Optimizing School Choice

John Mackenzie

Dept. of Food & Resource Economics, University of Delaware, Newark DE 19717
and Christina School District Board of Education, 83 East Main St., Newark DE 19711
e-mail: *johnmack@udel.edu*

Abstract

The Gale-Shapley algorithm can provide a stable, efficient match of students with their preferred schools in a school choice program. A GS match with students in the "proposing" role will maximize aggregate satisfaction for students, and will be immune to strategic misstatement of preferences by students. However the efficiency of the match is dissipated when students' preferences exhibit a high degree of conformity. Such bandwagon behavior is symptomatic of positional competition for schools. Positional competition will dissipate the potential benefits to students from a school choice program, and may worsen the disparities between successful and unsuccessful schools.

Optimizing School Choice

Abstract

The Gale-Shapley algorithm can provide a stable, efficient match of students with their preferred schools in a school choice program. A GS match with students in the "proposing" role will maximize aggregate satisfaction for students, and will be immune to strategic misstatement of preferences by students. However the efficiency of the match is dissipated when students' preferences exhibit a high degree of conformity. Such bandwagon behavior is symptomatic of positional competition for schools. Positional competition will dissipate the potential benefits to students from a school choice program, and may worsen the disparities between successful and unsuccessful schools.

Introduction

The market critique of public education maintains that conventional school assignment plans give schools local monopolies over students, so that schools have insufficient incentives to be efficient. The premise is that a choice program, in which parents get to select the public schools their children attend, will improve the overall quality of public education by forcing schools to compete for students. This paper explores a set of theoretical issues relating to the design of equitable and efficient school choice programs.

Instead of simply assigning students to schools based on fixed geographic feeder patterns, a growing number of public school systems now allow parents to choose the schools their children attend, subject to school capacities and other constraints. Proponents of school choice suggest that parents who actively choose their children's schools will remain more involved in schools, and will help maintain local political support for the school system. A choice program can also give school officials reliable feedback on relative levels of community support for the various schools they manage.

The standard theoretical arguments in favor of school choice are clear and compelling, but the empirical evidence that choice programs actually improve overall student achievement is largely insignificant (Cullen *et al,* 2003). Cullen *et al.* (2000) found that students in the Chicago Public School System who chose another school had a 7.6 percent higher rate of graduation than students who remained in their local school, but Levitt (2004) argues that the students who opted out of their local schools were already predisposed to succeed (were more motivated, had more involved parents), and that the exodus of these students from local schools likely reduced the local schools' graduation rates since the remaining students had fewer success-oriented peers to motivate them.

Districts with choice programs may see significant improvements in mean student test scores at highly sought-after schools, but there is little evidence to date that school choice programs yield *system-wide* improvements in student performance. In essence, school choice is not yet proved to be a positive-sum game, and skeptics argue that it is probably a zero- or even negative-sum game.

Recent experience with school choice programs shows that community consensus about the relative qualities of alternative schools causes some schools to be greatly oversubscribed while others are undersubscribed. Thus a successful choice program will generally require some formalized system for matching as many students as possible to their preferred schools, subject to school capacities, geography and cost constraints and perhaps racial/ethnic or socioeconomic diversity considerations.

Formally, the school choice problem can be abstracted to a type of two-sided match problem in graph theory (Arney and Bumcrot) or game theory (Shapley and Shubik). The first part of this paper reviews the history and mathematical properties of an efficient match algorithm first introduced in 1952 to match US medical school graduates to hospital internships, and first formalized in simplified form by Gale and Shapley in 1962. This algorithm has some desirable properties lacking in the match algorithms used in Boston, Seattle and some other school systems.

The second part of this paper addresses inefficiencies that arise when parents exhibit bandwagon behavior, all seeking to get their children into the "best" school. Such behavior suggests that schools are positional goods (Hirsch), valued for their perceived quality *relative to each other* rather than for their absolute quality. The essential feature of positional goods is that, besides conferring utility upon the owner, they impose disutility upon non-owners. This paper demonstrates how highly conformist preferences impair the efficiency of the Gale-Shapley two-sided match algorithm (and any other match algorithm as well), and derives some cautionary implications for school choice programs from this.

