

PRACTICAL MARKET DESIGN: FOUR MATCHES[†]

The New York City High School Match

By ATILA ABDULKADİROĞLU, PARAG A. PATHAK, AND ALVIN E. ROTH*

We assisted the New York City Department of Education (NYCDOE) in designing a mechanism to match over 90,000 entering students to public high schools each year. This paper makes a very preliminary report on the design process and the first year of operation, in academic year 2003–2004, for students entering high school in fall 2004. In the first year, only about 3,000 students had to be assigned to a school for which they had not indicated a preference, which is only 10 percent of the number of such assignments the previous year.

New York City has the largest public school system in the country, with over a million students. In 1969 the system was decentralized into over 30 community school districts. In the 1990s, the city began to take more centralized control (Mark Schneider et al., 2000), and in 2002, a newly reorganized NYCDOE began to reform many aspects of the school system.

In May 2003, Jeremy Lack, then the NYCDOE Director of Strategic Planning, contacted one of us for advice on designing a new high-school matching process. The NYCDOE was aware of the matching process for American physicians, the National Resident Matching Program (Roth, 1984; Roth and E. Peranson, 1999). They wanted to know if it could be appropriately adapted to the city's schools. The three authors of the present paper (and, at several crucial junctures, also Tayfun Sönmez) advised (and often convinced) Lack, his colleagues (particularly Elizabeth Sciabarra and Neil Dorosin), and the DOE's software vendor, about the design of the match.

[†] *Discussant*: Paul Milgrom, Stanford University.

* Abdulkadiroğlu: Department of Economics, Columbia University, New York, NY 10027; Pathak and Roth: Harvard Business School and Department of Economics, Harvard University, Cambridge, MA 02138.

I. The Prior (2002–2003) New York City Matching Procedure

There are seven specialized high schools in New York City whose places are allocated by entrance exam (one by auditions). Rising high-school students (mostly 8th-graders, but some 9th-graders) could also apply to up to *five* other programs, by ranking them on a preference list. (Different high-school programs, with separate applications and admissions, are referred to here, interchangeably, as schools or programs. There are over 500 programs.) Just over 50 percent of students in 2002 applied to the maximum allowable five programs. Schools receiving a student's application saw her list of preferences (and could see where they ranked on the list). How they processed applications varied by program type.

Unscreened programs admit students by lottery. *Zoned* schools give priority to students from the local neighborhoods, and many positions were filled this way. (One impetus for increasing school choice was to make sure students who lived in disadvantaged neighborhoods were not automatically assigned to disadvantaged schools.)

Screened programs can evaluate students individually. *Educational Option* (EdOpt) programs also can evaluate students individually for half their seats, subject to the restriction that 16 percent be allocated to students who were rated top performers in a standardized English Language Arts exam, 68 percent to middle performers, and 16 percent to lower performers. The other half of the seats are allocated by lottery, according to the same distribution of test scores. EdOpt programs were also subject to a special rule: any student with an ELA score in the top 2 percent would be automatically admitted if the program was ranked first on the student's list. (Screened and EdOpt schools

could use whether a student ranked them first as an admissions criterion for any student.)

Subject to these constraints, schools could decide which of their applicants to accept, place on a waiting list, or reject. Each applicant received a letter with the decisions of the schools to which she had applied, and applicants were required to accept no more than one offer, and one wait-list. This process was repeated: after the responses to the first letter were received, schools with vacant positions could make new offers, and after replies were received, a third letter with new offers was sent. New offers did not necessarily go to wait-listed students in a fixed order. Students not assigned after the third step were assigned to their zoned schools, or assigned via an administrative process. There was an appeals process, and an “over the counter” process for assigning students who had changed addresses or were otherwise unassigned before school began.

Three rounds of processing applications to no more than five out of more than 500 programs by almost 100,000 students was insufficient to allocate all the students. That is, this process suffered from *congestion* (Roth and Xiaolin Xing, 1997): not enough offers and acceptances could be made to clear the market. Only about 50,000 students received offers initially, about 17,000 of whom received multiple offers. When the process concluded, approximately 30,000 students had been assigned to a school that was nowhere on their choice list.

Three features of this process particularly motivated NYCDOE’s desire for a new matching system. First were the approximately 30,000 students not assigned to a school they had chosen. Second, students and their families had to be strategic in their choices. Students who had a substantial chance of being rejected by their true first-choice school had to think about the risk of listing it first, since, if one of their lower-choice schools took students’ rankings into account, they might have done better to list it first. (The 2002–2003 *Directory of the New York City Public High Schools* advises students (p. x) that, when ranking schools, they should “... determine what your competition is for a seat in this program.” A similar problem occurs in Boston schools (see Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005). Third, schools were also strategic: a substantial num-

ber of schools apparently managed to conceal capacity from the central administration, thus preserving places that could be filled later.

