multinomial logit regressions. This threshold may be zero percent. Set a factor for adjusting the vouchers add-ons in each iteration. The factor should be inversely related to the estimated multinomial logit coefficient for the category in question. (I give an example below.)

2. Run a multinomial logit equation using the data from each school's preference ordering of applicants. For each category of students, record the share of schools that show a statistically significant preference for (against) the category.
3. If the share of schools that show a statistically significant preference for a category exceeds the threshold, proceed with the following steps. If not, stop.
4. If the category is preferred in the majority of schools that had statistically significant coefficients, lower the add-on for the category using the adjustment factor. If the category is disfavored in the majority of schools that had statistically significant coefficients, raise the add-on using the adjustment factor.
5. Back out the basic vouchers that are consistent with the new set of add-on vouchers.
6. Require schools to re-order students, taking into account the students' new voucher amounts.
7. Repeat steps 2 through 6 until the standard for stopping in step 3 is met.

At this point, it is useful to consider the preference orderings generated by the algorithm. The algorithm guarantees that the preference orderings exhibit category-indifference (or near category indifference, if the thresholds and significance levels are lenient). That is, the basic voucher and add-ons must be such that schools do not systemically prefer to enroll one category of their applicants over another. The add-ons prevent discrimination against externality-imposing categories of students by counterbalancing (and thereby internalizing) the externality that schools perceive. There is no simple way

for a school to “game” this algorithm because, as we shall see, the allocation mechanism described in the next section guarantees that a school’s likelihood of enrolling a particular student depends on how highly it ranks him. If schools believe that a category of students is mis-perceived and that their add-on is too high, then the schools must give the category of students relatively high ranks in order to get them and arbitrage the add-ons. This arbitrage will drive down the add-on. If a school does not rank students according to its true preferences, it runs the risk of enrolling numerous students whom it does not prefer. (See the next section for more on this point.)

Although the algorithm prevents discrimination against categories of externality-imposing students, it does not attempt to prevent other preferences from revealing themselves. In particular, schools may prefer certain students within each category: John Smith may be universally preferred to Tim Foster within the basic category of students. Alternatively, only the school with the classical curriculum may prefer John Smith to Tim Foster because John’s family seemed more enthusiastic about the curriculum.

The following simulation shows the algorithm in action. Suppose that the following conditions prevail.

- (a) There are 500 students who must be allocated to 5 schools, each of which has 100 places.
- (b) Each student has a desirability factor that is constant across schools. This desirability factor is drawn randomly from a normal distribution with mean 5 and variance 2.
- (c) 100 of the students impose a small negative externality equal to -1. 50 of the students impose a large negative externality equal to -3. The remaining 350 students impose zero externalities.

Students are randomly assigned to an externality category.

- (d) Each student has an idiosyncratic match quality with each school. This match quality is drawn randomly from a uniform distribution over $[0,2)$.
- (e) Schools rank students based on the sum of their desirability factor, the externality they impose,

their idiosyncratic match quality, and their voucher add-ons.

Clearly, the externality-internalizing voucher add-ons in this situation are 1 for the students who impose an externality of -1, and 3 for the students who impose an externality of -3. Suppose that the initial add-ons are 0.6 and 3.9. Let the threshold for rejecting a school's ranking be a stringent 0.25 level of significance. Let the add-ons be adjusted according to the rule: $\text{add-on}_{\text{this iteration}} = \text{add-on}_{\text{last iteration}} \cdot (1 - \text{multinomial logit coefficient} \cdot 100)$, where the multinomial coefficient is the one that applies to the category for which the add-on is being adjusted.

The results of the algorithm are as follows. In the first round of multinomial logit testing, every school's ranking is rejected—for both types of externality—at the 0.10 level. In other words, rank is a statistically significant predictor of category at the 0.10 level, in every school. Both add-ons are adjusted: 0.6 becomes 0.79 and 3.9 becomes 2.70. In the second round of multinomial logit testing, one school's ranking is not rejected, even at the 0.25 level. One school has its ranking for the -3 externality rejected, but does not have its ranking for the -1 externality rejected, even at the 0.25 level. The remaining three schools have their rankings rejected for both externalities, at the 0.10 level. Both add-ons are adjusted: 0.79 becomes 0.87 and 2.70 becomes 3.21. In the third round of multinomial logit testing, one school's ranking is not rejected, even at the 0.25 level. Two schools have their ranking for the -3 externality rejected, at the 0.25 level. Two other schools have their ranking for the -1 externality rejected, at the 0.25 level. Both add-ons are adjusted: 0.87 becomes 0.93 and 3.21 becomes 3.10. In the fourth round of testing, none of the schools' rankings are rejected, even at the 0.25 level. The procedure therefore stops.

How well does the algorithm work? One measure of the schools' category indifference is whether externality-imposing students are spread evenly across schools' final rank-order lists. The following tables show how the students who impose externalities are spread.

Table 1: Number of -3 Externality Students in Each Quintile of Schools' Rank Order Lists

	school 1	school 2	school 3	school 4	school 5
rank of 1 to 100	11	9	10	10	9
rank of 101 to 200	7	13	10	7	8
rank of 201 to 300	11	7	10	12	13
rank of 301 to 400	14	12	12	12	13
rank of 401 to 500	7	9	8	9	7

Table 2: Number of -1 Externality Students in Each Quintile of Schools' Rank Order Lists

	school 1	school 2	school 3	school 4	school 5
rank of 1 to 100	22	15	26	23	20
rank of 101 to 200	25	23	18	15	24
rank of 201 to 300	17	25	16	22	15
rank of 301 to 400	18	18	16	22	17
rank of 401 to 500	18	19	24	18	23

Tables 1 and 2 show that the -3 and -1 externality students are spread quite evenly. In particular, externality-imposing students are not notably overrepresented in the lowest ranked quintiles, which is what one fears would happen if the add-ons were too small for category-indifference. They are also not notably overrepresented in the highest ranked quintiles.

C. Exemptions from Category Indifference

Fundamentally, the add-ons are supposed to allow students from every category to exercise school choice. Every student should have the chance to attend numerous schools and should only be crowded out by other students in his category, with whom he is competing on roughly equal terms.

A state may decide that it is unnecessary for *every* school to exhibit category indifference because it may deem it sufficient that a student have access to 90 or 95 percent of local schools. If so,

the state should set the threshold (for step 3 of the algorithm) at 5 or 10 percent, instead of 0 percent.

The advantage of setting a more lenient threshold is that the state allows for variation in schools' ideas about what the externality-internalizing amounts are; the (obvious) disadvantage is that students' school choice is slightly limited.

More importantly, a few schools may have good reason to violate category indifference even when the add-ons are such the vast majority of schools are category indifferent. For instance, a school that specializes in the hearing impaired may prefer students in the hearing impaired category "too much." The state may wish to exempt specific schools from specific types of category indifference, while holding the schools to other types of category indifference. For instance, the school for the hearing impaired would presumably be required to exhibit category indifference across poor and non-poor, black and non-black, etc. students who are hearing impaired.

