

Introduction to *The Economic Analysis of School Choice*

Caroline M. Hoxby

I recall attending a conference on education in Washington D.C. in 1994 that was attended by leading policy makers and expert researchers. At that time, the topic of school choice occupied only a small share of the discussion, and I was the only economist presenting research related to choice. In fact, although economist Milton Friedman is generally credited with spurring modern interest in school choice,¹ economists were contributing relatively little to the school choice debate at the time. The two practical choice proposals that were best known were authored by, respectively, a sociologist, Christopher Jencks, and two legal scholars, John Coons and Stephen Sugarman.² A few programs with choice features had recently been enacted (vouchers in Milwaukee, Minnesota's open enrollment plan, intra-district choice in Cambridge, and so on), but these programs had been initiated by politicians and courts with at least as much of an eye to politics (especially racial politics) as to school improvement. Moreover, *analysis* of school choice was largely out of the hands of economists. If a policy-maker asked for research on choice, he was likely to be referred to work by a political scientist—for instance, John Witte's comparison of Milwaukee's voucher students to Milwaukee public school students;³ studies of the short-lived Alum Rock choice program authored by Rand researchers (none of whom was an economist);⁴ or legal scholarship on church-state issues.

In the several years since 1994, the economic analysis of school choice has burgeoned. As economists have worked on choice, their areas of ignorance and confusion have narrowed enormously, and their discussion has become much more coherent. Perhaps this will be a surprise to the outside observer, who assumes that economists always disagree. However, the

rapid growth in economists' understanding of school choice should be no surprise, really. At its core, school choice relies on very basic economic theory about the effects of competition. Moreover, the tools needed for analyzing the more complex aspects of choice were at hand in 1994. They had been built for related economic problems and could be applied readily to school choice.

It would be optimistic to suggest that economists fully understand school choice and agree about all its intricacies. Nevertheless, there is now a consensus about what we know (and do not know) and about the sorts of evidence and analysis that we need in order to resolve uncertainties. This volume is a testament to the rapid growth of that consensus and to the richness of the economic analysis of school choice. Not only do the authors represent a good share of the economists who have written on choice, but each chapter was subjected to expert critique by other economists and authorities who work on the topic: Joseph Altonji, John Chubb, Chester Finn, Jane Hannaway, Thomas Kane, Helen Ladd, Charles Manski, Richard Murnane, Derek Neal, and Ananth Seshadri. The authors are very grateful for their comments and wish to acknowledge how they have shaped the book.

Why the *Economic Analysis of School Choice*?

What does it mean to perform an "Economic Analysis of School Choice?" It does *not* mean that the authors in this volume are interested only in the financial aspects of school choice. The authors are deeply interested in (and analyze) many non-financial aspects of choice, including student achievement, parental satisfaction, school segregation, mainstreaming of disabled children, and parents' choosing where to live. What it *does* mean is that the authors rely on methods that were originally developed for the purpose of economic analysis. (In fact, when I

refer to “economists,” I refer to people who practice such methods—thereby including some people who are not card-carrying economists.)

Perhaps this statement will leave readers still in doubt. Why, they will ask, should we think that economists are “naturals” when it comes to school choice? After all, learning the institutional details of elementary and secondary schools is not part of the typical economist’s training, and anyone who wants to make important contributions on school choice had better know what actually goes on in schools. There are also legal issues associated with school choice, and economists are generally not experts in school law. Economists *do* have some learning to do when they take up the topic. Nevertheless, economists are naturals in several other, arguably much more important, ways.

An Unabashed Apologia for the Economic Analysis of School Choice

I would argue that there is a simple reason why economists have and will continue make so much progress on school choice: tools. In a typical discussion of school choice, a variety of questions arise. Answering these questions generally requires the use of some analytic tools that are not necessarily complicated but that do require practice if they are to be used effectively. I will elaborate on these tools below. For now, all we need to know is that, when the school choice problem fell into the laps of policy analysts, economists were unusually well-equipped to start answering the questions that arose. As a consequence, economists quickly got immersed in practical tasks, unraveling questions about choice and refining economic methods so that they could applied to schools.

In contrast, many other commentators on education found that they could not make much headway against the questions that arose in a typical debate on choice. Discussions often ended

with the participants being more confused than they were initially. As a result, many commentators abandoned the idea of trying to find analytic answers to questions about choice and decided instead that it was essentially a matter of principle. Unfortunately, when commentators view school choice purely as a matter of principle, their positions (of support or opposition) tend to become hardened. Consensus is unlikely to grow.

Let us consider some of the tools—or, more properly, areas of familiarity—that economists bring to the analysis of school choice. This exercise is not merely a justification of economic analysis. It is way to see where the confusions arise in school choice debates, and how it is that research (like that contained in this volume) clarifies them.

1. Market structure makes the difference.

If school choice makes a difference (good or bad), it will be because it changes the *structure* of the market for K-12 education. When one says that school choice affects market structure, one means that it affects basic constraints that schools and students face. For instance, choice make it easier for students to be mobile among schools, and choice often makes a school's revenue depend directly on its attracting students.

There are two reasons why it is important that an analyst of school choice be comfortable with markets. First, we need to understand how market structure affects how market participants (schools and students) behave and how—in turn—their behavior affects outcomes (achievement, school productivity, and so on). This relationship is often summarized in a phrase familiar to every economist: “structure, conduct, performance.” Understanding this relationship is important because we *can* usually describe how a school choice program affects market structure. If an analyst is good at reducing a program to its effects on market structure, and knows how to

predict the results of that structure, he can make significant progress.

