Note: A version of this paper was published as:

Julian R. Betts, (2005), "The Economic Theory of School Choice," Chapter 2 in Julian R. Betts and Tom Loveless (Eds.) **Getting Choice Right: Ensuring Equity and Efficiency in Education Policy**, Washington, D.C.: Brookings Institution Press, pp. 14-39.

Chapter 2

Economic Theory of School Choice

by Julian R. Betts

In a variety of forms, school choice is making gains in America's K--12 school systems. At one end of the spectrum are open-enrollment programs that allow students to switch from their local school to other public schools in the same district (or in some cases, other districts) subject to space being available. In addition, almost all states now have provisions for charter schools. At the other end of the spectrum are experimental voucher programs in cities such as Milwaukee that provide public subsidies to students to enroll in local private schools.

School choice, especially its more radical variants such as vouchers, has attracted strong reactions, both positive and negative. In light of this public controversy, it makes sense to take a step back from the fray to think through just why school choice might improve education. Because it is clear that school choice is not about to disappear, it also makes sense for all those who care about K--12 education, regardless of their beliefs about the wisdom of school choice, to examine how school choice should be implemented. This chapter seeks to provide some insights on the why and how of school choice. It draws insights from economic theory, which has developed a mathematical model of outcomes in markets under various assumptions about the degree of competition in the given market. This is directly relevant to school choice because the idea of choice programs is to create a competitive market that forces schools to compete for consumers (families and their students).

The first question I address is the theoretical case for why school choice would improve on the traditional system in which students are compelled to attend their local neighborhood school. The second question I address is how to implement school choice. Theory provides insights here because it yields specific warnings about what can go wrong in markets. The chapter devotes considerable attention to what can go wrong with school choice and offers some concrete suggestions for ways to minimize the potential problems. A second implementation

issue concerns inequality in outcomes across different students and their families. Again, economic theory provides insights into mechanisms that might reduce some very real concerns that school choice in practice will lead to large variations in the quality of education received by all students.

While the chapter draws frequently from the well of economic theory, the goal of the analysis is to provide some new---and accessible---insights.

The Theoretical Case for School Choice

For the most part, public school students are required to attend their local neighborhood elementary school and the middle and high schools that are in the same feeder pattern. So roughly speaking families choose a school for their children in one of two ways, either by placing them into private schools or by choosing a public school through their residential decisions. For the last three decades only about 10--13 percent of K--12 students in the United States attended private schools, probably in part because the cost poses a barrier to some. Public school choice is also limited in the sense that most families have a limited number of neighborhoods in which they can afford to live. Taking this further, one could argue that the school district can act as a monopoly provider of educational services. What are the relative outcomes between this simplified world, in which families have virtually no choice among public schools for their children, and a world with school choice, in which parents could enroll their children in any of a large number of schools? Economic theory provides an analysis of these two extremes of monopoly or perfect competition. In the former, a single provider dominates the market; in the latter, there is a large number of providers and buyers (families and their students), none of which is sufficiently influential to manipulate the market outcomes.

Perfect Competition

When the market for a specific good or service is composed of a very large number of buyers and sellers, perfect competition is likely to emerge. Perfectly competitive markets have one marvelous property: if the conditions for perfect competition are met, then no buyer or seller could become better off without making at least one other agent worse off. Economists call this Pareto efficiency. The market is economically efficient in the sense that resources are not wasted and there is no way to make individual consumers of the service or product better off

without taking something away from another. The key to this outcome is that individual producers and consumers make optimal decisions in a highly decentralized way. To gain insight into how the market's "invisible hand" generates these outcomes, it is important to review the necessary conditions for a perfectly competitive market to exist:

---Both firms and consumers are sufficiently numerous that none can affect the product's price.

---The market is for a homogeneous commodity or service.

---Firms are willing to sell to all consumers equally.

---There is perfect information on price.

---Firms maximize profits, consumers maximize "utility", that is, their private interests.

---There is free entry and exit for both firms and consumers, so that, for example, new schools can start up without any barriers, and schools that are failing can close.

---There are no externalities in production or consumption. (Externalities refer to interactions between market participants that work outside of the normal market mechanism in which prices adjust to equate supply and demand.)[1]

If all of these conditions are in place, then, simply put, consumers as a whole could not do better. Each firm will have to compete so hard for market share that it makes only a normal rate of profit on its physical investments. If firms are making above-normal profits due to high prices, then new firms will be attracted into the industry. The overall supply of the product will then rise, driving prices down to the minimum that firms can afford to charge without going out of business or switching to another more lucrative industry. In addition, firms that are less efficient than other firms will be forced to shut down, as they cannot compete effectively.

It is easy to extend this model to a world in which consumers buy many goods and services in markets with perfect competition. In this case, the relentless forces of competition cause prices for some goods and services to rise and for others to drop, until just the right mix of goods and services is produced for economic efficiency. Again, changes in prices send signals to producers about what consumers want, inducing the entry of firms into booming sectors and the exit of firms from lagging sectors.

This outcome is a sort of gold standard against which to compare markets in the real world. That is all well and fine, but what does it have to do with public schools? At present, not much. We need to imagine a

situation in which school administrators have almost complete control over school size and the mix of services that they provide. We also have to imagine a situation in which parents, perhaps through a voucher program, have a much wider array of schools available for their child than at present. We can think of schools as producing multiple services. If there are enough schools to choose from, parents will find the best school for their children, given their needs and interests.

To give a simple example, suppose that schools provide two types of educational service that parents value to varying degrees: math achievement and language achievement. The curve in figure 2-1 labeled production possibilities frontier shows all the combinations of math and language achievement that a school can provide. The frontier has a negative slope, reflecting the idea that if school administrators decide to boost math achievement they will have to take resources out of programs that are directed at teaching students how to read. The production possibilities frontier is also curved, in a concave fashion. This curvature is based on the concept of diminishing returns: it is possible to increase language achievement further to the right at the cost of decreasing math achievement, but as language achievement is repeatedly increased by, say, one point, the school would have to give up more and more math achievement for each one point gain in language achievement.