Although this paper is basically a synthesis of two pieces of economic theory, I have attempted to minimize the economic jargon, explain whatever jargon remains, and avoid boring examples. The paper does not argue for or against school choice *per se*. It provides a more formal context for analyzing some equity and efficiency issues in the school choice debate.

Two-Sided Match Problems

For over a century, US hospitals have recruited graduating students from US medical schools to work and receive clinical training as interns. US medical schools did not graduate enough medical students to fill every internship position, so the hospitals found themselves competing for interns. Prior to 1950, hospitals resorted to high-pressure recruiting practices, such as “exploding” offers to candidates (the offer expired within 24 or 48 hours if not accepted) and contracting for internships more than two years ahead of the student’s completion of medical school. Successive agreements to adhere to uniform appointment dates for interns failed due to widespread cheating. This poorly-structured placement process, known as the “Scramble,” was clearly a negative-sum game with significant potential for regret on both sides.

In 1951, the Association of American Medical Colleges tested and subsequently implemented the intern “Match” system that is in use today. Under the Match, after interviewing at various hospitals, each graduating medical student submits to the National Resident Matching Program (NRMP) a preference-ranked list of hospitals she is willing to work at, and after reviewing student credentials, each hospital submits to the NRMP a preference-ranked list of graduates it is willing to hire. The NRMP performs a complete matching based on these preferences. All Match participants obligate themselves at the outset to contract with whatever match they are assigned.

Between 1952 and 1997, the Match algorithm treated each hospital's rank order list as a series of provisional job offers to candidates. In the first round, each hospital’s n internship positions were offered to the n highest-ranked candidates on its list. Each candidate receiving one or more offers provisionally accepts her most-preferred offer. In the second round, the unfilled positions are offered to hospitals’ next-highest-ranked candidates, including candidates that were provisionally matched in the first round. The candidates receiving offers are again matched to their most-preferred offers. Any candidate who gets an offer that she prefers to her best match from a previous round is re-matched to the preferred offer, and the position she vacates is reoffered in the next round. Unfilled positions are re-offered and interns are matched or re-matched in successive rounds. As successful matches eliminate the top choices on hospitals' and candidates' lists, lower choices move up on the lists until all possible mutually-agreeable hires have been made. Since 1998 the Match has been performed the same way, but with candidates in the offering role and hospitals in the responding role.

Hospitals and candidates are free to skip the Match and contract directly. A "Scramble" still occurs each year after the Match to fill remaining vacancies and place unmatched candidates, but the sustained high levels of participation in the Match (now well over 20,000 intern and resident placements annually) over the past fifty years suggest that the system is *stable* (meaning that there will be no situation where a hospital and a candidate both prefer each other to the match they got) and thus Pareto-efficient (meaning you can't give any participant a more-preferred match without giving another participant a less-preferred match).

Gale and Shapley published a formalized description of this matching algorithm in 1962, and subsequent authors have explored its properties further. The simplest analog is the "stable marriage problem" involving a one-to-one match of agents from two sides. This illustrates most of the important properties of the Gale-Shapley (GS) algorithm, of which the intern Match is a many-to-one variant.

Suppose you are the match-maker for young people in some remote community where arranged (monogamous) marriages are the norm. This year you have a set of four boys (A,B,C,D) to match with four girls (W,X,Y,Z). You want to marry these young people off as happily as possible, and you want to avoid those sticky situations where a boy and a girl are married to other people but prefer each other. So you ask each young person to tell you his or her utility or anticipated satisfaction from each possible match (Table 1).

Now suppose you follow the GS algorithm and instruct each boy to propose to his first-choice girl, and each girl to provisionally accept the best proposal she gets, with the understanding that she can accept any better proposal she receives in a subsequent round. In the first round, A and C both propose to W, who accepts A. X accepts B (her second choice); Y accepts D (her last choice). Then in the second round, C, rejected by W, proposes to Z who accepts him. With pairings AW,BX,CZ,DY the four boys got utilities 4-4-3-4, while the four girls got utilities 4-3-1-3. Given the competing preferences of A and C, this is the best possible stable match for the boys, and, it turns out, the worst possible stable match for the girls!