People who are not able to reduce a school choice plan to its effects on market structure tend to get distracted by its superficial details—the transportation plan, the school buildings currently in use, and so on. They do not distinguish between local idiosyncracies and phenomena that are *systemically* affected by the choice program, and—therefore—their analysis gets bogged down.

The second reason why it is important that an analyst of school choice be comfortable with markets is that he must, at a minimum, be *open* to the idea that market forces matter—that is, that people may alter their behavior in response to the pressures and incentives that the market generates. Economists are open to this idea. Although economists may not share the same prior beliefs about the *degree* to which people respond to market forces, they do share the belief that the degree of response ought to be measured. Many non-economists in the educational sector *assume* that market forces do not affect educators or students. Sometimes, this assumption is a matter of principle: it is ignoble to describe market forces in education, let alone measure them. Sometimes, this assumption stems from a belief that only for-profit firms respond to market forces. Economists have a long acquaintance with governments, individuals, and non-profit organizations responding to market forces, and therefore do not dismiss the task of measuring their responses.

2. *It helps to call a spade a spade.*

In some localities, the idea of school choice is more popular with the public than it is with interest groups in the education sector. As a result, one often sees programs that include “choice” in their title but that contain few elements that are recognizable as choice. Even more

confusingly, many Americans have grown accustomed to the idea that traditional public schools are wholly *public* (in the sense of being equally open to all people) when, in fact, the traditional system contains some strong market elements. Discrepancies between the nomenclature and the reality confuse many would-be analysts of school choice.

Because economists focus on how school choice programs affect market structure, they do not get hung up on the names of programs. They know that when no money follows a student, they should expect different outcomes than when money does follow a student. They know that a program in which schools are not allowed to contract or expand, to enter or exit, is different from one in which the suppliers of schooling are elastic. They know to look at the constraints that a school faces, rather than whether it is called “public,” “private,” “charter,” “community,” “magnet,” or something else. In short, economists eliminate myriad sources of confusion by knowing how to extract the market structure from the description of a school system (whether or not its name includes the word “choice”).

3. One cannot avoid the interdependence of school choice and school finance, so one might as well enjoy it.

Every school choice program contains provisions about money. For instance, voucher amounts must be set in some fashion and must be funded by some stream of revenue. Charter schools receive a per-student fee that must be related in some fashion to local per-pupil spending.

Thus, school choice inevitably intersects with school finance, which is the study of (a) how school districts spend and raise property taxes and other sources of revenue, (b) how state and federal aid affects school districts’ revenues and expenditures, and (c) the relationship between property tax rates and property tax revenue. This intersection is dreaded by many

education policy makers, whose eyes glaze over at the thought of learning more about taxes than they need to know in order to avoid being audited. As a result, financing is sometimes only an afterthought (and often a poorly designed afterthought) in school choice plans.

In contrast, economists are not only not repelled by the school finance issues implicit in school choice, they are actively drawn to the intersection between school choice and school finance. The intersection interests economists because school choice makes it possible to fund students at an *individual* level and to fund schools *flexibly*. That is, choice expands the set of financial instruments that are available to fulfil the goals of school finance. Nowhere has economic analysis been more productive than at the intersection of school finance and school choice.

4. *School choice is more interesting to people who are puzzled by the inefficacy of school inputs.*

The “man on the street” thinks that it is obvious that a school with more resources—higher per-pupil spending or smaller classes, say—will produce higher achievement. In fact, an examination of school data shows that this is far from obvious. There are literally thousands of economic studies that attempt to estimate “education production functions,”—that is, the relationship between school inputs and outputs (student achievement). For at least the past 30 years, since the 1966 publication of the influential Coleman Report,⁵ these attempts have focused on the question of whether there is *any* relationship at all. The measured productivity of school inputs is so low, even by the most optimistic estimates, that it would greatly shock the man on the street. For instance, the *most optimistic* widely accepted estimates of the effect of class size reduction suggest that lowering class size by 10 percent (approximately 2 students) for all the years that a student is in elementary school raises his achievement by 0.17 of a standard

deviation.⁶ Because the man on the street is ignorant of such facts, he remains unpuzzled. In contrast, economists are impatient to solve the puzzle of why money does not matter more in schools.

The lack of market forces in education is one of the most promising potential explanations of the puzzle that has yet been put forward. After all, market pressures are generally credited with stimulating firms to be productive. Thus, it is natural that economists are interested in the productivity consequences of choice: they know that there is a puzzle to be solved and they know that market pressures are a potential solution that is worth understanding.

5. You cannot predict the effects of school choice on student sorting without the tools of general equilibrium.

The most complicated effects of choice are on student sorting—how students will allocate themselves among schools when allowed to choose schools more freely. In popular parlance, such issues are described as “cream-skimming” or segregation, even though these are only two forms of student sorting that could arise. Debates about choice often run aground on such issues, as opponents and proponents find themselves getting confused about how choice would affect student sorting. Even if debaters do not admit to being confused, they find it hard to explain the logic behind their assertions that a certain type of cream-skimming or segregation will occur.

It is natural that confusion occurs. In fact, it is impossible to predict the effects of choice on student sorting without:

- (a) knowing numerous parameters about households and schools, such as how a child’s achievement is affected by his peers and what the efficient scale is for producing education with various types of children—such as disabled, limited English proficient, or

- gifted students;
- (b) having an accurate characterization of *current* student sorting, which is very strong in most metropolitan areas;
 - (c) having tools for finding equilibria in which students will be allocated to schools in a stable way. This is a general equilibrium problem that requires simultaneously solving for three equilibria: equilibrium in the market for schooling, equilibrium in the market for housing, and equilibrium in the labor market (solving for the income distribution). In practice, a combination of closed-form proofs, simulations, and computable general equilibrium techniques are required.