[Figure 2-1 here]

The other lines in the figure are indifference curves. The points on a given indifference curve show all combinations of math and language achievement that the family would view as equally good outcomes.[2] Families would prefer their child's math and language achievement to be as high as possible. So points on the indifference curve that is further to the northeast in the figure represent a higher degree of family welfare than any point on the lower indifference curve. If families could choose a school anywhere along the production possibilities frontier they would definitely prefer point *A* to point *C*. Furthermore, point *A* is the very best combination of math and language achievement that the family can obtain: there is no higher indifference curve that is on the production possibilities frontier.

No school would want to provide math and language achievement to the southwest of the production possibilities frontier in the figure, say, at point *X*. (A school might arrive at point *X* if its administrators have goals that do not coincide with what parents care about, thus spending money on something that parents do not value.) If the school does offer a combination of educational achievement given by point *X*, then parents would switch

their children to another school that, for the same cost, would provide improved teaching in math or language or both.

School choice, if it worked as described above, could accommodate different tastes among families for the various school services. Suppose that half of families believe that it is extremely important that their children learn technical skills like math, while the other half believe in a broadly based liberal curriculum with a foundation in reading and writing. Figure 2-2 illustrates this, showing the different indifference curves of the two types of family. Schools would adjust to these different preferences quickly or would lose all of their students. We would see, in the figure, half of schools operating at point *A* and the other half operating at point *C*. In both cases, schools are being efficient in the sense that either type of school produces the highest level of math achievement that it can without sacrificing language achievement, and vice versa. But the two types of schools differ considerably in their academic focus. Together they jointly serve the needs of the two types of family.

[Figure 2-2 here]

If all schools had to compete for students along the lines described above, then there would be no way to improve the outcomes for one student without hurting outcomes for another student. This is the beauty of perfectly competitive markets. In addition, it is important to note that this first-best outcome is arrived at in a very decentralized way. Thousands of decisions by families within a district---combined with perhaps hundreds of decisions by providers of educational services about whether to set up a school and what sort of services to provide at each school---work through the interaction of supply and demand to create an efficient outcome, in the sense that nobody can be made better off without somebody else becoming worse off.

It is important to note that the first-best result of perfect competition leaves each family as well off as possible in utility terms, without decreasing the welfare of other families. This is not necessarily the same thing as saying that student achievement is maximized. If families care about aspects of education other than achievement in math and reading, then choice will not necessarily maximize achievement. (I return to this point later.)

Monopoly

It is useful to compare the "first-best" case of perfect competition with its polar opposite, monopoly. Monopoly refers to a situation in which there is only one producer of a good or service in a given market. It produces worse

outcomes for consumers than does a market with more choice for consumers. Imagine that one producer dominates the market for widgets. As the producer increases production, the overall price of the widget will fall as the increased supply floods the market. While in theory this producer could keep producing more and more units until the price of widgets just equaled the cost of producing one more widget, it will not be profitable to do so. Instead, a monopolist's best course of action is to restrict the amount produced. Reducing output reduces sales volume, but on the other hand it increases prices. The former reduces profits while the latter increases profits. The monopolist will reduce output until these two forces just counterbalance. Consumers suffer compared to a perfectly competitive situation, because not only do they consume less of the product but they pay more.

At first this story about a monopoly widget manufacturer seems to have nothing to do with public education. But upon closer inspection there are close parallels. In the market for widgets, if the monopolist sets the price too high, consumers will not buy any widgets. School is mandatory for young children, but even facing a monopolist school district, parents have the option of "refusing to buy" by sending their children to private school, by home schooling, or by moving to a different locale. So a district that is a local monopolist cannot completely ignore the preferences of local families. But it can afford to do less than would a school in a more competitive setting without losing many of its students. Without the threat of competition, a district will not in general provide the same quantity of resources for a given amount of funding as will a school in a more competitive environment.[3]

The parallel between a monopoly firm and a school district that faces no competition is not faultless. Clearly, public schools and districts do not directly charge parents for the service that they provide. But through the political process, families do choose overall spending per pupil for the district, and the given district then decides how to spend these dollars. Another difference is that districts do not directly control the number of students they serve in the same way that the widget firm sets output levels. But in effect, districts already do decide on the quantity of educational services to provide along many dimensions: the number of teachers at each school, the length of the school day, the intensity of remediation efforts for students who fall behind, and so on. In either case, the lack of competition will remove the incentive for the provider to provide the amount of output consistent with making consumers as well off as possible.

Similar arguments apply to the case mentioned earlier in which half of families prefer one type of education while the other half prefer another. Unlike the competitive situation in which schools would be forced to adjust so as to provide exactly the right mix of services for one type of family or the other, a monopolist might find it cheaper to provide a one-size-fits-all set of educational resources. In figure 2-2 this is illustrated by point *B,* in between points *A* and *C,* which are the two types of education strictly preferred by the two types of family. But both types of family would be strictly worse off than in a system with highly responsive and specialized schools that provide an educational mix described by points *A* and *C* respectively. We can see this because at points *A* and *C,* as obtained under perfect competition, both family types are on a higher indifference curve than they would be if their only choice were a "compromise school" operating at point *B*.

There are other costs of organizing education around a monopoly provider. In the perfectly competitive world in which education is not a single service but a complex set of variables, families would move to schools that best matched their children's needs. Thousands of such decentralized decisions by parents prompt the providers of educational services either to respond to families' preferences or to perish. In the case of a school district without competitors, the administration must bear huge information-gathering costs if it desires to provide the best education possible for each student. Even in the unlikely situation in which a district could mimic the perfectly competitive outcome by providing a variety of schools, the cost of creating a bureaucracy that would gather and synthesize all of this information would be daunting.

A related possibility in a monopolistic situation is that the producer's own preferences can lead it to produce less efficiently than it might otherwise. The district would spend money on programs that are of interest to administrators or teachers but that hold no value for families, and the district would produce an outcome to the southwest of the production possibilities frontier, such as at point *X* in figure 2-1. A monopolist district could get away with spending on things that parents do not value, simply because parents have no viable alternative source of schooling. I argue earlier that, in the case of perfect competition, if a school chooses to provide English and math classes at this relatively low level, then parents would switch their children to another school that, for the same cost, would provide better math and language classes. But in the case of monopoly, the families at such schools have little recourse.