V. Obtaining Initial Amounts for the Vouchers

So far, we have simply assumed that the state has some way to obtain initial voucher add-ons. The initial voucher amounts are not crucial because, as has been seen, they are merely a starting point for the market-like mechanism in its first year of operation. Nevertheless, reasonable initial add-ons are desirable because they will ease the launch of the mechanism.

A. Sources of Empirical Evidence on Externalities: Cost Accounting Data and Capitalization

There are at least two sources of empirical evidence that can guide the state in constructing initial add-ons: cost accounting data and evidence from house price capitalization.

Cost accounting data are likely to be the best source of information on how much it costs a school to fulfil IEPs with well-defined service requirements. Accounting data also likely to be a good source of information on the cost of providing certain services to limited English proficient students. The better

defined are the services, and the rarer the category, the better will be cost accounting data relative to capitalization as a source of information. This is because cost accounting data are a more accurate source of information for well-defined services (such as how much the equipment for a hearing impaired student costs) and because capitalization is a less accurate source of information for categories that arise rarely. (The latter statement will become obvious in my discussion below.) A state could assemble cost accounting data on categories of disability and limited English proficiency by working with a committee of financial representatives from districts. Because the initial add-ons generated by the committee would be subjected to the market-like mechanism anyway, the committee would have little incentive to provide misleading data.

For other externality-imposing categories, such as poverty or minority status, accounting data will be of no help. Also, accounting data will be of little help for disabilities that are common but that have poorly defined services, such as general learning disabilities. The key to obtaining initial add-ons for such categories is recognizing that house prices capitalize externalities imposed upon residents.

Consider two school attendance areas *identical in all respects* except that area A has a learning disabled population that makes up 10 percent of enrollees and area B has a learning disabled population that makes up 0 percent of enrollees. We expect house prices to reflect the difference in the costs imposed upon (and, thus, the lower regular per-pupil spending or higher taxes experienced by) households in area A. More precisely, suppose that areas A and B have the same statutory per-pupil spending, but that households in area A pay \$100 less per year in implied rent and property taxes because their (identical) houses are worth less. The savings of \$100 per year should be exactly enough to compensate them for the disproportionate burden they bear (regular per-pupil spending is truly lower). If every household contains one school-aged child, then the “true” per-pupil spending experienced by a regular education student in area A is the school’s statutory per-pupil spending minus \$100. (Because some

households actually have no school-aged children and others have multiple school-aged children, a bit of algebraic manipulation is required to translate observed capitalization into its implications for add-ons and regular per-pupil spending. I present the algebra in the Appendix and proceed as though there were, in fact, one school-aged child per household.)

B. Hedonic Regressions for Obtaining Initial Add-Ons to the Basic Voucher

In the example above, it was easy to determine the value of the burden that special education imposes on local households because there was only one externality to consider and the two attendance areas were identical in every other respect. Needless to say, the real world is more complicated. One must try to sort out the effect of multiple externalities upon house prices while controlling for every other characteristic of houses. In other words, one must estimate a hedonic equation with annual housing costs (implied rent plus property taxes) as the dependent variable and measured externalities and house characteristics as the independent variables. Implied rent is usually estimated by multiplying the house's current market value by the sum of the interest rate and the house's depreciation rate. A sensible hedonic equation is:

$$(3) \quad [P_{ij} \cdot (r + \delta)] + P_{ij} \tau_j = \alpha_0 + \alpha_1 pps_j + \alpha_2 \%SpEd1_j + \alpha_3 \%SpEd2_j + \dots + \alpha_k SqFt_{ij} + \alpha_{k+1} Acres_{ij} + \dots + \alpha_{k+m} DistanceCity_{ij} + \alpha_{k+m+1} DistancePark_{ij} + \dots + \epsilon_{ij}$$

where i indexes the house and j indexes the school attendance area. P_{ij} is the house's current market value, r is the interest rate, δ is the depreciation rate of housing, and τ_j is the property tax rate. The omitted terms with coefficients α_4 through α_{k-1} are measures of the other externalities discussed in the previous section of this paper. The terms with coefficients from α_k through α_{k+m-1} record the value of house characteristics, such as a house's living area ($SqFt_{ij}$) and lot size ($Acres_{ij}$). The terms with

coefficients starting with α_{k+m} record the value of a house's locational and neighborhood characteristics, such as a house's distance from the city ($DistanceCity_{ij}$) and the nearest park ($DistancePark_{ij}$).

In the above equation, each externality is measured by the share of externality-imposing students in enrollment. Such measures assume that externalities are simple and that what matters to students is their *share* of the costs of an externality, their *probability* of experiencing an externality, or the *share* of their time that they spend experiencing an externality. There may be preferable ways of measuring externality-imposition in a school, but the measures in equation (3) are reasonable, especially given the fact that states will be interested in add-ons that treat externalities as though they were approximately simple.

To see how one can calculate add-ons from the coefficients in equation (3), suppose that a learning disabled student's presence raises a typical school's enrollment of learning disabled students (as a share of all students) by 5 percent. If one gives learning disabled students add-ons of $0.05 \cdot \alpha_3$, then their school fellows will typically be compensated for the externality. Once one has obtained the initial add-ons in this manner, one simply back outs the initial basic vouchers, as described in the previous section.

C. Solving the Identification Problem

There is an identification problem that plagues hedonic regressions like (3): omitted variables that are correlated with right-hand-side variables can generate biased estimates of the value of externalities, house characteristics, and local amenities. In particular, one worries that there are unobserved neighborhood characteristics such as neighborhood norms about the value of education and treatment of public spaces. Such omitted neighborhood characteristics affect house prices and can be correlated with the measured right-hand-side variables. Naturally, the more informative the set of right-hand-side variables is, the less bias one expects from omitted variables.

Because the purpose of the regression is obtaining *initial* add-ons that will be adjusted using the

market-like mechanism, one might decide that it was satisfactory to estimate regression (3) as it stands, using all available variables. Alternatively, as shown by Black [1999], one can restrict the sample of houses to those that are on the boundary of schools' attendance areas and can include boundary fixed effects in the regression:

$$(4) \quad [P_{jt} \cdot (r+b)] + P_{jt} = \beta_0 + \beta_1 PPS_{jt} + \beta_2 SpEd1_{jt} + \beta_3 SpEd2_{jt} + \dots + \beta_k SqFt_{jt} + \beta_{k+1} Acres_{jt} + \dots + I^{boundary} \gamma + v_{jt}, \quad \forall jt \in \{boundary 1, \dots, boundary N\}.$$

With this specification, one controls for all neighborhood characteristics, because houses on a boundary necessarily share the same neighborhood. Black, for instance, compares houses that are within 250 feet of a boundary. While equation (4) almost certainly produces estimates that are less subject to bias than equation (3) does, equation (4) is considerably more demanding of precise geographic information about the location of houses and boundaries. There is also a loss of precision when one estimates equation (4) because the restricted sample of houses is small.

D. An Empirical Example of Initial Vouchers

In this section, I present an example of estimating initial vouchers based on schools in the Detroit metropolitan area.