It goes without saying that the typical commentator on school choice lacks all of these requirements for understanding student sorting. Moreover, he often does not realize that he lacks them, and thus does not even try to acquire what he needs in order to answer questions about student sorting.

Economists attempt to fulfil requirement (a) by making a *variety* of assumptions about the parameters and obtaining a corresponding *variety* of predictions about student sorting. Empirical economists can supply some of the information for requirement (b), although other the information remains obscure because schools do not keep much data on their students' backgrounds. Economic theorists are girded with the general equilibrium tools listed in requirement (c), although these tools are often seriously strained by the complexities of the student sorting problem.

It would be optimistic to say that economists will soon be able to predict the student sorting consequences of any proposed school choice plan. Our theoretical machinery is too crude

to incorporate the actual complexity of households, schools, and housing markets, especially because the information fed into the machinery falls far short of what is needed. Nevertheless, compared to the typical commentator on school choice, economists have made significant strides on the student sorting problem: they grasp the structure of the problem, they understand what information they need in order to proceed, and they have some idea of how *various* the plausible outcomes are. Economists can at least maintain a proper sense of humility about predicting how school choice will affect student sorting.

6. School choice will affect labor markets for educators.

In a discussion of school choice, it is common to hear one of the following protests: “Teachers in my district cannot be paid that way. It is not in their contract.” or “Teachers cannot be paid that much: a lot of the budget has to be spent on other staff.” or “Voters in my area will not support teacher salaries that are that high.” or “Administrators in my district do not get to make that kind of hiring (firing) decision. Their decisions are constrained by the union.” or “Good teachers in my area leave the profession quickly. They say that there is no appreciation of their skills (room for advancement).” Such protests often bring a discussion of school choice to a halt, as participants shake their heads over the labor market for educators.

All such protests are excessively rigid when school choice is the issue. The market for educators is an upstream market (a provider of inputs) for the market for schooling. Anything that fundamentally changes the structure of the market for schooling, as school choice can, can deeply affect the way its upstream markets function. Indeed, the typical protests *reflect* the structure of the current market for schooling; they are not preconditions to which it is subject. For example, the salaries that parents are willing to pay reflect their satisfaction with schools, and choice

generally gives schools stronger incentives to satisfy parents.

Analyzing how a change in market structure affects upstream markets is a classic problem in economics. Thus, not only are economists not stymied by protests like those listed above; they can borrow experience from other sectors in which upstream labor markets changed in response to downstream market structure. Such experience is not derived exclusively from for-profit industries. Much recent experience comes from industries that include not-for-profit and government providers, like health care.

7. Evidence on school choice requires the latest methods in non-experimental empirical analysis.

As a rule, it is impossible to conduct controlled, double-blind experiments on schoolchildren, primarily because it is considered unethical to experiment on them. That is, if the policy under consideration is considered likely to be beneficial, then it is considered unethical to allow some children to experience the policy while forbidding it to others. Furthermore, controlled, double-blind experiments are impractical in many cases. It would be impossible, for instance, to let some families to exercise school choice and forbid other families to exercise it, but to keep them blind about the group to which they had been assigned (or even to keep schools ignorant of the families' group assignments). In addition, if we were to allow only some families to exercise school choice (keeping others in the "control" group), we could not observe some of the effects of a full-blown school choice program. For instance, the labor market for educators might not change much if educators could easily avoid being in a choice school (by working at a control school).

The lack of a laboratory-like experimental setting is nothing new to economists, given

that they cannot experiment on workers, the unemployed, trainees, or any number of other people affected by economic policies. Particularly over the last 30 years, economists have worked intensely to develop statistical techniques for analyzing policies in non-experimental settings. Empirical economists now have methods for exploiting policy enactments (“policy experiments”) and accidental policy changes (“natural experiments”). These methods work because they extract from the real world those events that most closely mimic the controlled setting of a laboratory experiment. Economists are acutely aware of problems like selection (the families who use a school choice program may be different from those who do not), policy endogeneity (school choice policies may be enacted in response to problems that continue to affect students after enactment), and the inadequacies of partial equilibrium analysis (the effect of a choice plan in which only a fraction of students can participate may differ from the effect of one in which all students can participate).

Given the rapid improvement in these techniques during the 1980s and early 1990s, the school choice debate could not have arisen at a better time for empirical economists. Put another way, it is a quirk of timing that has made economists (and others using their methods) the best empirical analysts of school choice: when the need for analysis arose, they were simply the most experienced users of the right tools for the job.

8. It is important to know which students are likely to be affected by school choice.

The current make-up of private schools is a source of considerable confusion in school choice debates. Children who currently attend private schools belong to several, diverse groups: central city children who attend inexpensive schools that are charitably subsidized by a denomination (such as the Roman Catholic church) to which most of the children do not belong,

suburban children who attend private schools that are affiliated with their own religious group, disabled children who attend private schools that cater to their special needs, and children of affluent parents who attend college preparatory schools. The first two groups account for approximately 85 percent of private school students in the United States.