Theory and Reality: Lessons for Implementation

This section summarizes some of the most likely deviations between the theory of perfectly competitive school choice and the reality of school choice. The goal here is not to poke holes in the rationale for school choice. On the contrary, the goal is to catalogue some of the obstacles to the first-best outcome of perfect competition and then to offer at least partial policy solutions. Seven assumptions are required for perfect competition to flourish (listed above). We discuss likely deviations from each of these assumptions.

Assumptions One and Two: Numerous Suppliers and Buyers of a Homogeneous Product

Two of the key assumptions required for perfect competition are that the good or service is homogeneous and that there are numerous suppliers and buyers. The assumption of homogeneity in education is clearly incorrect but it is easy to fix in the theoretical model: I already argued that school choice is probably best thought of as a series of markets for different flavors of education. If there are enough schools providing each type of education, then school choice may operate in nearly perfectly competitive conditions within each of these markets. However, because schools in this view operate in a series of closely related markets, the assumption of numerous suppliers of each type of schooling in a city is probably far-fetched. But this misses an important point: competition theory shows that, in a continuum between perfect competition and monopoly, some competition is almost always better for consumers than none, because it increases the quantity or quality of service provided for a given price.

Monopolistic competition is a type of imperfect competition that seems particularly likely in widespread school choice schemes. This situation borrows aspects of both monopoly and perfect competition. The seller gains some market power because consumers are not equally willing to buy from all sellers. For instance, most parents would probably prefer to send their children to a school that is relatively close to their home. This gives local schools some market power over these parents. In the short run the outcome is somewhat similar to a monopoly, with the schools providing lower quantity or quality of services for a given price, without fear of losing all of their students. What about in the long run, after new schools enter the market? Outcomes for students will start to improve due to this competition. However, in the end each neighborhood may have "too many" schools in the sense that each school enrolls too few students to be efficient. Just as grocery stores rarely operate at full capacity, the same could happen to schools in a school choice system with free entry of new schools. It is certainly true that

this is inefficient in an economic sense. However, excess capacity at most schools would generate real choice for students. (In contrast, in most school districts that currently offer open enrollment as a form of school choice, there are long waiting lists to get into the best schools.)

Assumption Three: Schools Accept All Students

Earlier I stated that firms must be willing to sell their good or service to any consumer for perfect competition to develop. If firms discriminate in some way against certain types of consumer, then the full benefits of competition will not emerge.

This problem seems particularly likely in any system of school choice that uses test scores to hold schools accountable, as this creates an incentive for schools to accept only the highest achieving students. In addition, many existing private schools prefer students from specific religions. A more radical form of school choice such as educational vouchers would need to take this into account.

In either case, the designers of school choice would need to find ways to minimize the ability of participating schools to limit enrollment based on student traits such as religion and initial test scores. This is not to say that specialized schools should not screen their applicant pools to find students best poised to gain from the offerings of the school. But safeguards need to be found to ensure that school choice does not become school choice only for the most advantaged.

Assumption Four: Perfect Information

To achieve a perfectly competitive outcome, all buyers and sellers must have full information on price. In a richer model with several types of school performing in separate markets (vocational, college prep, and so on), we would also need all market participants to have full information about the quality of schools in each of these separate markets. Lack of information among families will obviously lead them to make imperfect choices for their children. Families need two types of information: information about the choices available to them and information about the academic needs of their children. Both types of information are needed for parents to find a good match for their child. Similarly, if public and private producers of educational services lack full information about the education market, it is likely that pockets of imperfect competition will persist in certain areas, even

though it would be in the interests of both local families and education providers to have one or more new schools open up locally.

Assumption Five: The Objectives of Parents and Producers

For perfect competition, consumers must maximize their utility, or welfare, while firms maximize profits. In the case of schooling, parents typically make decisions for their children, so we must assume that parents have their children's long-term interests at heart. The vast majority of parents probably fit this description well, but it does raise the issue of how school choice would affect the small minority of children whose parents are indifferent or even hostile to their children's best interests.

The standard assumption that firms maximize profits applies fully for the small number of schools that are run for profit, such as the Edison schools. But what about schools run by school districts, or charter schools that, while operating fairly independently from school districts, are typically nonprofit? The key here is that in a world with widespread school choice, even nonprofit schools will have to compete with each other for students. Those that fail to produce results in line with parents' expectations will lose students and may ultimately be forced to close. The central point here is simple: in competitive markets consumer sovereignty prevails.

Thus, ultimately, the key question is, What are the goals of parents? If all parents actively seek the best education possible for their children given their children's interests, talents, and needs, then it does not particularly matter whether schools are profit-driven. Even nonprofit schools will have to find ways to satisfy parents in a world with widespread school choice.

Assumption Six: Free Entry and Exit

The assumption of free entry and exit implies that parents are free not only to move their child from one public school to another but also to send their child to a private school or even to homeschool their child. Typically, parents already have these rights, although in some cases, such as open enrollment, district administrators limit the number of opportunities for parents to change their child's school---for instance offering the opportunity for school switching only at the end of the school year or semester.

A more serious issue concerns free entry and exit of sellers---that is, schools. Free entry requires that parties wishing to start a new school face few obstacles. But obstacles, unfortunately, abound: land zoning regulations, for example, or considerable paperwork. This is not to say that paperwork and public scrutiny are always bad things. But the higher the hurdles for potential entrants, the less likely parents will enjoy an abundance of choices.

Free exit of schools is perhaps an even more difficult issue. The logic of perfect competition is that if initially there is imperfect competition, with higher than average profits and lower than average outcomes for consumers, new firms will be enticed to enter. This lowers the price for any level of services provided. Inexorably, the least efficient firms will shut down as they lose clients. Now, translating this to school choice, what will happen when a local school performs so poorly that it loses most of its students? Will it shut down or will government step in to subsidize the school further? Most states have some provision for state takeover of failing schools, but to the best of this author's knowledge there is no evidence that such takeovers produce markedly better schools. An actual closing of a school is rare and would represent a major change in the way we think about public education.