The GS algorithm has a number of attractive properties. It will always terminate, since no boy proposes to the same girl twice, and in each round there is at least one proposal. The order in which proposals are made within a round doesn't influence the outcomes. If each agent submits a complete strict preference ordering for all possible mates, the match of one side to an equal-sized other side will be complete. The matches are stable, meaning that there are no destabilizing "blocking pairs" where a boy and a girl matched to others both prefer each other instead. (A and W are a blocking pair since they would both reject anyone else for each other.) A boy might prefer another girl to the one he matched to, but she would have rejected him in a prior round, and a girl might prefer another boy to the one she matched to, but he matched to a girl he preferred more. The match is thus Pareto-efficient, since no agent can be reassigned to a mate he/she prefers more without someone else being assigned to a mate she/he prefers less.

If the girls do the proposing under the same preferences, the GS algorithm yields a different stable, Pareto-efficient match. W and Y both propose to A, who accepts W. C accepts X (his third choice); D accepts Z (his second choice). Then in the second round, Y, rejected by A, proposes to B. With pairings AW,BY,CX,DZ the four boys got utilities 4-3-2-3, while the four girls got utilities 4-4-3-4. Given the competing preferences of W and Y, this is the best possible stable match for the girls, and the worst possible stable match for the boys. In fact, Gale and Shapley proved that under any set of strict agent preferences, *the GS match algorithm produces the stable Pareto-efficient match that maximizes net benefits to the proposing agents and therefore minimizes net benefits to the responding agents.*

The following extreme case illustrates some principles that derive from this theorem. Suppose the agents have the strict preferences shown in Table 2, where the boys all state identical (non-cooperating) preferences while the girls rank the boys differently. If boys propose and girls respond, all four boys propose to W first, who accepts A. The remaining three boys propose to X, who accepts B. The remaining two propose to Y, who accepts C, and Z accepts D. The boys get utilities 4-3-2-1 while the girls get utilities 4-4-4-4. If this is the best possible stable match for the proposing boys and the worst stable match for the responding girls, then it is the *only* stable match in the core of the game. (And it is the same match that results from girls-propose/boys-respond.) Here the strictly non-cooperating pattern of boys' preferences shrinks the core to a single match, and the girls' cooperating preferences capture the entire utility surplus. In fact, there is no match algorithm that yields anything but a permutation of the 4-3-2-1 utility allocation for the boys.

Strategic incentives under the GS algorithm are ambiguous. In the non-cooperative boys' game above, if the boys are aware of each others' true preferences, one boy's max-min strategy of being guaranteed his second-choice girl (X) by listing her as his first choice will only work if the other boys all list their true preferences. In fact, Roth (1984) has proved that "No stable matching procedure exists for which it is a dominant strategy for all agents to state their true preferences," but also "The [GS] procedure gives no [responding agent] any incentive to misrepresent his true first choice," and "The [GS] procedure makes it a dominant strategy for all [proposing agents] to state their true preferences." Thus the GS algorithm is as immune to strategic misrepresentation of preferences as any stable match algorithm can be, with only respondents' second- and lower-ranked choices subject to possible strategic misrepresentation.

It is not always obvious which side in a match situation has the advantageous proposing role. Gale and Shapley discuss the case of college admissions, where student applicants might seem to have the proposing role but don't. In the usual spring admissions cycle, colleges make offers to their top-ranked applicants based on yield projections, and wait-list some acceptable next-best applicants to insure an adequate yield, and each applicant accepts her most-preferred offer. The result is that the colleges get the best students that are attainable for them. The best offer that many applicants receive will be their second- or lower-ranked college, and they will displace some wait-listed applicants for whom that college was a first choice. In the fall early-decision cycle, however, applicants can take the proposing role by offering to attend their first-choice colleges if accepted, and the colleges assume the responding role.