There are two things to note about all of these students. First, they are *by definition* not constrained to attend the public school to which they would otherwise be assigned. Second, they are unusual. Central city private school students are unusual because they are children who have had the good fortune to land one of the highly rationed places in a subsidized non-public school. Places in such schools are not rationed on tuition (that is, such schools would be glad to offer more school places if the marginal students could pay tuition equal to the per-pupil costs). Instead, such schools maintain low tuition in order to remain accessible to poor families, and they distribute their limited places on relatively arbitrary bases, such as “first-come, first-serve.” The second group of students is unusual because they come from families who place so much weight on religious education that they are willing to make considerable financial sacrifices for it. Not only do they pay tuition on top of paying taxes for public schools that their children do not use, their private schools also spend only about half of what their public schools spend (and their children consequently experience fewer amenities). The third group—disabled students—is obviously unusual, and the fourth group (who make up only about 1 percent of American K-12 enrollment) is unusual because the vast majority of affluent parents simply live in an affluent area and send their children to a school that is public but nevertheless caters almost exclusively to people like them.

Many commentators, when attempting to envision a world with school choice, turn to the current private school sector for illumination, and they predict that children like those described

above will be the children who are most affected by choice.

Economists, in contrast, understand that what choice programs do is *relax constraints* on students' mobility among schools. Therefore, the students who will be most affected by choice are those for whom the current constraints are most binding. None of the students described above fit into this category, since they have already overcome the constraints that would be relaxed by school choice. (This is not to say that they would be unaffected by choice, since school choice could change the availability of private schools.) Economists, therefore, focus on students who are currently constrained to attend a school that unconstrained families avoid. Economists also focus on students who live in an area that can support multiple schools (since allowing a rural student to "choose" a school that is far away is not a meaningful relaxation of constraints). Thus, the students who are most likely to be affected by choice are *urban* students who (a) are either sufficiently poor or sufficiently discriminated against that their parents are constrained to live where the schools are unappealing, and (b) who are too poor to pay tuition equal to some private school's per-pupil costs.

9. *School choice is all about school supply.*

The school choice debate is also plagued by confusion about the *supply* of schools of choice. A common misapprehension is that, under school choice, students would have to be allocated among each existing school's current number of places. Another common misapprehension is that, under a voucher program that allowed religious private schools to accept vouchers, approximately 85 percent of private school enrollment would be in religious schools because that is the current composition of private schools. Such misapprehensions stem from the belief that the supply of schooling is inelastic.

Economists realize that such an assumption is extreme and very unlikely to be true. In every sector, there are factors that determine supply, and economists know that understanding such factors is the key to predicting supply accurately. Economists focus on factors that would determine what the supply of schools would look like under choice: the cost of school inputs, economies of scale, and the features on which parents are willing to “spend” their vouchers.

For example, it is useful to know how much it costs to build new schools and how much it costs to refurbish current schools so that they can be used for reorganized or new schools. Those who believe that the supply of schools is inelastic apparently believe that such costs are prohibitively high. They are not, as is demonstrated by the ability of school management companies that now routinely build new schools and renovate current schools for their use (Edison Schools, Advantage Schools, et cetera). Moreover, many school inputs are in elastic supply and can be purchased at a price that can be readily established: classroom equipment, school accounting software, computers, and so on. There are numerous economic studies of how the quantity and quality of teachers responds to salaries and benefits, and we can use estimates from such studies.

If we wish to understand what the supply of schools would look like under choice, it is also useful to know the preferences of the parents who are most bound by the constraints that a school choice program would relax. For instance, unless constrained parents have a great taste for religious education, it is unlikely that the supply of choice schools would end up being dominated by schools with religious affiliation. In short, current private school parents have unusual tastes, so their preferences tell us little about the preferences of future choice parents. Economists realize that it is more useful to examine the stated preferences, from the National Household Education Survey or opinion surveys, of parents who are likely to be constrained.

Such survey evidence is imperfect, but it is a better guide to future supply than are the preferences of current private school parents.

10. It is the threat of competition that matters.

Part of understanding markets is understanding that the schools' conduct and performance will depend on the availability of alternative schools, not on whether the parents actually use the alternatives. That is, it is the threat of competition that matters, not whether the threat results in (a) the incumbent school improving so much that parents do not want to use the alternative school or (b) parents leaving the incumbent school for a better alternative.

There are two important implications of realizing that it is the threat of competition that matters. First, if we observe a public school that loses many students to choice schools (voucher, charter, and private schools), we should realize that the school faces competition *and* is responding poorly to it. A similar school that loses few students to choice schools is not necessarily facing less competition—it may just be responding better. Therefore, cross-sectional comparison of schools that do and do not lose students to choice schools is *not* good evidence about the effects of competition. What one needs to do, in order to obtain evidence on competition, is find schools that are and are not subjected to the threat of competition. Economists understand the distinction (between the threat of competition and actual loss of students) and can design an empirical study of the *threat* of competition.

Second, once we recognize that it is threat of competition that matters, we see that students who do not attend choice schools may benefit from competition just as much as students who do. Indeed, if we want to know whether choice matters, then students at seriously threatened incumbent schools are not a good control group for students at choice schools.

Insights from the Chapters that Follow

Each of the chapters in this volume takes up a different aspect of school choice and is the culmination of a research agenda. As a result, opening each chapter is like opening a door into a body of research. The authors themselves are the best guides to their particular areas, and the authors are certainly best at presenting their own results. Therefore, I will not attempt to summarize their chapters here, but will use this as an opportunity to draw attention to insights and features of each chapter that I found to be particularly striking.

Eric Hanushek and Steven Rivkin, in “Does Public School Competition Affect Teacher Quality?” explore the hypothesis that choice will force schools to employ teachers of a more consistent, high quality. The study is central to the question of how choice (the structure of the downstream market for education) will affect teaching (the key upstream market). Ordinarily, teacher quality is very difficult to measure. For instance, in a related study, I had to use the selectivity of a teacher’s college and whether she had a subject-area degree as proxies for quality. Hanushek and Rivkin, however, are able to measure a teacher’s quality by her systematic effect on student achievement. Their clever empirical strategy requires very detailed and complete data, which they have obtained for the entire state of Texas.