Clearly, any plan for widespread school choice would have to contend with how to deal with failing schools. The genius of competitive markets is that such schools would go out of business. But because even in the most radical of choice schemes schools are likely to be important local providers of education, a choice plan needs to stipulate how local students will be cared for if their school shuts down.

Assumption Seven: Externalities

Perfect competition will not result if there are externalities in production or consumption. Production externalities seem particularly likely in education, in the sense that one student's rate of learning is likely to depend on the characteristics of other students in the classroom. There is a large and increasingly convincing literature that points to the importance of classroom and grade level peers on student achievement.[4]

If it is indeed true that a student learns more quickly if there are many highly achieving students in the classroom, the societal ideal may be to mix high and low achievers. It all depends on two things. First, the extent to which society cares about the *distribution* of student achievement as opposed to average achievement will

dictate whether it is socially best to mix low and high achievers in the classroom. Second, there are important and unsettled questions about whether moving from a system of achievement grouping to completely mixed-achievement classrooms would hurt the high-achieving students more than it would help the low-achieving students.

The issue of peer groups has obvious implications for school choice. If all parents want to surround their child with high-achieving peers in the classroom, then the most elite schools in a choice system will be flooded with applications, and if these schools serve parents' wishes, they will probably skim the highest-achieving students from the application pools. Ultimately, this would either recreate or perhaps exacerbate the large gaps in achievement seen among public schools today.

Potential Solutions

Above are listed several challenges to the notion that school choice will yield the first-best outcomes promised by the theory of perfect competition. None of these problems necessarily renders the idea of school choice unworkable. However, each challenge deserves a careful response by current and future architects of school choice. I here suggest some partial solutions to these possible problems.

Imperfect Competition without Free Entry and Exit

With parental and student tastes for education as varied as they apparently are, I argue above that it may be best to think of school choice as operating in several markets for various types of schooling. This heterogeneity, compounded by barriers to free entry by potential school providers, suggests that it is unlikely that perfect competition will ever be attained in the market for schooling. As noted earlier, imperfect competition is better than no competition at all. But policymakers need to design school choice in a way to maximize the extent of competition among schools for students. Obviously, this goal is intricately linked to the goal of creating low-cost entry and exit for school providers. Some suggestions:

---School districts could consider leasing out existing school buildings to other providers of education. This would reduce the large fixed costs of constructing new school buildings and of satisfying local zoning ordinances, both of which currently create barriers to entry. The leasing of school sites to independent providers

of educational services would also go a long way toward solving the political issue of whether regulators will ultimately have the nerve to close failing schools. Under a leasing system, closing a failing school means replacing the management of a school with another management under a new lease. The school need not physically shut down. Edison and other private providers of educational services have in some cases taken over existing schools in a similar way.

---Large schools, especially the largest high schools, could be split into parts, creating two or more schools, with perhaps 500--1,000 students in each school. The idea here is for competing schools to share the costs of building upkeep, food, and transportation, while operating under the same roof. The Gates Foundation has embraced a similar approach by subsidizing the conversion of large low-performing schools into multiple "learning communities" on a single campus.

---The paperwork required to open a new school should be the minimum needed to provide sufficient public scrutiny. This suggestion stems from the simple observation that bureaucracy can create formidable barriers to entry for new schools.

---Any district that grants a school charter should clearly guarantee the amount of funding that the school will receive for busing its students. Further, state funding mechanisms should allow for choice schools to receive above-average busing support. The idea here is that parents can send their students to a given school if there exists a means of public transport to get the student from home to the school. Obviously, American cities vary in size and in the density of their public transit networks. In many cities, though, especially those that lack a subway system, it would prove very difficult for a student to attend a distant school without adequate funding for school busing.

Schools Accept All Students

The two central issues here have to do with whether schools will avoid accepting students who lag behind academically and, in the case of private school voucher schemes, whether private schools will exclude students based on student characteristics such as religion and gender.

There is a straightforward solution to the latter problem, even though it might reduce the number of private schools willing to participate in a voucher program: any school choice system, and particularly one that

includes vouchers allowing students to attend private schools, must explicitly indicate that schools cannot select students on the basis of gender, race, sexual orientation, or religious belief.

But what about the problem of schools skimming off the cream? As mentioned in the last section, it is probably unwise to insist that every school participating in a choice system must accept every student who applies. The diversity of charter schools already in existence suggests that one benefit of choice is the variety of scholastic emphases. In this case it probably is reasonable for specialized schools to screen student applicants to ensure that students' interests and talents are aligned with the goals of the school. For instance, a school devoted to the performing arts probably should be able to screen student applicants for talent and interest in this direction. Below are some suggestions that might strike the right balance between granting schools flexibility while preventing their skimming off the most highly achieving students.

---While some schools participating in a choice system may screen student applicants for academic achievement, mechanisms such as quotas on the number of students who can be accepted from the most affluent neighborhoods, or school funding formulas that reward schools for drawing from a wide distribution of academic performance, can ensure that no school accepts only the very top students.

---If a school is oversubscribed, then the school should be required to conduct a lottery in a public setting to determine which new students enter the school. While the school should have the right to screen applicants for interests, motivation, and academic achievement, the lottery should include a minimum percentage of applicants. For example, if a school receives A applications for S spaces, where $A > S$, then setting the minimum lottery pool size to $S + (A - S) / 2$ would ensure that at least half of the excess number of candidates would be included in the lottery. However, this stipulation alone could create an incentive for schools to recruit exactly the set of students that they want to enroll ($S = A$), removing all chance from the lottery. So a second stipulation could be that a school can accept no more than, for example, 80 percent of applicants (that is, $S <= 0.8A$).