[Preliminary drafts of this paper do not show this example, in order to minimize the possibility that preliminary estimates are used by the public. The example is presented in oral presentations.]

VI. Implementing School Choice and Rewarding for Efficient Schools

With all the discussion of add-ons, it is easy to lose sight of the fact that the purpose of a voucher program is to allow students to exercise choice among schools. More precisely, the goal is to allow

students to exercise maximum choice, subject only to the constraint that they may crowd one another out of certain schools. (Of course, the point of the market-like mechanism is to guarantee that no category of student systemically crowds out another.) In addition to allowing students to match themselves to school programs, maximizing school choice is supposed to encourage a school to expand if it is unusually good at attracting students, given access to the same student population and (therefore the same possibilities for remuneration) as other schools. School choice is also supposed to encourage bad schools to contract or exit. In this section, I show how students, each associated with a voucher generated by the market-like mechanism, can exercise maximum school choice. It turns out that their choices produce inducements for good schools to expand and bad schools to contract.

At this point, it is necessary that families take the first and only action required of them: they must rank local voucher schools, in order of preference. (It is expected that families would take this action once for each level of schooling that their child required—elementary, middle, secondary—unless they moved or otherwise wished to reenter the choice arena.) Notice that this action does not require them to know the voucher add-ons. Indeed, although it is reasonable to expect that families would know their district's basic voucher and their child's category, there is no reason for them to know the add-on formulae (just as families do not know the categorical formulae in their states' school finance laws). Although the action that families should take is straightforward, the more effort that families devote to assessing schools, the better school choice will eventually work. Indeed, while an optimally designed voucher program should minimize the complexity of actions that families are required to take, it should encourage families to exercise maximum concern and diligence in taking the simple, required action. The fundamental idea of school choice is to give school correct incentives by exploiting families' information and unconflicted motives. (By unconflicted motives, I refer to the fact that families' utility is increasing in their children's utility. The same cannot necessarily be said for other constituencies in the education

sector.) Naturally, there is more to exploit when families are more informed and concerned. In a later section, I discuss a back-up system for students whose parents cannot or will not take an active role in choosing a school for them.

Having required families to rank schools, it is straightforward to implement school choice in such a way that families' choice is maximized. This is because, armed with parents' rankings and schools' rankings (which are necessarily category-indifferent), I can rely on a known mechanism to turn the preference orderings into an allocation of students to schools. The mechanism is often called the Gale-Shapley mechanism, after Gale and Shapley's description of the mechanism with applications to college admissions and marriage. The mechanism is equally associated with Roth [1984], who demonstrated that the National Intern Matching Program used by hospitals and medical interns since 1953 is an example of the mechanism. Indeed, Roth presents evidence that the mechanism works well in practice. It is good that the mechanism enjoys both familiarity and practical use because policy makers would be averse to exposing children to a mechanism that might produce unexpected outcomes.

The mechanism is simple to describe. Because the purpose of school choice is giving families maximum choice among schools, I describe the version of the mechanism that produces an outcome in which each student is assigned to the school that he prefers the most among his achievable alternative schools. (In the alternative version, each school gets the students that it prefers the most among its achievable alternative students. The alternative version would be appropriate if the spirit of school choice were to maximize schools' choice among students.)

First, examine only the first choice of each student. A few schools may be especially sought after and thus heavily oversubscribed, looking just at the first choices. Most schools will probably be undersubscribed (looking just at the first choices). If school j is oversubscribed, then assign it its N_j most preferred students, where N_j is the total number of students it wants to enroll. If school j ' is

undersubscribed, then assign it all students whose first choice it was. Reset each school's "counter" for the number of school places it has left. A few schools will have no places left; most schools will have some, but fewer than N_j places left; a few schools may have all of their N_j places left.

Some (perhaps many) students will not have been placed in the last step because their first choice school was filled before it got to their name on its preference ordering. Examine the second choice of all unassigned students. Again, some schools may be oversubscribed, looking at students' second choices and their remaining school places. Some schools will probably be undersubscribed. If school j is oversubscribed, then assign its most preferred students to its remaining places. If school j' is undersubscribed, then assign it all students whose second choice it was. Reset each school's "counter" for the number of school places it has left.

Repeat the last step as many times as necessary, continuing at each stage with the next choice (third choice school, fourth choice school, and so on) of students who have not yet been placed. Needless to say, this mechanism would actually be performed by a computer, nearly instantaneously.

The basic properties of the resulting allocation are known. For proofs, see Roth [1984].

A. Students' Choice is Maximized

The allocation is such that each student k is assigned his most preferred attainable school (where a school is attainable if there is a stable allocation in which the student could be assigned to the school). Put another way, every student does at least as well under this allocation mechanism as he could do under any other stable allocation mechanism.

This property is important because it tells us that the mechanism maximizes students' school choice, subject only to the constraint that students must inevitably crowd one another out at schools that are oversubscribed.

Note, as well, that nothing restrains students from making their choices freely. If many bilingual

students want to attend a school with a special language emphasis, there is nothing to prevent all of them from listing it as their first choice. If many students want to attend a school with a core knowledge curriculum, there is nothing to prevent their listing it first. To the extent that students do not all get their first choices, it is purely because they displace one another.

The market-like mechanism guarantees that one category of students can never systemically displace another category. That is, no category of students can be segregated involuntarily. Voluntary self-segregation can exist, however, and may have benefits, as in the case of a school that emphasizes bilingual instruction or the arts.

What will the allocation look like? All schools will tend to have similar shares of students in each category, with any dissimilarities driven by voluntary self-segregation. Because vouchers vary with a student's category, all schools will tend to have similar per-pupil spending—with dissimilarities also driven purely by voluntary self-segregation.

B. Stability

The allocation is stable. That is, there does not exist a student k and a school j such that the student k prefers school j to the school to which he was assigned, and school j prefers student k to one of the students that it actually admitted.

By stable, I do not mean to imply that a student will not reallocate himself (re-join the allocation mechanism for the following year) if his family discovers that the school is not as it purported to be. Indeed, that type of stability would be undesirable. If a school purports to offer an arts curriculum or to hold students to high academic standards and does not fulfil its mission, then it is appropriate that families seek other schools.

Because of the market-like mechanism, the allocation is resistant to another, more interesting, type of instability. Suppose that a school offered a curriculum that was unexpectedly popular with poor

students. Then (both poor and non-poor) students at the school might find themselves with many more poor students than they expected when they ranked the school. If students dislike this situation and want to reallocate themselves, then one of two situations has occurred. First, it may be that poor students impose a simple externality but that schools consistently underpredicted how much it would take to internalize that externality. Thus, the market-like mechanism set too low an add-on for poor students, and the inadequacy of the add-on (though felt in every school) is particularly obvious at the school with many poor students. This problem will presumably be solved in subsequent years, as schools raise their estimate of the add-on necessary to make them indifferent between poor and non-poor students. Thus, the allocation will be relatively stable—because the add-ons remain responsive to revelations about supply and demand.