Hanushek and Rivkin exploit variation in the most common, traditional form of public school choice: parents choosing among schools by choosing where to live. Two interesting insights arise in relation to this strategy. First, the evidence suggests that, although both choice *among schools* (within a district) and choice *among districts* are meaningful, they are not the same. When a parent chooses a district, there are financial implications of his choice because districts are financially autonomous and depend on local property values. When a parent chooses a school within a district, his choice does not have the same implications. Second, the Texas

data shows clearly that choice is more meaningful for metropolitan students than for rural students.

David Figlio and Marianne Page's chapter, "Can School Choice and School Accountability Successfully Coexist?" examines Florida's voucher system, in which students are offered vouchers if they attend a school that consistently fails to meet Florida's achievement standards. Proposals for similar programs (choice limited to students who would otherwise attend failing schools) enjoy considerable political popularity.

Figlio and Page's study exposes a fundamental difference between such choice programs and more conventional programs in which eligibility is based mainly on students' own characteristics (not his school's failure). Conventional choice programs rely on the idea that parents are inclined to choose better schools for their children when they can. Thus, conventional programs attempt to make escape easier for parents who are currently *constrained* (by their incomes, job locations, or other factors) to choose a bad school. Conventional programs often include provisions that increase the information available to parents, in the form of school report cards and so on. Nevertheless, choice plans usually depend on parents to filter and judge the information they receive about schools, and parents are the main source of discipline for underperforming schools.

In contrast, the Florida program relies on the idea that the state is better than parents at determining whether a school is underperforming. In the Florida program, a school's being attractive to parents is no guarantee against state sanctions, nor is a school's being unattractive to parents a guarantee of vouchers. Figlio and Page demonstrate empirically that the Florida system, regardless of whether it depends on achievement levels or value-added measures of achievement, does not give vouchers to many students who are currently constrained and does

give vouchers to many students who are currently unconstrained. A deep inconsistency in the Florida program is that its success depends on parents using the vouchers wisely, even though its structure implies that parents are poor judges of schools compared to the state.

In short, Figlio and Page's paper should make us think carefully about the fundamental claims on which choice programs are justified. Their paper is a wake-up call for people who ignore differences between Florida-style and conventional choice programs.

Julianne Berry Cullen and Steven Rivkin study the ticklish intersection between school choice and special education in their chapter, "The Role of Special Education in School Choice." The authors bring to light several important questions about school choice and special education. Can a choice program ensure that all (or most) schools have the resources to fulfil a student's individual education plan (IEP)? If the answer is "no," then special education students may have limited choice in practice, either because schools will attempt to exclude them or because they themselves will stay away from schools that do not have adequate resources. Also, would other students attempt to avoid special education students in a system with greater choice?

Using data from the Texas public school system, the authors conclude that there is little evidence that regular education students attempt to avoid special education students. However, the data do suggest that special education students are disproportionately likely to make use of opportunities to choose among public schools. Moreover, special education students already use private school vouchers in a number of states. Thus, one of the key things that we learn from Cullen and Rivkin is that choice does not merely generate unique risks for special education students, it also presents them with unique opportunities.

The authors' results lead them to ask, "In a system of school choice, who should decide what a student's IEP is?" Cullen and Rivkin show that special education families value school

choice because it allows them to seek not only sympathetic environments for their children, but also the IEP that they feel is appropriate. This is the positive side of IEP seeking. Cullen and Rivkin also suggest that some IEP seeking may be less positive: a family may switch schools until it finds an administrator whom it can bully into writing an inappropriate IEP for its child. Moreover, Cullen and Rivkin inform us that, empirically, it appears that schools already over- or underclassify students in response to financial incentives to do so. Gaming of the system can be exacerbated by a poorly designed choice program. For instance, the authors describe Minnesota's open enrollment plan, in which receiving schools were inclined to overclassify students because the sending schools were responsible to paying the costs of the IEP. Designers of choice plans will come away from the Cullen and Rivkin chapter with many ideas about how (and how not to) design the special education provision of their programs.

Paul Peterson (with William G. Howell, Patrick J. Wolf, and David E. Campbell) compares the outcomes of students who are randomly given and not given vouchers in "School Vouchers: Results from Randomized Experiments." The randomization occurs when applicants for vouchers put their names into lotteries, because the number of available voucher is fewer than the number of applicants. The randomized design has features that are obviously desirable. In particular, it is extremely plausible that the randomly selected "treatment" and "control" groups of students have similar unobserved characteristics. Because we cannot check students' unobserved characteristics, it is valuable to have a design that guarantees similarity to the maximum extent possible.

Readers will naturally want to focus on Peterson's comparison of the standardized test achievement of voucher and control students. In his chapter, Peterson describes these results for several voucher programs, all of which target poor children. Readers should be encouraged,

however, to look further than the achievement results, because Peterson's data are rich on other dimensions. We can derive what anthropologists call a "thick description" of the entire voucher experience.

For instance, we learn that younger children adjust more readily to voucher schools and consequently improve their achievement more quickly after the choice opportunity is made available to them. Older children are more likely to complain about tough new discipline or academic demands, at least in the immediate aftermath of the transition. (As we might expect, the parents of older children do not always share their children's dissatisfaction.)