The advantage to the proposing side is a form of institutional bargaining power, although it may not have much practical significance. In the case of the intern Match, a coalition of medical residents suspected that the hospital-proposing rule might be helping hospitals collude in holding down intern salaries, and they got the Match switched from hospital-proposing to candidate-proposing in 1998. The NRMP had several prior years of matches recalculated under the candidate-proposing rule to test the effect of this rule change. The simulated intern-proposing matches turned out to be virtually identical to the actual hospital-proposing matches (Roth, 2002). In this case, the core set of stable intern matches was so small that there was no significant economic surplus inherent in the proposing role. Thus the Match *per se* has no demonstrable influence on intern salaries. A likelier explanation for low intern salaries is noted in the next section.

In practice, matching processes such as courtship, college admissions or the intern Match typically involve prior recruiting, information-gathering and signaling between sides that help clarify participant preferences and identify the best options that may be attainable for each participant. Because of information and transactions costs, prospective college students and medical interns each apply to a small number of possible programs, and programs only rank candidates who actually apply. Thus the core set of stable matches may be relatively sparse.

Ergin and Sonmez (2003) demonstrate that Boston's school-proposing match is likely to motivate strategic misrepresentation of preferences by student-respondents, and argue that Boston should switch to a GS student-proposing match. In Boston's current system--which is used with minor modifications in Seattle, Minneapolis and other cities--each school determines a priority ranking of applicants: siblings of current students and living in the school's walk zone, then siblings of current students living outside the walk zone, then other children within the walk zone, then other children outside the walk zone; children are ranked within each category by random lottery. Each student submits her preference ranking of the schools she has applied to. In the first round of the match, only students' first choice requests are considered: schools offer seats to these students in order of priority, and may run out of seats before they can accommodate all first-choice requests. In each subsequent *n*th round, only unplaced students' *n*th choices are considered, and they are offered available seats in order of priority. This system essentially treats students' rankings of schools as bids, so that a student who didn't get her first choice will be outbid at her second-choice school by all students who ranked that school first, outbid at her third-choice school by all students who ranked that school first or second, etc., and may well wind up at her least-preferred school. This system creates strategic incentives for students to misstate their true preferences and makes the first-choice bid critically important. A risk-averse satisficing (good-enough) strategy under the Boston system would be for students to rank their walk zone school first if it is merely adequate, which rather defeats the purpose of a choice program and conveys false information to school officials about true levels of parent satisfaction with schools.

The design of a school choice program needs to balance the preferences of students and their parents with equity and/or cost concerns that may be formalized in schools' prioritizing various categories of applicant, e.g., disadvantaged minorities, low-income students, siblings of currently-enrolled students, children within a designated walk zone, etc. A school-proposing GS match will enable schools to satisfy their priorities as applicant numbers allow, but it may be vulnerable to strategic misstatement of preferences by applicants. A student-proposing GS match will optimize overall applicant satisfaction, and it gives applicants no reason to disclose anything but their true preferences.

Positional Goods in the Match Problem

Many real-world match problems actually involve goods that are positional or have positional attributes. A positional good gives its owner some advantage that leaves others at a relative disadvantage within a finite realm. The examples discussed above all involve goods that may have some positional aspect (e.g., Ivy League colleges, prestigious teaching hospitals and "trophy" wives). The essential feature of a positional good is that it violates Pareto-efficiency by increasing the owner's utility while reducing the utility of others. By definition, positional competition cannot be win-win, only win-lose. Under ordinary concavity of utility (where the utility cost of a loss exceeds the utility gain from the equivalent win), it is unlikely to satisfy any weak (e.g., Hicks-Kaldor "winners could compensate the losers") social efficiency criterion either.

Positional goods, along with "status goods," "snob goods" and "bandwagon goods," all involve interdependencies among different consumers' utility functions. Veblen's *Theory of the Leisure Class* (1899) describes competition for social status via "conspicuous consumption of valuable goods" that provide "reputability to the gentleman of leisure." In some cases the practical function of a status good becomes entirely irrelevant, and its true function is simply to advertise its own expense. A $20,000 Rolex doesn't really need to keep more accurate time than a $20 Timex, but it does need to be readily identifiable as a Rolex.