Such lottery mechanisms have other advantages. First they put the onus on each school to recruit widely if they are to meet their enrollment targets. Second, they ensure that schools' admission standards are not rigid but instead fluctuate with student demand, but in a controlled way. This in a sense mimics the manner in which the price of a good or service adjusts in a regular market to equate the quantity demanded with the quantity supplied.

As the reputation of a school rises, the number of applicants will rise, but such schools cannot simply skim off the very best candidates, because the size of the lottery pool will rise with the number of applications.

---A new school should be exempted from the requirement to have a lottery ($A > S$) for a short period of time, perhaps one or two years, until it has become well enough established to fill at least a minimum number of slots.

Perfect Information

Obviously, much can be done about providing information:

---Assuming that the school district remains the focal point of school choice, then each district should distribute a pamphlet detailing the characteristics of all schools available under each component of choice (open enrollment, traditional busing, magnets, charters, and any private schools participating in a vouchers program). The information should include the academic focus of the school, admissions procedures, information on transportation options, and summaries of the teachers' qualifications. This information should also be available on the web.

---As Laura Hamilton and Kacey Guin argue in chapter 3 in this volume, school choice will only succeed insofar as parents can obtain detailed information on how successfully schools are meeting state standards. Accessible and up-to-date information on levels and gains in student performance on state-mandated tests is crucial in this regard.

---Private schools participating in a school voucher program should administer the same tests given to students in regular public schools, to provide parents with a valid comparison of student outcomes.

---A second role of testing is to provide teachers and parents with information on each student's academic strengths and weaknesses. In reality, state testing systems vary markedly in their diagnostic power. States should consider adopting tests that give parents and teachers not only data on overall achievement in subjects such as math and reading, but also specific information on areas within these subjects in which the student needs to concentrate. This sort of information will increase the probability that parents can find a school well matched to their child's needs.

Anecdotal information based on conversations with charter school administrators suggests that sometimes regular public schools do not welcome requests for meetings between charter school representatives and the public school's parents or students. This is particularly true when the charter school is recruiting students from a grade that is not the "exit" grade at the public school, because administrators at the latter school view this as a threat. This insight leads to the following suggestion:

---Each public school should institute and publicize a systematic method for disseminating information to parents about choice schools available to their children. Examples might include an annual choice fair or a regular series of mailings.

The Objectives of Parents and Producers

The most worrying issue here is the minority of parents who are indifferent to their children's academic achievement. Chapter 6 in this volume, by Dan Goldhaber and co-authors, carefully assesses the implications of expanded choice for students who are nonchoosers. However, I do offer two suggestions in this regard:

---In schools that are not meeting state academic standards, districts should actively seek out parents to alert them to the choice options available to their children. Ideally, an advocate outside the given school should be assigned to the school to communicate these choice options to parents. It is unlikely that this alone will overcome the problem of parents who are indifferent or otherwise unable to seek out the best school for their children, but it could help.

In addition, districts might want to identify parents who are least engaged in their children's education and focus these outreach activities on those families. Here is one possible mechanism:

---Toward the end of each school year parents should be asked to fill out a school choice form that asks parents to make an explicit decision about whether to leave their child in the same school next year. Parents who fail to return the form could be selected for follow-up by the school choice advocate described above.[5] Notably, the federal No Child Left Behind law calls for schools deemed as underperforming to notify their students that they can opt to be bused to another school. This represents a first step in the direction being espoused here.

Externalities

If peer group effects exist, as a growing body of evidence suggests, then there is no easy solution to the issue of how a school choice system will redistribute the most highly achieving students across schools in a district. While it is certainly possible that grouping of students by achievement across schools could become more accentuated in a choice system, some of the suggestions listed earlier, such as admission lotteries, can mute the tendency for this to happen.

Unequal Educational Quality

As mentioned earlier, competition produces a Pareto efficient outcome, which is a situation in which no person can be made better off without making at least one other person worse off. No resource goes wasted. This result is known to economists as the first theorem of welfare economics. While this is a salutary outcome, society typically places value not only on the total level of production but also on its distribution among consumers. It is extremely unlikely that, for instance, a school voucher program would lead to identical educational outcomes for all students. Indeed, many opponents of school choice fear that it would aggravate the existing achievement gap between various racial groups and between students who live in affluent neighborhoods and those who live in less affluent neighborhoods. The biggest reason for unequal outcomes under a choice regime is that relatively affluent families will be able to "top up" the educational voucher provided by the government, enabling their children to attend more elite private schools than will relatively disadvantaged families. In addition, families with relatively high socioeconomic status also are likely to have better information about the quality of schools than other families, which may further skew the quality of education. This problem will also afflict more conservative versions of school choice, such as charter school and open-enrollment programs.

These realizations lead to an important question: Is there any reason to think that an equitable education for all could result from the creation of a competitive market through school choice? Again, economic theory provides a distinct answer. The second theorem of welfare economics states that, subject to some technical conditions, *any* Pareto efficient outcome can result from perfect competition. What is required to shift the market from one Pareto efficient outcome to another is to change the distribution of buying power among consumers. In other words, the outcome of school choice can be quite inequitable or highly equitable. What determines the outcome is the relative amounts of resources put in the hands of each family initially.

One way to see this point is to examine a variant of the production possibilities frontier shown in figure 2-1. Figure 2-3 shows a production possibilities frontier for average student achievement for students living in two different neighborhoods in a city. Suppose that neighborhood 1 is relatively disadvantaged compared to neighborhood 2. Point *A* shows the outcome in the status quo with all students attending local schools. Average achievement among students from neighborhood 1 is lower than that for students from neighborhood 2 because of the direct impact of poverty on learning opportunities, the impact of peers' achievement on learning, and the sorting of teachers between schools in the two areas.

Now, suppose that a fixed quota of seats in schools in neighborhood 2 is made available on a lottery basis to students from neighborhood 1. Through increases in peer group achievement and teacher quality, achievement of students from neighborhood 1 who attend schools in neighborhood 2 may rise. Overall, we might move from point *A* on the production possibilities frontier to point *B*: students from neighborhood 1 become much better off on average, while students from neighborhood 2 become only slightly worse off.[6]

Perhaps point *B* is not satisfactory to society because students from neighborhood 1 are still learning far less than students from neighborhood 2. What the second theorem of welfare economics tells us is that we could do even better, perhaps arriving at point *C,* by finding a way to reallocate "buying power" toward the families living in neighborhood 1.