A large share of the students in Peterson's data are black, owing to the location and eligibility criteria of the programs that he studies. Thus, the results for white students are less precise than they are for black students, and we should be circumspect about interpreting differences between the black students' and white students' results. Nevertheless, the data suggest that black students benefit more from vouchers, perhaps because black families are more constrained without choice (they are poorer or have more limited housing choices) or because black students suffer from more discrimination in public schools than in voucher schools. These results are intriguing, and suggest "leads" for unraveling the puzzle of America's significant black-white achievement gap.

Even with a randomized design, some empirical problems arise in Peterson's study, and he shares his solutions to them. For instance, are the families who apply for the vouchers very different from non-applicants? The randomized design does not have a natural way of comparing applicants to non-applicants. Peterson solves this problem with additional data, which he gathered through representative surveys of local families. Another tricky question is what to do about families who win vouchers but do not use them. After all, the randomization is only over

voucher receipt, not voucher *use*. Instrumenting for voucher use with voucher receipt allows Peterson to recover an unbiased estimate of the effect of voucher use—that is, the effect of voucher use uncontaminated by biases that might creep in if a certain type of families were more likely to leave the voucher unused.

The inclusion of Peterson’s study in this volume deserves special comment, not because he is a political scientist (which is unimportant, given that he uses modern econometrics), but because economists’ initial response to a study like Peterson’s is nearly always, “Why are such studies needed at all?” Economists reason that if parents use the vouchers (when they could easily continue to use the public schools or return to the public schools), then it is obvious that “treatment” parents are more satisfied than they would otherwise have been. Is not it obvious that their children are doing better, in some parentally-defined sense if not in terms of standardized tests? That is, economists tend to take a revealed preference view: they are as unwilling to believe that families will choose their less-preferred school (when another is freely available) as they are to believe that families will shop at their less-preferred grocery store. In addition, economists wonder how to use the study’s results: they are disinclined to extrapolate from *some* voucher schools to *all* potential voucher schools—just as they would be disinclined to say that all grocery stores are preferred to A&P just because certain grocery stores are preferred to A&P in some localities. Finally, economists are loath to believe that families can be made worse off when they are given an option that they did not have before, except when their default public school deteriorates because other families were given the option to leave it too. (The possibility of such deterioration is the reason why economists are interested in how choice affects student sorting and school finance.) Since Peterson examines voucher programs that are tiny relative to the public school systems from which they draw students, it is very unlikely that the

families' default public schools could deteriorate significantly because of the voucher program. Thus, economists reason, families *cannot* be made worse off in the situations that Peterson studies.

Given economists' initial reactions to studies like Peterson's, why do they return again and again to such studies and where do such studies belong in the economic analysis of choice? The answer is that Peterson's study helps us to learn about *fundamental* parameters that are necessary for predicting the response to school choice. For instance, commentators on education sometimes doubt whether parents—especially central city, poor parents—take achievement into account at all when they judge a school. Commentators often suggest that poor parents are drawn to schools mainly on the basis of ethnic concerns (they like ethnocentric curricula regardless of the effect on achievement), laziness (they choose whatever school minimizes their effort), or superficial attributes (they like attractive uniforms). Not only does Peterson's study allow us to estimate how much parents seek achievement when they make school choices, the study also gives us a relatively rich picture of the families whose behavior will be important in a choice environment.

A final note on Peterson's analysis is in order. It is not subject to the criticism, mentioned above, that students at incumbent schools are invalid "control" students because their schools may be affected by the threat of competition, even if they do not exercise choice. The voucher programs that Peterson studies enroll just a few percent of local students and they are *privately funded* (so the public schools do not lose any money when they lose a student). Thus, the incumbent schools attended by the control students are unlikely be significantly changed by the voucher programs.

Thomas J. Nechyba, in "Introducing School Choice into Multi-District Public School

Systems,” demonstrates how important it is to model the *current* school system realistically before attempting to predict the effects of school choice. Nechyba manages to reduce the complexity of the current system to its essentials, without losing any of those essentials. In his model, people simultaneously choose their public school district, residence, and whether to send their children to private schools. Voters of each school district choose the level at which their schools will be funded, with the tax base being the property they actually own. In short, he manages to incorporate a housing market, private schools, and political economy in a tractable model. He uses data from actual school districts to bound the parameters for his model.

Nechyba emphasizes that the current system pressures families to live in school districts with other families who have similar demand for housing (in other words, similar incomes). This force is responsible for much of the inequality in school spending that we see in the current system. Vouchers sever the link between housing and schooling, and thereby reduce the incentive for families to segregate themselves on the basis of income. For instance, well-off couples who currently leave the central city for the suburbs as soon as they have school-aged children would be more likely to remain under a voucher program. Nechyba demonstrates that, because vouchers induce greater *income integration*, they do not raise the inequality of schools’ resources much, if at all. He demonstrates how incorrect a commentator’s predictions will be if he (the commentator) characterizes the current system as an idealized system in which all public schools are equal and no areas are segregated on the basis of income.

In a series of studies, Raquel Fernández and Richard Rogerson have examined how school finance affects the growth and distribution of income in a nation. Each of their studies has wrestled with the same fundamental problem. The market for investments in human capital is flawed because little children do not borrow against their future earnings to finance the

investment in education that would be optimal for them. (Little children cannot borrow for several reasons. They are not competent to sign contracts that would bind them for life to a schedule of repayments; their parents cannot sign such contracts on their behalf; families know little about a child's abilities and about the future labor market; and human capital is not collateral that could be repossessed by a lender if the borrower were to default on his education loan.)