II. Design of the New System

Initial discussions focused on whether the medical match was a good model for New York City schools, or whether another kind of clearinghouse might be more appropriate. The medical match applied to schools would be a two-sided model in which both schools and students have preferences, with the object of implementing a *stable* assignment, that is, an efficient assignment such that no school and student not matched to one another would both prefer to be. Thus, the question was, are the students the only real players in the system, with choices by schools merely a device for allocating scarce spaces? If this were the case, there might be appropriate *one-sided* clearinghouse models in which only student preferences determine efficient allocations (cf. Boston Public Schools; Abdulkadiroğlu et al., 2005).

Two things convinced us that New York City schools are a two-sided market. The first was that schools withheld capacity to match with students they preferred. Stable assignments would eliminate the main incentives for this. Second, discussions indicated that principals of different EdOpt schools had different preferences even for students with reading scores in the lowest category, with some schools preferring higher scores and others preferring students who had good attendance. If schools have different comparative advantages, allowing scope for their preferences seemed sensible. Also, the fact that school administrators gamed the system indicated they were strategic players.

The medical match employs an applicant-proposing deferred-acceptance algorithm (David Gale and Lloyd Shapley, 1962; Roth and Peranson, 1999). Ignoring for the moment the details of New York City schools, this could be applied as follows. Students and schools rank each other (schools do not see students’ preferences), and the clearinghouse processes these lists so:

- (i) Each student applies to her highest ranked school, and each school rejects unranked applicants and “holds” its highest ranked

applications (up to the number of positions it has) and rejects the rest.

- (ii) At any stage at which a student has been rejected, she applies to her next most preferred school if one remains. Each school holds its most preferred set of applications and rejects the rest.
- (iii) The algorithm stops when no rejections are issued, and each school is matched to the applicants it is holding.

No student would receive multiple offers. We discussed whether this was an unmixed benefit: students who received multiple offers in the old system might benefit from this (e.g., in increased decision time). But relatively few such students chose a school different from their indicated preferred choice, so this seemed like a bearable cost, considering that in a system without excess capacity the cost of giving some students multiple offers is that multiple students get no offers.

The next design choice was whether the algorithm should be student-proposing, or another stable mechanism (e.g., schools-proposing). A student-proposing algorithm was selected because this has the best welfare properties for students and (in sufficiently simple environments) makes it a dominant strategy for students to state true preferences (and since no alternative stable mechanism gives schools straightforward incentives [Roth and M. Sotomayor, 1990; Sönmez, 1997]).

Of course, adapting the mechanism to the regulations and customs of New York City schools involved departures from the simple algorithm (cf. Roth, 2002). Schools that allocated seats by lottery are assigned randomly generated preferences. Each (half) EdOpt program is treated as three different programs whose preferences must reflect the 16/68/16 reading score distribution. If a student ranked an EdOpt school, this was treated in the algorithm as a preference for one of the random slots first, followed by a preference for one of the slots determined by the school's preferences.

The EdOpt automatic admit for top-2-percent students who choose it first could not be changed, although this adds strategic risk to the decisions of students who are eligible to use it. Another decision that makes some students not have a dominant strategy is that preference lists

were limited to a maximum of 12 schools. Over 22,000 students listed 12 in the first year, so this was a binding constraint. These choices are therefore candidates to be revisited when modifications are considered.

NYCDOE wanted students who are offered specialized-school positions also to be given an offer from a nonspecialized school. Therefore students who applied to specialized schools submit a preference list of nonspecialized schools along with all other students, and a first round of the algorithm is run with all students. Students who receive a specialized-school offer receive a letter giving them a choice between that exam school and a nonspecialized school. After they respond, capacities are adjusted, students who accepted offers are removed, and the algorithm is run again. Only after this second round are students who did not receive specialized-school offers told their assignment.

We would have preferred to integrate these two rounds into one, by having applicants include the specialized schools in their preference lists. (The two-round design creates a possibility of unstable allocations, as when a student gets an offer from a specialized school, but not from a nonspecialized school he prefers that would have had a place for him after the specialized-school students have declined places.) However, if students who are offered specialized-school places generally rank high in all schools' preferences, this may not be a big problem.

Students who were unassigned after the second round were informed of the schools with empty places and asked for another preference list of up to 12 schools. The NYCDOE felt there would be insufficient time to elicit new preferences from schools (a decision that might usefully be reviewed in the future), and so these students were ordered in a single random list that was used as the preferences for all schools in a third round of the algorithm. (This compares favorably with alternative methods of randomization.) The small number of students who remained unmatched were assigned administratively.

III. First Year of Operation

The new match matched over 70,000 students to a school on their initial choice list, an increase of more than 20,000 students compared

to how many received one of their choices the previous year. An additional 7,600 students who were unmatched based on their initial preferences were assigned to schools based on the preferences they submitted over schools that still had vacancies.

Of over 90,000 students who submitted preferences, approximately 8,000 students withdrew from the New York City public schools, and more than 2,000 remained in their current school either in a 9th-grade program or through failure to graduate, leaving approximately 3,000 students who did not receive any school they chose. This compares to the approximately 30,000 who NYCDOE reports were administratively assigned, mostly to zoned schools, the previous year.