Finding an exact mechanism for this reallocation of resources will prove difficult. A simple but ham-fisted way to help more students from neighborhood 1 is to set quotas for students from each neighborhood who will attend a school. For instance, instead of one lottery for admissions to a choice school in neighborhood 2, there could be two lotteries, with *x* percent of admissions assigned to neighborhood 1 applicants and the rest of admissions assigned to students from neighborhood 2. In practice this could become specific, using even zip codes to define the separate lottery pools.

Other mechanisms of achieving a mix of students in each school would be to take a leaf from the recent annals of U.S. deregulation of the airwaves and environmental regulation. In both cases, auction mechanisms have been used by the U.S. government to allocate scarce resources (for instance, sections of the radio frequency and rights for factories to emit specific types of pollutants). At first glance these examples seem far afield from the issue of school choice. But imagine a system in which each school is granted rights to enroll a fixed percentage of

students with very high achievement. These rights would be individually tradable among schools. How would this work in practice? Let's begin with a summary of a hypothetical status quo before the choice system is introduced.

Suppose, as we have seen in so many urban school districts, that high-achieving students typically come from the most affluent areas of the city and attend highly ranked schools in their local neighborhoods, while inner-city schools have far fewer highly achieving students. At the same time, the pattern seen repeatedly is that the most highly educated and experienced teachers tend to migrate to the suburban schools in the most affluent areas over time. In both of these ways, inner-city students get the short end of the stick.

Now imagine that at the start of the new school year a tradable market in enrollment rights is created. Some of the affluent schools find themselves with far too few rights to claim highly achieving students. Inner-city schools find themselves with a surplus of these rights. For example, suppose that the district average of high-achieving students is 50 percent and that one particular high school has 1,000 students, of whom 600 are high achieving and 400 are low achieving. It has been granted 500 permits but needs 600. The school could avoid buying any additional permits by accepting 200 new low-achieving students who have applied from other schools to bring it to the district average of a 50--50 split. Alternatively, the school could buy 100 permits on the open market from one or more schools at which more than 50 percent of students were low achieving.

The beauty of this idea is that when an inner-city school sells a right to enroll a high-achieving student it banks the money paid by the urban school. Through competition among schools with high percentages of high-achieving students, the price of each right would rise to the point at which the suburban schools are indifferent between enrolling one additional high-achieving student and buying one more right. At one extreme, then, there could be no exchange of students whatsoever. Instead, there would be an exchange of funds from schools with high-achieving students to less fortunate schools. At the other extreme, there could be no trading in rights, with high-achieving schools recruiting low-achieving students from other schools. In reality we would expect some of both: some students would leave low-scoring schools to attend high-scoring schools, and the high-scoring schools would also provide financial compensation to the low-scoring schools for refusing to enroll a fully proportionate share of low-achieving students.

The inner-city schools would receive every dollar received from the purchase of their rights. They could spend this fund on any number of academic aids, such as supplemental reading materials, after-school classes, and

instructional aides. The improved resources at the inner-city schools would, we would hope, go some way toward removing the gap in performance between inner-city and suburban schools. In turn, this might induce fewer students to leave inner-city schools and could even induce some students to leave their suburban schools for better-funded urban schools. In addition, it might attract some of the most highly experienced and educated teachers in the district back to inner-city schools from the suburban schools to which they gravitated as they gained seniority. In practice, what would likely result is that most suburban schools would opt to accept more low-achieving students from elsewhere while sacrificing some funding to pay for the right to continue to enroll an above-average percentage of high-achieving students.

There is a potentially fatal flaw in this proposal for tradable enrollment rights. Suppose that a superb school in an urban area wants to thwart a district's choice program so as to "keep the school local." It could instruct its students intentionally to perform poorly on the test used to label students as high achieving or low achieving. Then it would not come close to filling its quota of high-achieving students. In an extreme example, instead of having to buy additional enrollment rights it could fraudulently obtain surplus rights by qualifying as an underperforming school and then sell its surplus rights to other schools. So it would keep its high-achieving students and unethically make a profit at the same time!

Two solutions to this form of cheating are evident. First, if these test scores matter to students and school administrators for other reasons, then it becomes less likely that administrators will encourage students to perform poorly and even less likely that individual students will play along. For example, in California, as in many other states, high school students will soon have to pass a high school exit examination in order to graduate from high school. Neither administrators nor students will want bad outcomes on this test. Similarly, some states' accountability systems and the federal No Child Left Behind law impose sanctions on schools that underperform. In some cases, states have also created financial carrots for teachers and administrators to reward high performance or gains in performance on the state test. Typically, these sanctions and rewards are quite weak, but they could mitigate the temptation for school administrators and teachers to discourage student effort on the state test. Furthermore, in California strong individual student performance in grades 9 through 11 up to 2002 earned students college scholarships, providing a direct incentive to at least some students to take the state tests seriously.

[7] In a similar vein, Georgia offers its HOPE scholarships to high school graduates who maintained a GPA of at least 3.0, or B, in core subjects in high school.

Financial incentives may not be enough to prevent schools from attempting to mislabel students as low achieving. A more bulletproof solution technically is to assign tradable enrollment rights using some characteristic of students that is less manipulable. For instance, each school could receive permits to enroll the district average percentage of students who are not eligible for lunch assistance. Those schools with almost no students eligible for lunch assistance would have to bus disadvantaged students in, pay inner-city schools for their excess permits, or a combination of both. Other alternatives include using parental education or even student race as a proxy for achievement. While these student characteristics are less susceptible to manipulation, it is clearly much more politically intractable to allow schools to sell enrollment rights for "rich students" or even more explosively, students of a given race. For these reasons, classifying students by their achievement levels, on tests that matter, may prove a more realistic if still imperfect solution.