The second possibility is that poor students actually impose a complex externality. In particular, the externality may be bigger if poor students are in high concentrations. Thus, although the add-on is adequate for a typical school, it is inadequate for a school that offers a curriculum with a peculiar draw for poor students (that is, a school in which poor students choose to self-segregate). There are two ways for the voucher program to respond to such a situation. It can explicitly recognize the complexity of the externality and design an optimal school composition voucher around it (see below). Alternatively, the program may make no explicit response, thereby passively discouraging schools from offering programs in which poor students want to self-segregate and discouraging poor families from giving high ranks to schools in which they are self-segregated. In other words, by being passive, the program makes an implicit statement that society does not wish to pay the cost of internalizing externalities that are more expensive than it deems necessary. That is, society may prefer poor students to impose smaller externalities by spreading themselves more evenly. In short, the program can adjust, at a minimum through passive discouragement of certain behaviors, to initial instabilities generated by complex

externalities.

C. Rewards for and Expansion of Efficient Schools

The allocation is such that each school j is assigned its N_j lowest attainable students (where a student is attainable if there is a stable allocation in which the student could be assigned to the school). If school j is undersubscribed at the end of the allocation process and enrolls only N_j' students, then it is assigned its N_j' lowest attainable students.

This property is very useful because it ensures that rewards differ among more and less efficient schools, while still imposing a “level playing field” on schools. (Also see the section on topping-up, below, for another means by which schools obtain rewards for efficiency.) Schools can get more and less preferred allocations of students *within* each category, even though the voucher adjustment mechanism induces schools to be category-indifferent. A school that gets its more preferred students is rewarded in a myriad of ways relative to a school that get its less preferred students. More preferred students will tend to fit better with the school’s mission, be more engaged, and generally be more pleasant to teach. In other words, getting its more preferred students is the *reward* for a school that is systemically preferred by parents in spite of its (the school’s) getting no more resources than any other school. Such rewards are useful both because they give schools an incentive to raise productive efficiency and because they encourage efficient schools to expand and inefficient schools to contract.

Generating higher productive efficiency should be a key goal for school choice, given American schools’ parlous level of efficiency—compared either to other countries’ schools or compared to their own productive efficiency in the recent past. Schools need incentives to be efficient—that is, to produce the maximum value given the budgets that they spend. If a school is not satisfied with its rewards under the allocation mechanism, it can improve its rewards, but only by improving its pool of attainable students. A school improves its attainable pool by raising its efficiency—that is, by making itself more attractive given

the money it spends—so it ends up higher on students’ preference orderings.

A school that is being rewarded is likely to find expansion attractive. The marginal students whom it would admit if it were to raise N_j would be students relatively high up in its preference ordering. A school that is not being rewarded is far less likely to expand; indeed, it may contract. Its marginal admits tend to be far down on its preference ordering. Notice that, if all schools reach high levels of productive efficiency and are approximately equally preferred by parents, then the allocation mechanism will assign approximately equal enrollees to each. There are two methods by which all schools might reach relatively high and equal productive efficiency. Inefficient schools might improve by offering more attractive programs, hiring better teachers, and reforming poor administrative rules. Alternatively, efficient schools might expand (perhaps to multiple campuses) and inefficient schools might contract.

At this point, it is useful to resolve a common confusion: are not the penalties are “unfair” to bad schools because a bad school’s badness is compounded when it gets less preferred students? On the surface, this question is nonsensical: it is necessarily impossible to penalize inefficiency without taking something away from schools that already have low performance. One could similarly complain that it is “unfair” to take revenue away from firms that are inefficient relative to their competitors. The question is probably not nonsensical, but instead stems from confusion over three possible cases. In case one, a school is inefficient because its administrators are ineffective: they hire low quality teachers given the salaries they pay, make poor curricular decisions, and so on. In this case, the school should be penalized by getting its less preferred students. Its administrators are capable of improvement (by hiring better teachers, say) or getting out of the schooling business altogether. School choice should encourage them to take one of the two actions. In case two, a school is inefficient because there is some circumstance, exogenous to the school, that makes families consistently shy away from the school. For instance, the school might be located in such a dangerous neighborhood that even local families would prefer to send

their children to more distant schools. Strictly speaking, the school is not to blame for the circumstance; but the school does have an attribute that families dislike. School choice will naturally encourage growth of schools that are located in safe neighborhoods and will discourage schools from locating in unsafe neighborhoods. The outcome will be most, perhaps all, students attending schools in safe neighborhoods—though some students, because it is their families' choice, may travel relatively long distances. In case three, the school is again inefficient because of some exogenous circumstance but “society” knows better than families do and wishes to induce them to give higher preference to the school than they would be inclined to do. I take up this case, in which families cannot be counted upon to pursue children's best interests, in the final section of the paper.

D. Revealing True Preferences

The last property of the allocation is that students and schools have little incentive to give false preference orderings. More precisely, the allocation procedure gives no school an incentive to misrepresent its true first choice. It also gives no student an incentive to misrepresent his true first choice. If schools and students do not have much information about the preferences of other schools and students (as seems likely), then they will not be able to construct strategic preference orderings that will allow them to do better than their true preference orderings. Furthermore, in the case in which students and schools do have enough information to strategically misrepresent a choice on their list, the allocation that obtains if they distort a choice in a rational, strategic manner is stable, given their true preferences. That is, if schools and students are highly informed about others' preferences, then it is not a dominant strategy for all of them to state all of their choices truly (except their first choices, for which it is always a dominant strategy to state truly). If, however, they misrepresent their preferences in order to be strategic, then the Nash equilibria which result will be stable, even when they scrutinize the equilibria in the light of their true preferences.

Would a school misrepresent its preferences in the hopes of getting the add-on for some category to be larger than it really needs to be to make the school indifferent between the categories and other categories? First, one school's deceptions would probably not translate into a change in the add-on, particularly if the state sets the threshold such that only 90 or 95 percent of schools have to exhibit category indifference on each category. Second, if the school were able to raise the add-on through deception, it would end up being undercompensated for all of the students it enrolled from other categories (whom it had to have ranked with category indifference if it were to succeed in its deception). It would necessarily be undercompensated because there is an adding up constraint: basic vouchers are the residual after all of the add-ons are subtracted. A school with a powerful self-selective draw for a particular category could perhaps engage in more strategic behavior, because it could count on families to guarantee that it would end up with few students from the categories for which it would be undercompensated. However, such strategic behavior would be fairly obvious and easy for the voucher program to discern. For instance, suppose that hearing impaired students had such a strong preference for a school that specialized in hearing impairment that their rankings would make them end up at the school, even if the school gave them low rankings. Then the school might be tempted to try to raise the add-ons for hearing impairment by putting hearing impaired students very low on its ranking. Such a mismatch (the school is highly ranked by hearing impaired students, but hearing impaired students get low rankings from the school) would imply that there were substantial *diseconomies of scale* in the treatment of hearing impairment and that these diseconomies were wholly unperceived by families of hearing impaired students. My example of hearing impairment is the best case scenario for strategic behavior. Of course, such strategic behavior would like backfire on a school: after getting a particular add-on raised, it might find that the category of students, armed with larger vouchers and thus higher in other schools' orderings, went elsewhere.