The term "positional good" was coined by Hirsch (1976), who worried that increasing competition for positional goods would divert more and more of society's resources from productive uses, causing eventual economic stagnation. Hirsch cited examples such as rare paintings and beach-front houses—relatively functionless repositories of wealth that are in fixed supply. This is basically a neo-Ricardian world view, with wealth embodied in a finite supply of differentiated symbolic goods rather than a finite supply of land with differentiated productivity.

It is not immediately obvious that pure status goods are socially inefficient, or that their consumption should be taxed at higher rates or otherwise discouraged. People are free to opt out of competition for status, as many psychologists encourage us to do. If we recognize envy as a true psychological cost, should we therefore prohibit some people from buying nice things just because it makes other people jealous? Perhaps jealousy is a very fitting punishment for the jealous! In a counter-argument to Hirsch, we might argue that envy is a necessary motivator of continued economic growth in wealthy economies, replacing the traditional motivators such as hunger and cold that drove most economic activity for millennia.

Kahneman *et al.* have shown that income itself is a status good. They asked survey respondents to choose which of two imaginary worlds they would prefer to live in, where prices in the two worlds are the same, and World A is simply richer than World B:

World A: You earn $110,000 per year; the average salary is $200,000

World B: You earn $100,000 per year; the average salary is $85,000.

Classical theory predicts that people will prefer the higher absolute wealth of A, but most survey respondents opt for the higher *relative* wealth of B even though it means a lower real standard of living (Frank 1997, 2003; Carlsson *et al.*). This result is consistent with longitudinal survey research in various developed nations showing that rising incomes generally do not translate to rising levels of happiness (Easterlin 1974, 1995).

A second category of positional goods imposes more tangible social costs on others than simple jealousy. For example, sport-utility vehicles typically weigh a lot more and ride higher than cars. This probably doesn't affect the rate of vehicle collisions much, but it does make SUV's relatively safer in highway collisions, and thus makes cars relatively less safe. Of the 5,259 highway fatalities resulting from SUV-car collisions in the US in 1996, 81 percent were occupants of cars (National Highway Traffic Safety Administration, 1997). This is a zero-sum game in which people in SUV's have improved their chances of surviving a collision by transferring risk to automobile occupants.

There are many similar examples. If a few spectators stand up in the front rows at a sporting event or concert to see a little better, everyone behind them has to stand up to see at all. This is a negative-sum game because everybody is less comfortable standing rather than sitting, and nobody's view is improved. Similarly, if a few NFL players or professional cyclists can get away with using steroids, other athletes will need to use them to remain competitive, injuries will increase and the long-term health of athletes will be impaired. In these contexts, positional effects are analogous to externalities. Garrett Hardin's overgrazed common is another negative-sum competition for position, where each individual's incentive to keep adding his sheep to the finite common fails to account for the collective costs of overgrazing that leaves everyone's sheep hungry and stunted. In such cases, a Pigovian tax or other policy mechanism may be justified to correct individual incentives and improve social efficiency.

Education can have obvious status good characteristics, and may have more substantive positional aspects as well (Adnett and Davies). Some univerrsities enjoy an entrenched prestige that has persisted for generations. The prestige is self-perpetuating because it keeps attracting top-quality students and faculty who maintain the academic standards for which these universities are known. The same principle applies to internships at prestigious teaching hospitals. Students compete for entry into these programs because their prestige signals the quality of graduates to prospective employers and spouses. In fact, a status hierarchy is often discernible at every level of education. Parents may compete to have their children attend elite high schools that will signal the quality of their children to admissions officers at prestigious colleges, or elite middle schools that will signal a child's quality to elite high schools. Really ambitious parents will want to get a jump on things by enrolling their children in the most prestigious day-cares, pre-schools and kindergartens.