The flawed market for human capital investments generates investments that are systematically inadequate for certain groups, especially poor families who are unable to provide internal (family) financing of education. Such inoptimal investment translates into slow growth for a country and unnecessary inter-generational transmission of economic status. Fernández and Rogerson first remind us of how well current school finance systems remedy the market flaws. They treat pure local finance as their benchmark, but also discuss some popular state aid systems. They go on to examine three different types of vouchers that would be straightforward to enact: lump-sum vouchers (every student gets the same voucher), means-tested vouchers (every student with less than a certain income gets a voucher), and power-equalizing vouchers (every student with less than a certain income gets a voucher that rises as his family's income falls and as the *share* of his family's income that is devoted to schooling rises).

Using a calibrated theoretical model, Fernández and Rogerson demonstrate that the vouchers, especially the power-equalizing vouchers, generate large increases in a nation's income and well-being. One key insight is that vouchers, because they are specific to individual students (rather than entire districts) generate more optimal investments in education than any version of current school finance (which operates at the district level) could generate.

Dennis Epple and Richard Romano, in "Neighborhood Schools, Choice, and the

Distribution of Educational Benefits,” study *intra*-district choice. The distinctive feature of *intra*-district choice is that there is no channel by which a school gains per-pupil spending when it attracts a student (or loses per-pupil spending when it loses a student). *Intra*-district choice plans are used in some large cities (for instance, Chicago). Some states, if they were to adopt public school choice plans, would effectively adopt *intra*-district plans: Hawaii has one school district for the entire state; California is approximately one district for financial purposes; and a few other states (such as New Mexico) allow very little variation in per-pupil spending among their districts.

Epple and Romano compare *intra*-district choice to typical neighborhood schools, in which each student attends his assigned local school. They demonstrate that there is an important trade-off between *intra*-district choice and neighborhood schools, so long as every student benefits from being with students of higher ability (a crucial assumption). The trade-off is as follows: *intra*-district choice encourages segregation of students on the basis of ability, compared to neighborhood schools; but *intra*-district choice discourages segregation of students on the basis of income, compared to neighborhood schools.

Another insight to take away from Epple and Romano’s chapter is that the effects of *intra*-district choice depend very much on whether surrounding districts use *intra*-district choice too (or retain neighborhood schools) and whether surrounding districts are close substitutes for the district with *intra*-district choice. For instance, if one district in the midst of many others unilaterally enacts *intra*-district choice, then well-off families who prefer to send their child to an income-segregated neighborhood school will move out of the district with *intra*-district choice. The systematic departure of such families will drive down per-pupil spending in the district with *intra*-district choice.

In “School Choice and School Productivity: Or, Could School Choice Be a Tide that Lifts All Boats?”, I begin by reviewing some facts that suggest that American schools could be substantially more productive. (The productivity of schools is measured by dividing a standardized measure of students’ achievement by per-pupil spending.) In the last 30 years, the productivity of American schools has fallen by between 45 and 75 percent, depending on how one controls for changes in the socio-demographic composition of the population and for the increased cost of hiring well-educated female workers. Therefore, even if choice could make American schools recoup only one-third or a half of the productivity that they have recently lost, then the productivity effect might easily swamp other effects of choice. For instance, suppose that choice, through a channel like student sorting, has negative effects on some students’ peer groups. What we know about the scale of peer effects suggests that regaining half of the lost productivity would easily outweigh such negative effects.

In the chapter, I examine how three recent choice reforms affected the productivity of *incumbent public schools*. The three reforms (Milwaukee vouchers, Michigan charter schools, and Arizona charter schools) were selected because the number of eligible students and the fees associated with them were large enough that incumbent public schools might begin to feel threatened. I show that, as one would expect, the schools that lost students under choice were schools that were underperforming when the choice programs were enacted. In fact, it appears that charter schools, which have some discretion about where they will locate, picked locations at which they would have access to a population of dissatisfied families.

I find that incumbent schools reacted surprising quickly and positively to the threat of competition. For instance, public schools in Milwaukee, which began to be seriously threatened only in 1998-99, raised in their test scores remarkably in first two years they were threatened

(with no relative increase in spending), compared to similar, but unthreatened, public schools elsewhere in Wisconsin. School superintendents sometimes claim that they have a backlog of changes that need to be made, and that a serious competitive threat allows them to make several changes at once. Because the quick, large reactions may reflect such backlogs, the chapter also reviews evidence of long-term productivity reactions to the availability of traditional forms of choice (choosing a public school by choosing a residence and sending a child to a regular private school).

Looking Forward

In the coming years, we can look forward to further advances in the economic analysis of school choice. Empirical evidence will grow in proportion to the enactment of reforms, and—fortunately for research—the number and variety of charter school, voucher, and public school choice programs are continually increasing. It would be particularly helpful, from the research point of view, to have one state or one large metropolitan area enact a choice plan that is both stable and relatively universal. Arizona’s charter school program is the nearest approximation to this that we currently have, and it has been in place only since 1994. So far, we have had to rely on traditional forms of choice in order to get evidence on the long term, general equilibrium effects of choice—such as the effects on the market for teachers, student sorting, residential segregation, school finance, and the housing market.

There is a great deal of variation in the financial arrangements of the experimental choice plans enacted over the past several years, but much of the variation is not useful for empirical research. Many of the experimental plans have financial arrangements are obviously unsuitable for full scale choice, and we can learn only a limited amount from them. The future is brighter:

the happy merger between research on school choice and school finance is already leading to more thoughtful construction of the financial side of choice plans. As the plans get refined, it will be easier for researchers to learn about how choice changes the school finance environment.