In short, schools and students are not very likely to misrepresent their true preferences. Strategic behavior is likely to be limited because it is not possible in the case where the behavior of the many students and schools is hard to predict and it would be fairly obvious to the voucher program in the rare case in which student behavior was easy to predict. Moreover, if anyone should engage in successful strategic behavior, it will not destabilize the allocation.

E. A Simulation

Recall the simulation of the last section, in which category-indifferent lists were obtained for five schools over 500 students. Recall also that the students had only idiosyncratic preferences, based on match quality, for the five schools. No school had higher productive efficiency and was therefore systematically preferred. If one runs the students' and schools' preference orderings through the allocation mechanism described above, 95 percent of the students are placed with their first choice school, 3 percent are placed with their second choice school, and the remaining 2 percent are placed with their third or fourth choice school. Externality-imposing students are spread quite evenly across schools in the final allocation of students, as shown in Table 3.

Table 3: Allocation of Students to Schools

	Allocation if Students were Uniformly Distributed	school 1	school 2	school 3	school 4	school 5
-3 Externality Students	10	8	10	10	12	10
-1 Externality Students	20	17	21	16	25	21
0 Externality Students	70	75	69	74	63	69

Now consider the somewhat more interesting situation where the schools do not have equal productive efficiency. In particular, suppose that the value of school 5 to students is 4 plus their match

quality with school 5 (randomly drawn from the uniform distribution over $[0,2)$). Suppose that the value of school 4 to students is 3 plus their match quality with school 4 (randomly drawn from the uniform distribution over $[0,2)$). And, so on. School 5 is generally preferred to school 4, which is generally preferred to school 3, and so on down the list of schools.

If one runs the students' new preference orderings and the schools' preference orderings through the allocation mechanism, 31 percent of the students are placed with their first choice school, 18 percent are placed with their second choice school, 18 percent are placed with their third choice school, 15 percent are placed with their fourth choice school, and remaining 18 percent are placed with their fifth choice school. Fewer students are placed with their first choice school because their preference rankings are highly correlated, owing to the large differences in productive efficiency among schools. Nevertheless, externality-imposing students are spread quite evenly across schools in the final allocation of students, as shown in Table 4. Table 4 also shows the average value of students in a school, from the school's point of view. School 5 is rewarded with the highest average value students, school 4 with the next highest average value, and so on down the list of schools.

Table 4: Allocation of Students to Schools

	Allocation if Students were Uniformly Distributed	school 1	school 2	school 3	school 4	school 5
-3 Externality Students	10	11	10	11	9	9
-1 Externality Students	20	22	21	17	23	19
0 Externality Students	70	67	69	72	68	72
Average Student Value		4.58	5.91	6.23	6.91	8.06

F. Summary

In summary, the allocation has many desirable properties. The procedure is simple for students and schools. Students' choice is maximized. Through voluntary self-segregation, students can take advantage of targeted curricula or special programs that need economies of scale to be sustainable. Yet, students' voluntary and decentralized choices cannot generate involuntary segregation of a category of students. Schools are rewarded for productive efficiency, and efficient schools find expansion more attractive than inefficient schools do..

VII. Passive Parents

In principle, school law could mandate that all parents submit preference orderings for their children. Indeed, in districts with mandatory intra-district choice (such as Cambridge, Massachusetts), the share of parents who initially resisted making active decisions about preference orderings dwindled after only a few years. Parents quickly became more active once they realized that, by being passive, they were effectively making a choice, but just not one that they necessarily liked. (For instance, some inactive parents appear to believe that their children will automatically be assigned to the closest school if they take no action. Once they find that their children can be arbitrarily assigned an inconvenient school, they realize that active choice is ultimately less trouble than passive choice.)⁸ However, even if one is optimistic about passive parents becoming active, there ought to be safeguards in a voucher system for children whose parents cannot or will not create a preference ordering for them. Consider how one might construct such a safeguard.

Suppose that a child whose parent has given up his right to construct the preference ordering is assigned to a counselor. The counselor should be able to construct a preference ordering, taking into

⁸ See Petronio [1995].

account the child's education needs and locational constraints. On the one hand, counselors have an advantage over parents, in that they may know more about programs and staff at area schools. On the other hand, one worries that counselors could neglect their duty and construct arbitrary rankings for their student charges, so that the children with passive parents could end up in poor matches or, even worse, in schools largely populated by other children with passive parents.

The appropriate response to such a worry is a statistical test of whether counselors rank schools similarly to parents. Counselors' rankings could be tested before being submitted to the allocation mechanism, and a counselor whose rankings differed significantly from those of parents could be asked to justify his choices.

An appropriate test is based on the difference between F_1 and F_2 where F_1 is based on the coefficient of concordance among parents' rankings and F_2 is based on the coefficient of concordance once group of rankers includes not only parents but also one or more counselors. The coefficient of concordance for the parents alone is:

$$W_1 = \frac{12 S_1}{K_1^2 J (J^2 - 1)}$$

where K_1 is the number of parents ranking the schools (that is, the number of students whose rankings are being concorded), J is the number of schools being ranked, and S_1 is a sum of squared deviations, as follows. For each of the J schools, compute the sum of the ranks given by the K_1 rankers. Compute the mean of these J sums-of-ranks. Obtain S_1 , the sum of the squares of the deviations of the J sum-of-ranks around their mean. Note that, if there were perfect agreement among parents, then W_1 would be equal to 1. Repeat the exercise when the pool of rankers is the K_2 , the number of parents plus one or more counselors. Obtain W_2 . In general, coefficients of concordance (W) range from 0 to 1. F_1 and F_2 are

given by

$$F_1 = \frac{(K_1 - 1) W_1}{1 - W_1} \quad F_2 = \frac{(K_2 - 1) W_2}{1 - W_2}$$

and are distributed according to the F distribution with $K-1-(2/J)$ degrees of freedom for the numerator and $(J-1)(K-1-(2/J))$ degrees of freedom for the denominator.⁹

The key element of the test is the sum-of-ranks for each school. Its importance is intuitive. Simply put, one is testing whether adding the counselors' ranks makes for greater deviation from the mean sum-of-ranks. Brook and Upton [1974] and Fligner and Verducci [1988] present refinements of this test, which allow rankers to have characteristics that legitimately influence their rankings (location is the obvious influence in this case).

VIII. Optimal School Composition: Complex Externalities and

Situations in Which Parents Cannot be Trusted to Choose What is Best for Children

So far, I have considered simple externalities—that is, externalities that are a function of a child's presence, but not a function of how his presence interacts with the presence of other children in the school. There are, however, a set of issues in which the presence of students is thought to interact such that there are socially optimal school compositions. A few examples will make this point clear. Consider neighborhood schools. It may be social welfare enhancing to have schools that draw students disproportionately from their localities because such schools are a locus for neighbors to interact and bond so that they support a variety of non-school local public goods, such as neighborhood safety and recreational opportunities for children. If neighborhood schools are important, then the presence of one

⁹ This test is due to Kendall [1948], a classic treatment of rank statistics.

local child in a school is worth more if other local children are there too.