It is logical for institutions to exploit their endowment of prestige, charging people a high price for access to it, and the high price may itself reinforce the institution's prestige. This may explain why college tuitions in the US have steadily risen twice as fast as the CPI for the last 25 years, and why teaching hospitals can get top interns to work 80+ hours a week for less than $40,000 per year—equivalent to little more than minimum wage. The cost of your education signals both its (your) quality and the level of your commitment to it (even if you majored in beer), just as the amount you spend on a bouquet of flowers signals your esteem for your dinner hostess (even if you secretly dislike her).

A good's positional value depends on a broad social consensus regarding its relative position in a preference hierarchy, and the consensus typically exhibits substantial inertia. In effect, school A is better than school B simply because everyone agrees it is. You might do your own research and discover that school B would actually be better for your child, but you need to consider the social cost to her of going to the "wrong" school, and you would not send her to B unless its relative advantages for her exceed these social costs. More typically, when the social hierarchy of public schools can be readily decoded from a handful of observable metrics (e.g., racial/ethnic composition as a subcode for poverty, published mean scores on standardized tests) there is little incentive for parents to research school quality more deeply.

In the context of public school choice, a high degree of conformity in preferences is problematic because it generates bandwagon behavior where everyone applies to the same schools in the same preference order. The second match example in the previous section illustrated the consequences of identical preferences: even if the boys are given the advantageous proposing role, conformity of their preferences dissipates the advantages of choice, collapses the core toward the respondent girls' optimum, and establishes a clear hierarchy of winners and losers among the boys.

The principal conclusion of this paper is that positional competition for schools will dissipate many of the potential benefits of a student-optimal school choice program. If schools are strongly positional, then most of the utility gains that students could obtain from an efficient student-proposing match will disappear, and the match will simply conform to whatever equity and/or cost priorities the responding school system has defined. It won't matter much whether the algorithm is student-proposing or school-proposing. In fact, positional competition will undermine the social efficiency of any match algorithm. The match will simply pick winners who will continue to support winning schools, and losers who will not support losing schools. Position is self-reinforcing. Achievement gaps between schools will persist.

The challenge for school officials is to insure that a choice program doesn't degenerate into positional competition for schools. We want parents to stay involved in making their children's current schools better, rather than focusing their efforts on getting their children into better schools. Unfortunately, there are no obvious policy prescriptions to counter positional competition in school choice.

While it is more or less axiomatic for economists that more information results in smarter choices, this does not imply that more information reduces positional competition. One possible antidote to bandwagon behavior by parents might be to require that parents of choice applicants attend an open house at each of several schools, so that they are better informed about all the schools they list in their children's choice applications and can provide substantive justifications for the choices they make. Open houses and the like may give undersubscribed schools a chance to showcases their successes and attract more applicants. Or they may simply enhance the positional sensitivities of parents and reinforce the positional advantages of oversubscribed schools.

A more promising tactic may be for districts to improve the efficiency of choice programs by encouraging schools to develop comparative advantages in particular programs or subjects. Schools could develop different specialized programs in arts, languages, etc., to address particular interests, abilities or needs of children. Undersubscribed high-poverty schools might be given extra funding to support smaller class sizes, extended day programs and the like.

Summary and Conclusion

The Gale-Shapley algorithm is known to be a stable, efficient matching procedure appropriate for matching students with schools in a school choice system. A GS match with students in the proposing role will maximize aggregate satisfaction for students, and will be immune to strategic misstatement of preferences by students. The surplus satisfaction obtained from the GS match is dissipated when the proposing side exhibits conformity of preferences.

A school choice program may elicit positional competition for schools that would be manifested in a high degree of conformity in applicants' stated preferences for schools. A choice program that supports this positional competition will provide less satisfaction to participants, and will likely increase the disparities between successful and unsuccessful schools.

The concept of public education as a positional good is quite antithetical to Dewey's vision of education as America's great democratizing force, and many education leaders are understandably reluctant to put public schools in a positional competition arena where a few elite schools may thrive while the rest may fail. A school choice program will be most beneficial when it encourages well-informed, constructive parent engagement in all schools.