We can learn from foreign school choice programs, as well as American ones, but on this front, I am obliged to raise warning flags as well as hopes. Empirical evidence may be helpful if it comes from a foreign country that has similar school finance, market orientation, and culture as the United States does, but very few countries fulfil one of these criteria and no foreign country fulfils all the criteria. (Canada is the closest by far.) Some research on foreign school choice assumes that the results will translate readily to the United States. This is naive. School choice plans are layered on top of America's current system of public and private schools, which is an outlier in the world. American schools are more locally controlled, more reliant on local funding, and more entwined with the housing market (because of property tax revenue and Americans' greater residential mobility) than any other schools in the world. Inequality and disability exist everywhere, but America has a unique legal history regarding school desegregation and special education. In short, many of the questions that arise in a typical American debate on school choice are *peculiar* to the United States.

More importantly, school choice is fundamentally a market-based reform, and Americans have very different experience with markets than most people in the world. By world standards, Americans are confident consumers who negotiate markets well but take dictation poorly from social planners, even when the social planners are benevolent. In France, for instance, an educational elite chooses the curriculum for the entire nation and determines which students will be able to attend academic high schools and college. This system may work well for the French, but Americans have consistently resisted letting an elite group decide how their children should

be educated and which children deserve further education. That is, Americans appear to *like* making choices about their children's education. Also, the American labor market is less regulated than other countries'. The typical American teacher is not only familiar with jobs in which pay and promotions are market-oriented (unless she has a strange lack of acquaintances among the non-union, private sector workers who make up most of the workforce), she would probably be shocked to learn that most foreign teachers are on *nation-wide* contracts and are *assigned* to a school by a national ministry. She is already more market oriented than her foreign counterparts and probably values her ability to change her pay or working conditions by moving to a new district. In many other countries, government provision of services and elite decisions about education are not only more common, they enjoy far stronger social support. It is poor inference to assume that, because foreign families or teachers react in a certain way to a choice program, Americans will necessarily react similarly. We will learn the most from research on foreign school choice when we take the trouble to articulate and quantify the differences in institutions and typical economic behavior.

I can return to a more optimistic tone with theoretical economic analysis of school choice. On this front, there is little that constrains research from continuing at a good rate of progress, except for the energy of researchers. As a theory problem, school choice sits neatly at the intersection of several fields: labor economics (because investment in education is a human capital problem), public economics (because schools must be financed), and industrial organization (because schools must compete with one another under choice). In addition, analysis of school choice draws upon general equilibrium and overlapping generations models most often used by macroeconomists. We have not yet come close to exhausting *existing* tools' capacity for analyzing school choice. However, nearly every existing tool was initially

designed for another purpose than the analysis of school choice, and thus each tool must be modified to fit American educational institutions. Also, analysts find themselves having to learn economics outside their “home” field. In short, though progress is very likely, it must proceed at a somewhat measured pace.

Theory is likely to advance especially along the lines of *optimal design* of choice programs. In the first years of analysis, economists were (not surprisingly) absorbed in analyzing existing choice plans or proposals. Such analysis led to greater understanding of the issues, but the natural consequence of such understanding was that economists began to envision programs that dealt better with choice problems than existing programs do. Design of more optimal programs is a productive agenda, and recent papers (including the one by Fernández and Rogerson in this volume) illustrate its usefulness. We are just beginning to explore the potential of school choice to solve long-standing problems in school finance, racial and income segregation, and special education. We will continue to learn about the capacity (as well as the risks) of school choice to make schools more effective and to match students better to schools.

1. See Milton Friedman, "The Role of Government in Education," in Robert A. Solo, ed., *Economics and the Public Interest*, Rutgers, NJ: Rutgers University Press, 1955. Friedman also discusses school choice in Milton Friedman and Rose Friedman, *Free to Choose*, New York: Harcourt, Brace, Jovanovich, 1980.

2. Jencks's proposal may be found in a variety of writings, including Christopher Jencks, "Education Vouchers: Giving Parents Money to Pay for Schooling," in *New Republic: Journal of Politics & the Arts*, Volume 163, Number 1 (July 4, 1970), pages 19 to 21. For the work of Coons and Sugarman, see John E. Coons and Stephen D. Sugarman, *Family Choice in Education: A Model State System for Vouchers*, Berkeley: Institute of Governmental Studies, University of California, 1971. Also see John E. Coons and Stephen D. Sugarman, *Education by Choice: The Case for Family Control*, Berkeley: University of California Press, 1978.

3. See, for instance, Witte's report in John F. Witte and William H. Clune, eds. *Choice and Control in American Education*, New York : Falmer, 1990.
4. Alum Rock was the only school district in the United States willing to volunteer for the Office of Economic Opportunity's study of school choice (later managed by the National Institute of Education). No private schools were involved in the study, and schools were guaranteed that no money or jobs would be lost. The Rand reports were published as Rand Corporation, *A Study of Alternatives in American Education*, prepared for the National Institute of Education, Santa Monica, CA : Rand, 1978-1981.
5. The report that is generally known as the "Coleman Report" is: James S. Coleman, Ernest Q. Campbell, Carol J. Hobson, James McPartland, Alexander M. Mood, Frederic D. Weinfeld, and Robert L. York, *Equality of Educational Opportunity*, Washington, D.C.: United States Government Printing Office for the National Center of Educational Statistics, 1966.
6. The estimates described come from Alan Krueger, "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics*, CXIV (1999), 497-532. Other modern, credible research finds no effect of class size reduction: see Caroline Hoxby, "The Effects of Class Size on Student Achievement: New Evidence from Population Variation," *Quarterly Journal of Economics*, Vol. 115, No. 4 (November 2000), 1239-1285; and the review by Eric Hanushek, "Measuring Investment in Education," *Journal of Economic Perspectives*, X (1996), 9-30.