A closely related concern is that even if schools do not cheat by encouraging students to perform poorly on tests the first year, they still have an incentive to encourage students not to improve much over time. After all, a school that initially has mediocre test scores and that then works hard to improve student achievement would suddenly have to buy permits to keep its students whose scores had improved. This is a rather perverse incentive. One solution that would work well for high schools and middle schools is that the district would categorize students as high achieving or low achieving based on a battery of test scores in the students' final year in the feeder school before he or she graduated to middle or high school. Solutions for elementary schools are less obvious. Here, perhaps pretests of cognitive development given by the district when the students begin elementary school would reduce a school's opportunity to manipulate skills downward.

A final concern with this market for enrollment rights is that the sellers of rights (the schools with below-average percentages of high-achieving students) may have an incentive to collude, just as energy producers apparently conspired in 2000 to restrict natural gas and electricity supplies to California, thereby driving prices sky-high. Drawing a lesson from this attempt at energy deregulation gone awry, the school district could set a cap on the maximum price for which a single enrollment right could sell. One way to do this is to tell high-achieving schools that they can either buy the required number of permits on the open market or pay a fine of, say, $500 for

each high-achieving student lacking a permit. This penalty would guarantee that no school would be willing to pay more than $500 for an enrollment right. The proceeds from any penalties paid could be spread across all schools equally or used to subsidize the busing programs that school choice would require. An advantage of the tradable permits proposal over the quotas idea is that the former would shift money from schools with a larger share of high-achieving students to schools with fewer high-achieving students. This transfer of cash would mitigate the concerns many have that choice will hurt students who continue to attend failing schools.

Both of these proposals, quotas combined with lotteries and an interschool market, may appear to some readers to be monstrously radical notions that bear some resemblance to Jonathan Swift's "A Modest Proposal," without Swift's tongue-in-cheek delivery. But in the case of both quotas and markets for tradable enrollment rights the goal is the same: to improve outcomes for the least advantaged students in a school district through a choice mechanism. Moreover, flavors of at least the first proposal already exist. The federal No Child Left Behind legislation calls for schools that states identify as failing for two consecutive years to allow some of their students to be bused at public expense to better schools. In this plan, the low-achieving students have the first right to transfer out of their school. Similarly, lotteries have become a common way to admit students to charter schools, and in California state law requires this in cases where charter schools are oversubscribed.

Concluding Remarks

This chapter lays out the theoretical justification for school choice by tracing the logic behind the economic model of perfect competition and explaining why competition leads to a first-best efficient outcome. It then assesses numerous ways in which school choice systems might violate the assumptions needed for perfect competition to work. It will take real-world policy experiments to understand whether these potential problems are mere flies in the ointment or quite severe barriers to school choice programs. Each of these potential obstacles to school choice is addressed and steps are offered that policymakers might take to mitigate these problems.

It seems apparent that it will not be possible to remove all of the inherent tensions in school choice. For instance, school choice will probably create greater heterogeneity in teaching styles and curriculum than currently exists, and this balkanization of the "market" for education stands in direct opposition to the idea that school choice will improve schools through large numbers of schools competing to produce a homogeneous service.

There is also a tension between the goal of increasing the number and range of schools available to each student and the goal of reducing the large achievement gap between disadvantaged and more affluent students and among racial and ethnic groups. However, as the penultimate section indicates, with a few technical restrictions competitive forces can produce any Pareto efficient outcome policymakers seek, whether equal or unequal, depending on the initial distribution of buying power among consumers. Therefore, those who are concerned that choice will widen existing inequalities in outcomes should favor choice programs that level the playing field in favor of the poor. Two ways of achieving this in some part are either to set geographic or other quotas for each school's student body or to use more sophisticated market mechanisms that force schools in more affluent areas to compensate schools in less affluent areas for the right to enroll above-average shares of high-achieving students.

Neither of these tensions can be resolved completely, but careful implementation, perhaps using some of the suggestions outlined here, might alleviate these and other negative consequences that could accompany a system of school choice.

Finally, it seems clear from the above analysis that unregulated school choice is not consistent with good public policy. We have consumer protection laws and government oversight in all sorts of markets; the market for schools should not be an exception to this rule.

Figure 2--1 Tradeoffs between Math and Language Achievement for a Specific Family

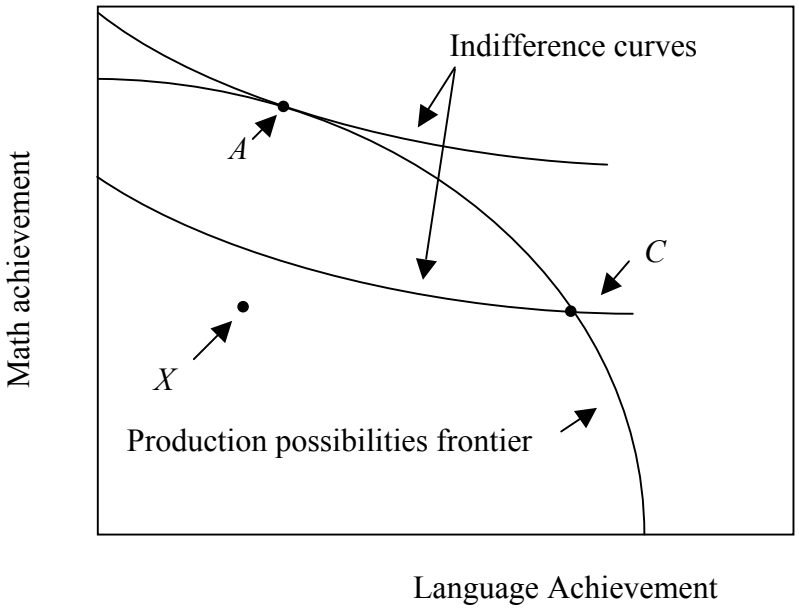

Figure 2--2 Tradeoffs between Math and Language Achievement with Heterogeneous Tastes for Education