Table 1: Ordered Preferences for Possible Marriage Partners

	----------------BOYS---------------				---------------GIRLS---------------			
Utility	**A**	**B**	**C**	**D**	**W**	**X**	**Y**	**Z**
4	W	X	W	Y	A	C	A	D
3	X	Y	Z	Z	B	B	B	C
2	Y	Z	X	W	C	A	C	B
1	Z	W	Y	X	D	D	D	A

Table 2: Ordered Preferences for Possible Marriage Partners (Non-Cooperative Boys)

	---------------BOYS---------------				---------------GIRLS---------------			
Utility	**A**	**B**	**C**	**D**	**W**	**X**	**Y**	**Z**
4	W	W	W	W	A	B	C	D
3	X	X	X	X	B	C	D	A
2	Y	Y	Y	Y	C	D	A	B
1	Z	Z	Z	Z	D	A	B	C

References

Adnett, Nick and Peter Davies. 2002. Education as a positional good: implications for market-based reforms of state schooling. *British Journal of Educational Studies* 50(2):189-205. (June 2002)

Arney, C. and R. Bumcrot. 1996. *Principles and Practice of Mathematics.* The Consortium for Mathematics and its Applications (COMAP). Springer-Verlag, NY. (Chapter 5: Graphs and Algorithms).

Carlsson, F., O. Johansson-Stenman and P. Martinsson. 2003. Do you enjoy having more than others? Survey evidence of positional goods. Working Papers in Economics #100 (May 2003) Department of Economics, Goteborg University, Sweden.

Cullen, Julianne B., Jacob, Brian and Levitt, Steven D. 2000. The impact of school choice on student outcomes: an analysis of the Chicago public schools. NBER Working Paper No. W7888. National Bureau for Economic Research, Washington DC. (September 2000)

Cullen, Julianne B., Jacob, Brian and Levitt, Steven D. 2003. The effect of school choice on student outcomes: evidence from randomized lotteries. NBER Working Paper No. W10113 (November 2003)

Easterlin, R.A. 1974. Does economic growth improve the human lot? In David, P. and M. Reder (eds.) *Nations and Households in Economic Growth*: *Essays in Honor of Moses Abramowitz.* Academic Press, NY.

Easterlin, R.A. 1995. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization* 27:35-47.

Ergin, Haluk and Tayfun Sonmez. 2003. Games of school choice among Boston students. Staff paper, Dept. of Economics, Princeton University. (April 2003)

Frank, Robert H. 1997. The frame of reference as a public good. *Economic Journal* 107:1832-1847.

Frank, Robert H. 2003. Are positional externalities different from other externalities? Draft for presentation at "Why Inequality Matters: Lessons for Policy from the Economics of Happiness" The Brookings Institution, Washington, DC, June 4-5, 2003

Gale, David and Lloyd Shapley 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69:9-15. (January 1962)

Hardin, Garrett. 1968. The tragedy of the commons. *Science* 162:1243-1248

Hirsch, Fred 1976. *The Social Limits to Growth,* Routledge & Kegan Paul, London.

Kahneman, Daniel, E. Diener, and N. Schwartz (eds.) 1998. *Understanding Well-Being: Scientific Perspectives on Enjoyment and Suffering*, Russell Sage, NY.

Levitt, Steven D. 2004. The economics of education. NBER Reporter: Research Summary, Winter 2004

National Highway Traffic Safety Administration 1997. *Traffic Safety Facts 1996: A Compilation of Motor Vehicle Crash Data from the Fatality Analysis Reporting System and the General Estimates System.* DOT HS 808 649, Washington, DC: U.S. Department of Transportation; National Center for Statistics and Analysis, December, 1997. (Chapter 3, p. 64, Table 37).

Roth, Alvin E. 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy* 92(6):991-1016. (December 1984)

Roth, A. and M.A. Sotomayor 1990. *Two-sided Matching: a Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, Cambridge, England.

Shapley, L.S. and M. Shubik. 1972. The assignment game I: the core. *International Journal of Game Theory* 1:111-130

Veblen, Thorstein. 1899. *The Theory of the Leisure Class: an Economic Study of Institutions*.