Alternatively, it may be social welfare enhancing to have siblings attend the same school. Having siblings attend the same school minimizes certain parental costs, such as transportation costs, transaction costs, and opportunity costs of participating in school activities. Parents whose children are grouped in a school may be more likely to volunteer or become a trusted repository of information about the school, on whom other parents can rely. Thus, the presence of one child in a school may be worth more if his sibling is there also.

Finally, there may be socially optimal racial, disability, or ability compositions of schools. At least in the past, a popular racial desegregation goal was for all the schools in a district to have the same black-to-white ratio of students as the district as a whole. Although the courts appear to be increasingly wary of such goals, it may be presumed that matching the district's black-to-white ratio was considered socially optimal school composition. Today, many districts have the goal of mainstreaming disabled students to the maximum extent possible. In other words, many districts apparently believe that the optimal school composition has disabled students spread as evenly as possible across classrooms. It is not the purpose of this paper to take a stand on optimal racial composition, optimal mainstreaming, or any other type of optimal school composition. I merely wish to point out that some current school policies already amount to an implicit claim about optimal school composition. From now on, I take for granted that some districts, through a judicial or political process, have resolved on what optimal school compositions are. Can the voucher procedure be adapted to take such resolutions into account?

A. The Need for Optimal Composition Vouchers

First, it is not clear that the voucher procedure needs to be adapted. To the extent that optimal school compositions are generated by voluntary self-segregation (or the lack thereof), parents' orderings have already taken optimal school composition into account. Consider, for instance, neighborhood

schools. Families themselves face many of the costs of sending their children to distant schools. Even if school bus transportation is entirely free, children must spend time in transit and parents must travel further to get to school events. Moreover, families directly receive many of the benefits associated with neighborhood schools. Thus, it is far from obvious that families' preference orderings will be much different than the preference orderings they would form if they were to internalize all of the social benefits of sending their children to neighborhood schools. Similarly, families are directly hit by most of the costs of dispersing siblings and directly receive most of the benefits of grouping them. Families' preference orderings may, therefore, be close to what is socially optimal.

Moreover, the voucher adjustment mechanism makes involuntary segregation of a category of students impossible. Thus, white families cannot, simply through careful construction of *their* preference orderings, keep their children's schools relatively free of black children. Families of regular education students cannot construct their preference orderings so as to keep their children out of schools in which disabled children are mainstreamed. Schools will differ in racial composition (mainstreaming of disabled children) only if families with black (disabled) children *elect* self-segregation by constructing preference orderings that differ systematically from those of white (regular education) families.

In short, one ought *not* to adapt the voucher system if one believes that (1) families can calculate the costs and benefits of having their children attend schools with various enrollment compositions and (2) *among families with protected category students*, families' private costs and benefits are similar to the social costs and benefits.

B. Optimal Composition Add-Ons

There are two cases to consider with regard to optimal school composition. In Case One, the policy maker has an optimal school composition that he wants every school to satisfy equally. Examples would be each school having a black/white ratio and disabled/non-disabled ratio equal to those of the

district. Policies that are equally desirable for all schools need not be only about the composition of enrollment. Case One also covers policies such as: it is socially good for schools to locate in certain neighborhoods that parents dislike (case three from the discussion above), all schools should administer standardized tests, all schools should publish school report cards, all schools should offer algebra by the seventh grade.

Suppose that the district makes the basic voucher small (perhaps zero) and gives a *school-specific* add-on that depends on the degree to which the school satisfies the district's policy resolutions. For instance, the add-on could vary (continuously) with the inverse of the absolute value of the difference between the school's percent white and the district's percent white. The more punitive is the add-on with respect to deviations from the percent white goal, the more families' preference orderings will be driven by the school-specific variation in the voucher. In the extreme, all families' preference orderings will be the same, with the school rankings perfectly correlated with the school-specific vouchers. In this extreme, all schools will achieve approximately the district's black-white ratio. The universally top-ranked school will take its N_j most preferred students, who necessarily approximate the district's black-white ratio because schools' rankings are required to be category indifferent. The next most highly ranked school will then take its N_j most preferred students out of the pool that remains. They, again, will approximate the district's black-white ratio. And so on. Thus, all schools will tend to achieve the district's goal.

In this extreme case, although vouchers exist, parental choice is not meaningful. Parents' choices are dictated by the district's goals, and schools are rewarded only for fulfilling the district's goals, not for achieving higher efficiency in parents' eyes. That is, although vouchers can be used to achieve social goals that are odds with parents' goals, there is a trade-off between advancing such goals and being able to exploit the information and unconflicted motives embodied in the parents. Of course, it is only in the extreme case that parents have no meaningful choice. More generally, districts can use add-ons that

have the structure mentioned (the add-on for which a school qualifies rises as the school approaches the goal), but that fall short of being dictatorial. That is, parents can be given some, but not overwhelming, incentives to make choices according to social preferences, rather than their own preferences. A district might have one add-on for schools' approaching the desired racial composition, another for approaching the desired mainstreaming of disabled children, and yet another for administering and publicizing standardized exams designed to test compliance with the district's curriculum.

Case Two is optimal school composition that varies with the school, such as a preference that each school enroll its local children or a preference that siblings attend the same school. For case two, the district need only create add-ons that are student-school-specific. For instance, a family could get a sibling add-on for sending a second child to a school or get a neighborhood add-on that would vary inversely with family's distance from a school. Such add-ons would induce siblings (neighbors) to have similar preference orderings—in the extreme, identical preference orderings. *Within-family* or *within-neighborhood*, one is back in Case One.

C. Cautious Use of Optimal Composition Add-ons

I have already emphasized that optimal composition add-ons are not necessary if the families of students in protected categories enjoy costs and benefits of school composition that are similar to social costs and benefits. Put another way, *many problems that initially appear to be optimal composition problems really involve only simple externalities, and do not require optimal composition add-ons.*

Moreover, there are several reasons to be cautious about introducing optimal composition add-ons. First, creating add-ons for optimal composition problems is fundamentally different than creating them for simple externalities. In the case of simple externalities, we can rely on market-like mechanisms (Section III) and the housing market (Section IV) for evidence on the value of the externalities. Problems of socially optimal composition do not arise, however, unless families' private calculations differ from

those of society. Thus, markets cannot generate evidence about the value of various compositions. While one might propose grand social experiments that would reveal the value of school various compositions, such experimentation is unlikely (and has no legal basis, if it is not strictly illegal).

We might turn to non-experimental evidence on the value of various school compositions, but even the best such evidence provides little guidance. There is paucity of credible empirical evidence about optimal school composition because studies are plagued by self-selection. Families self-select into schools, and it is difficult to disentangle the effects of group composition from the effects of the factors that lead families to group together in the first place. Non-experimental studies that credibly eliminate self-selection biases necessarily rely on natural experiments. Almost by definition, natural experiments are rare and not controlled by evaluators. Thus, many interesting school compositions are never tested by natural experiments.¹⁰ Also, even the best non-experimental estimates have standard errors that are so large that one would hesitate to conclude that optimal group composition could not be achieved by parents' choices, once vouchers were adjusted to account for simple externalities.