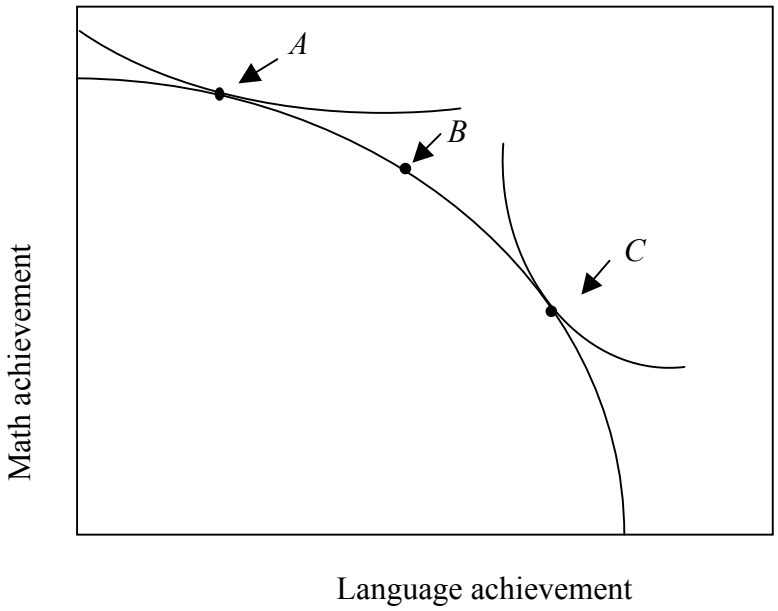

Figure 2--3 Tradeoffs between Average Achievement of Students Living in Neighborhoods 1 and 2, with and without Choice

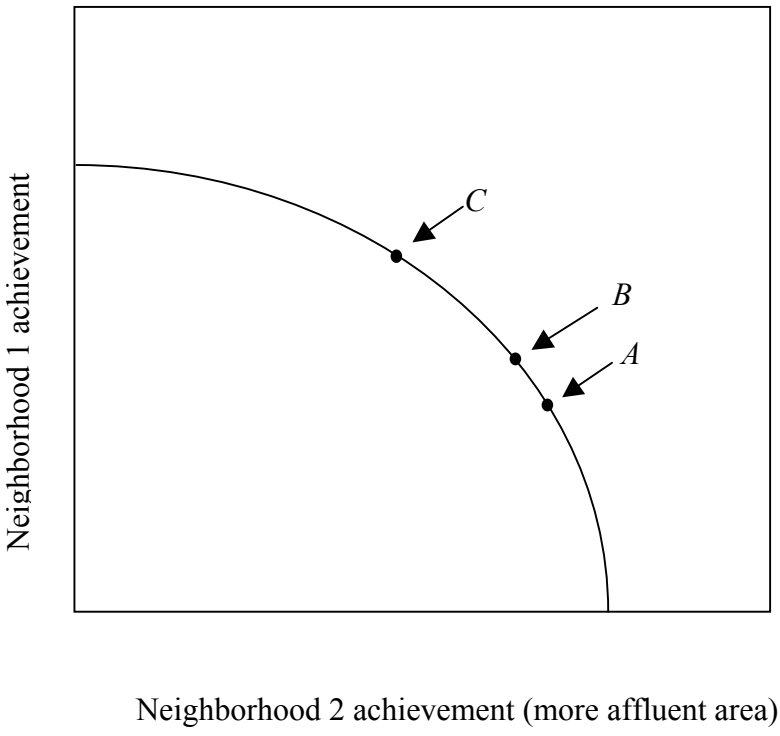

Endnotes

1. The textbook example of a negative externality is a firm that pollutes water, worsening outcomes for both firms and consumers who are located downstream.

2. These curves are U-shaped, indicating that if a child's achievement in one domain is very low, the family would not be willing to see the child's achievement in that domain fall any further unless achievement in the other domain rose by a very large amount to compensate.

3. In addition, in some important senses, teachers' unions affect districts' spending decisions. For instance, see Julian R. Betts, Kim Rueben, and Anne Danenberg, *Equal Resources, Equal Outcomes? The Distribution of School Resources and Student Achievement in California* (San Francisco: Public Policy Institute of California, 2000). These authors find examples of teachers' collective bargaining agreements in major California school districts that set limits on class size, that restrict how teachers are allocated across schools, and of course that set teachers' pay. As a co-monopolist, teachers' unions have distinct incentives to place restrictions on both the price and quantity of teachers, which on the whole will increase the cost of education beyond what would be seen in a fully competitive system with many independent producers.

4. See for instance Caroline Hoxby, "Peer Effects in the Classroom: Learning from Gender and Race Variation," Working Paper 8502 (Cambridge, Mass.: National Bureau of Economic Research, 2000); Eric A. Hanushek and others, "Does Peer Ability Affect Student Achievement?" Working Paper 8502 (Cambridge, Mass.: National Bureau of Economic Research, 2001); Julian R. Betts, Andrew Zau, and Lorien Rice, *Determinants of Student Achievement: New Evidence from San Dieg*, (San Francisco: Public Policy Institute of California, 2003).

5. I thank Rick Hess for this suggestion.

6. This simple example assumes something like a zero-sum game, in which students in more affluent areas can only lose from school choice. This is certainly a possibility that policymakers must consider when designing school choice systems that must meet a political litmus test. What is missing from this example is the possibility that school choice could enhance the quality of all schools in all neighborhoods, through the competitive effects discussed earlier.

[7]. Recent experience in California provides a cautionary tale on these strong forms of accountability that contain financial or other tangible rewards and sanctions. The high school class graduating in 2004 was supposed to be the first subject to the California High School Exit Exam requirement, but the requirement was waived after legislators learned that a large minority of students in the class of 2004 had not passed the test by spring 2003. Instead, the state decided that the class of 2006 will be the first class required to pass this test to graduate from high school. Similarly, although in 1999 and 2000 California passed legislation to introduce financial bonuses for schools as well as a generous pay incentive for both teachers and administrators at schools that either had high test scores or large improvements in test scores, financial cutbacks forced the abandonment of these programs after a single year of operation. Further, the scholarship program was discontinued after 2002.