Finally, although the availability of optimal school compositions add-ons could tempt policy makers to undertake ambitious social interventions, they must remember that they are ultimately constrained by parents' willingness to remain within the current system. School composition add-ons attempt to move parents away from the choices that they (parents) consider optimal. At some point, parents will move to less interventionist districts or use private schools with no voucher, rather than use voucher add-ons and accept the conditions attached to them.

In short, vouchers add-ons can be used to accommodate resolutions about socially optimal school

¹⁰ For recent evidence on the effects of group composition, see Zimmerman [1999], Sacerdote [2000], Hoxby [2000], and Katz, Kling, and Liebman [2000]. All of the papers listed go to considerable lengths to remedy the selection problem, but none of the papers could be said to identify socially optimal group composition.

composition. The structure that the add-ons should take is clear: a school-specific (student-school-specific) add-on that varies continuously with the inverse of the distance between the school (student-school combination) and the goal. However, there are a variety of reasons why such add-ons may not be as necessary as might first appear and why it is difficult to size such add-ons correctly.

IX. Taking Vouchers Across District Lines

Thus far, I have assumed that vouchers are financed by individual school districts, and I have set aside questions like whether vouchers can be taken across district lines and/or topped up. A related issue that I have set aside is whether vouchers should be financed at the state, rather than the district, level. In a few states, some of these issues do not arise because all (or most) school finance is conducted at the state level and the distinction among districts is therefore unimportant. These states include Hawaii, California, New Mexico, and Michigan. In most states, however, local and state governments share responsibility for financing schools.

State governments typically bear the major responsibility for financing programs for externality-imposing categories of students: the disabled, those with limited English proficiency, the poor, and minorities. Delegating such responsibilities to the state is a straightforward application of basic fiscal federalist reasoning: districts have difficulty financing programs for externality-imposing students because a generous program is likely to draw additional externality-imposing students into the district. The resulting incentives could generate a “race to the bottom,” in which every district is unduly stingy to externality-imposing students so that it gets no more than its fair share of them.

Local districts typically bear the remaining responsibility for setting per-pupil spending. Again, this responsibility is justified by basic local public finance reasoning. District finance results in increased allocative efficiency to the extent that families, by choosing among districts, can choose the level of

human capital investment that is most appropriate for their child. District finance allows variation in per pupil spending (and thus human capital investment) among families, but it also provides a means by which each family can spread the burden of its investments over its lifetime, rather than having them entirely concentrated in the period in which its children attend schools. Problems in allocative efficiency do arise with district finance—for instance, there is a missing capital market for poor families who want to finance optimal human investments in their children. Such allocative efficiency problems are best solved by policies, such as school finance equalization, that are effectively behind the scenes in this paper. Such policies are automatically incorporated in a district’s basic voucher.

Consider the voucher procedure described above. How should one finance the vouchers--at the state or local level? And, should students should be allowed to carry vouchers across district lines or top them up? In addition to wishing to incorporate known results from fiscal federalism and local public finance, I wish to describe a voucher procedure that causes as little disruption as is possible. In particular, current property values reflect districts’ current financing of local public schools. A voucher scheme that shocked school finance would shock the property market.

It is fairly obvious that add-ons for externality-imposing categories of students should be determined and financed at the state level. Making the state responsible for externality-internalizing add-ons would minimize disruption by keeping categorical programs at the state level. Moreover, if districts were to determine the add-ons, each district would have an incentive to create programs that served only the most marginal externality-imposing students—for instance, a program for limited English proficient students that worked only for students who already had some English skills. The families who would want to live in the district would naturally have children who were only marginally eligible for the externality-imposing categories. As a result, the district could “legitimately” claim that its market-like mechanism required only very small add-ons to make schools category-indifferent. Families whose children had more

serious needs would be deterred from moving to the district because the add-on for their child would fall.

In short, a race to the bottom (in program intensity) would be an equilibrium.

States are also the logical level for financing add-ons that encourage schools to use the state's curriculum, testing, *et cetera*. If such programs are socially beneficial, their benefits are presumably internalized at the level of the state, which is therefore in the best position to calculate their value.

What about the basic voucher, however? Except in states with state-level school finance, district-level finance of the basic voucher would minimize disruption. District-level financing of the basic voucher would also allow human capital investments to vary among students, increasing allocative efficiency. (Note that we are *not* considering how human capital investments might vary among protected student categories, only within categories.) However, district-financed vouchers require topping up if they are to significantly increase parental choice without creating a race to the bottom in basic vouchers.

Existing inter-district choice programs that do not require topping up, such as that of Massachusetts, can clarify these issues. In such plans, schools are asked to accept out-of-district students who bring with them per-pupil spending that is a function of their home district's spending. Such programs give families an incentive to move to districts with low per-pupil spending even if they intend to send their children to schools in districts with high per-pupil spending. Indeed, some inter-district choice programs did begin races to the bottom in local per pupil spending. These races were halted only when districts were able to decline participating (so that inter-district choice effectively ended).¹¹

Suppose that students are allowed to carry vouchers across district lines, so long as their own families top up the voucher so that it equals the voucher in their destination district. Giving families this option expands their choices and may be particularly important for families in which different siblings need different human capital investments. Yet, this option does not allow a family to underpay by living in a

¹¹ See Armor [1997].

district that finances a basic voucher that is lower than the regular tuition of the schools they intend to use. In fact, because the “topping-up” payments are due when the child is in school (rather than being spread over the parents’ taxpaying life), families still have an incentive to live in a district that finances basic vouchers close to the regular tuition of the schools they intend to use. A family loses the consumption smoothing properties of a district-financed voucher if it attempts to underpay by living in a district with a small basic voucher.

X. Schools that Require Topping-Up

So far, I have focused on schools that accept the vouchers of the district in which they are located, without requiring topping up. This is the condition under which publicly financed voucher schools and charter schools now operate. Should one allow vouchers to be used at schools that *require* topping up? An advantage of allowing this possibility is that schools can contemplate significant rewards to becoming very desirable. Only a very efficient school with a educational program that deserved expansion could expect demand for its education to remain healthy if it required topping up. Moreover, the topping up might enable (and give incentives for) expansion.

There is, however, a potential disadvantage to allowing schools to charge regular tuition in excess of the basic voucher. In a school that required topping up, the additional fee could be set so that the school was given a low ranking by all except affluent families. Although the school’s own ranking of students might be category-indifferent, families’ ranking of the school would be so driven by their ability to pay the additional fee that the school’s actual enrollment might end up being an exclusive group that included few or no externality-imposing students.

On the one hand, this disadvantageous situation cannot be inherently wrong because it is the natural extension of district-financed vouchers. Simply consider a district that supports only one school

and that has one school-aged child in each household. In such a district, financing a higher basic voucher is no different than setting a higher topping up amount. In practice, however, there are reasons to distinguish between a school that accepts a district's voucher and a school in which every student must top up the voucher. First, districts in metropolitan areas are generally significantly larger than a single private school, and they therefore necessarily contain a greater variety of students than a single school's applicant group would have to contain. Second, families pay for the voucher over their entire taxpaying lives, whereas the topping up payments fall in a short period of time, when parents are generally too young to be in their peak earnings years. Thus, by setting a high topping up payment, a school can easily discourage low income families from applying (compared to district's ability to use a high basic voucher amount to discourage low income families from moving in). Third, one reason for district-financed vouchers is continuity with the current system: this argument cannot be used in favor of schools that require significant topping up from every applicant.

If one wants to allow exist voucher schools to exist that require topping up, it is reasonable to install a safeguard that ensures that the topping up amount is not used primarily to make the school inaccessible to protected categories of students, especially poor students. A simple safeguard would be following: a school could accept vouchers *and* require topping up only if its enrollment, not just its ranking, exhibited category indifference when tested against its district's population. In order to pass this test, a school that required significant topping up might have to offer scholarships (that is, waive the additional fee) to poor students who would otherwise give the school a low ranking in order to avoid having to top up the voucher.

XI. Conclusions

In this paper, I describe a procedure for constructing vouchers that—on the one hand—greatly

increase parents' choice and the competition faced by schools, and—on the other hand—remedy several externality-related problems as well as they are remedied by current school law. In fact, the voucher procedure described is likely to remedy certain problems better than current law does because it explicitly internalizes externalities through a mechanism that retains the flexibility of market-set prices. There is even a natural way for misperceived externalities to come correctly “priced” over time. For instance, if misperceptions generate add-ons for minority students are initially more than is really necessary to internalize the externalities, schools will eventually figure this out and—in trying to arbitrage the add-ons—bring down the add-on for minority students.

The procedure described in this paper is practical and accommodates many of the school laws and institutional details of the current system. Parents and schools both have relatively simple tasks. Parents need only rank schools according to their preferences (although it is always important for school choice that they take this simple task seriously). Schools must be willing to re-rank students a few times, especially in the first year of the program. This ranking need not be elaborate, however, as the state will only be checking the rankings for schools' indifference among the *categories* of students. Once the market-like mechanism generates externality-internalizing vouchers, the procedure relies on a tried allocation mechanism that maximizes student choice among schools. I suggest how the procedure may be modified to provide for circumstances like passive parents, optimal school composition, incentives for schools that follow state guidelines, and the need for “topping up” so that vouchers can cross district lines.

References

- Armor, David J., and Brett M. Peiser, *Competition in Education : A Case Study of Interdistrict Choice*. Boston : Pioneer Institute, 1997.
- Black, Sandra E., "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics*, Vol. 114, Number 2 (May 1999), pp. 577-600.
- Brook, D., and G.J.G. Upton, "Biases in Local Government Elections Due to Position on the Ballot Paper," *Applied Statistics*, 23 (1974), 414-419.
- Cullen, Julie B, "The Impact of Fiscal Incentives on Student Disability Rates," NBER Working Paper No. 7173, June 1999.
- Epple, Dennis, R. Filimon, and T. Romer, "Equilibrium among Local Jurisdictions: Toward an Integrated Treatment of Voting and Residential Choice," *Journal of Public Economics*, 17 (1984), 281-308.
- Fligner, M.A., and J.S. Verducci, "A Nonparametric Test for Judges' Bias in an Athletic Competition," *Applied Statistics*, 37.1 (1988), 101-110.
- Hanushek, Eric A. *Making Schools Work*. Washington, DC: The Brookings Institution, 1994.
- Hoxby, Caroline M., "All School Finance Equalizations Are Create Equal," *Quarterly Journal of Economics*, forthcoming.
- Hoxby, Caroline M., "Peer Effects in the Classroom: Learning from Gender and Race Variation," NBER Working Paper 7867, 2000.
- Hoxby, Caroline M., "School Choice and School Productivity (Or, Could School Choice be a Rising Tide that Lifts All Boats?," in Caroline M. Hoxby, ed., *The Economics of School Choice*. Chicago: University of Chicago Press, forthcoming.
- Kendall, Maurice G., *Rank Correlation Methods*. London: Charles Griffin and Company, 1948.
- Katz, Lawrence, Jeffrey Kling, and Jeffrey Liebman, "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment," NBER Working Paper 7973, 2000.
- Ladd, Helen F., Rosemary Chalk, and Janet S. Hansen, eds. *Equity and Adequacy in Education Finance: Issues and Perspectives*. Washington, DC: National Research Council, 1999.
- Orfield, Gary, and S. Eaton, *Dismantling Desegregation: The Quiet Repeal of Brown v. Board of Education*. New York: New Press, 1996.
- Petronio, Maureen Allenberg, "Kindergarten Shopping: Parents' Response to a Controlled School Choice

- Plan,” Harvard University Graduate School of Education thesis, 1995.
- Randles, Ronald H., and Douglas A. Wolfe, *Introduction to The Theory of Nonparametric Statistics*. Malabar, FL: Krieger Publishing Company, 1991.
- Reber, Sarah, “Did Court-Ordered Desegregation Plans Benefit Black Students,” Harvard University typescript, 2001.
- Sacerdote, Bruce, “Peer Effects with Random Assignment: Results for Dartmouth Roommates,” NBER Working Paper 7469, 2000.
- Tiebout, Charles, “A Pure Theory of Local Public Expenditures,” *Journal of Political Economy*, 64 (1956), 416-24.
- Zimmerman, David, “Peer Effects on Outcomes: Evidence from a Natural Experiment,” Williams College typescript, 1999.

Appendix

Let a household's ratio of school-aged-child-years to living-in-district-years be a variable denoted as φ . Note that φ is a function of both the household's number of children and its expected stay in the district.

Let φ_j be the average value of the variable φ for households in district j . If the average household expects to experience school-related externalities half of the years it lives it in the area, say, then an externality that it values at \$50 per year must be worth \$100 per school year *experienced*. Thus, equation (3) should be modified as follows:

$$(3a) \quad [P_y \cdot (r + \delta)] + P_y \tau_j = \alpha_0 + \alpha_1 \frac{PPS_j}{\varphi_j} + \alpha_2 \frac{\%SpEd1_j}{\varphi_j} + \alpha_3 \frac{\%SpEd2_j}{\varphi_j} + \dots + \alpha_k Sqft_y + \alpha_{k+1} Acres_y + \dots + \alpha_{k+\pi} DistanceCity_y + \alpha_{k+\pi+1} DistancePark_y + \dots + \epsilon_y,$$

and equation (4) should be modified as follows:

$$(4a) \quad [P_y \cdot (r + \delta)] + P_y \tau_j = \beta_0 + \beta_1 \frac{PPS_j}{\varphi_j} + \beta_2 \frac{SpEd1_j}{\varphi_j} + \beta_3 \frac{SpEd2_j}{\varphi_j} + \dots + \beta_k Sqft_y + \beta_{k+1} Acres_y + \dots + I^{boundary} \gamma + \nu_y, \quad \forall j \in \{boundary 1, \dots, boundary N\}.$$