Much of this difference is due to allowing students to rank 12 instead of five choices, and to giving each student a single offer, rather than multiple offers to some students. Interestingly, it appears that at least 3,000 more students received one of their first five stated choices than in the previous year, under the old system.

Just over 5,100 students appealed their assignments, around 2,600 were granted on a case-by-case basis. (Around 300 appeals were from students who received their first choice. Some of these may have had to do with bad information about new programs. But some may reflect the difficulty of soliciting preferences involving 13–14-year-olds.) Designing an efficient appeals process remains a priority.

Thus despite some significant first-year problems of communication and implementation, the new match seems to have achieved many of its goals.

IV. Conclusions

New York City needs more good schools. But for a given stock of school places, more students can be admitted to schools they want if the matching process is free of congestion, so that students' preferences can be fully taken into account. The new clearinghouse, organized around a stable matching mechanism, has helped relieve the congestion of the previous offer/acceptance/wait-list process

and provides more straightforward incentives to applicants.

REFERENCES

- Abdulkadiroğlu, Atila; Pathak, Parag A.; Roth, Alvin E. and Sönmez, Tayfun.** "The Boston Public School Match." *American Economic Review*, 2005 (*Papers and Proceedings*), 95(2), pp. 368–71.
- Abdulkadiroğlu, Atila and Sönmez, Tayfun.** "School Choice: A Mechanism Design Approach." *American Economic Review*, 2003, 93(3), pp. 729–47.
- Gale, David and Shapley, Lloyd.** "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 1962, 69(1), pp. 9–15.
- Roth, Alvin E.** "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 1984, 92(6), pp. 991–1016.
- Roth, Alvin E.** "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." *Econometrica*, 2002, 70(4), pp. 1341–78.
- Roth, A. E. and Peranson, E.** "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 1999, 89(4), pp. 748–80.
- Roth, A. E. and Sotomayor, M.** *Two-sided matching: A study in game-theoretic modeling and analysis*, Econometric Society Monograph Series. Cambridge, U.K.: Cambridge University Press, 1990.
- Roth, Alvin E. and Xing, Xiaolin.** "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists." *Journal of Political Economy*, 1997, 105(2), pp. 284–329.
- Schneider, Mark; Teske, Paul and Marschall, Melissa.** *Choosing schools: Consumer choice and the quality of American schools*. Princeton, NJ: Princeton University Press, 2000.
- Sönmez, Tayfun.** "Manipulation via Capacities in Two-Sided Matching Markets." *Journal of Economic Theory*, 1997, 77(1), pp. 197–204.

The Boston Public School Match

By ATILA ABDULKADİROĞLU, PARAG A. PATHAK, ALVIN E. ROTH, AND TAYFUN SÖNMEZ*

After the publication of “School Choice: A Mechanism Design Approach” by Abdulkadiroğlu and Sönmez (2003), a *Boston Globe* reporter contacted us about the Boston Public Schools (BPS) system for assigning students to schools. The *Globe* article highlighted the difficulties that Boston’s system may give parents in strategizing about applying to schools. Briefly, Boston tries to give students their first-choice school. But a student who fails to get her first choice may find her later choices filled by students who chose them first. So there is a risk in ranking a school first if there is a chance of not being admitted; other schools that would have been possible had they been listed first may also be filled.

Valerie Edwards, then Strategic Planning Manager at BPS, and her colleague Carleton Jones invited us to a meeting in October 2003. BPS agreed to a study of their assignment system and provided us with micro-level data sets on choices and characteristics of students in the grades at which school choices are made (K, 1, 6, and 9), and school characteristics. Based on the pending results of this study, the Superintendent has asked for our advice on the design of a new assignment mechanism. This paper describes some of the difficulties with the current mechanism and some elements of the design and evaluation of possible replacement mechanisms.

School choice in Boston has been partly shaped by desegregation. In 1974, Judge W. Arthur Garrity ordered busing for racial balance. In 1987, the U.S. Court of Appeals freed BPS to adopt a new, choice-based assignment plan. In

1999 BPS eliminated racial preferences in assignment and adopted the current mechanism.

I. The Current Boston Mechanism

BPS has over 60,000 students from grades K–12 in almost 140 schools in three zones: East, West, and North. During the first registration period in January, students who will be entering a new school in grades K, 1, 6, and 9 are asked to rank at least three schools in order of preference. Although most assignments are made in the first registration period, Boston has other registration periods in February, March, and April.

For elementary and middle school, parents are asked to consider schools in their zone plus five schools open to all neighborhoods. High school admissions are city-wide for 18 schools. There are also 13 high schools that require special admissions and three special-education programs that are not part of the centralized allocation process.

In 2004, at the end of the first registration period, there were about 4,800 students entering kindergarten, 4,000 entering grade 1, over 4,300 students entering grade 6, and about 4,000 entering grade 9.

Boston assigns students if possible to their first-choice school, allocating over-demanded seats by a system of priorities. First, a younger sibling has priority to attend the same school as an older sib. Next in priority for half of each program’s seats are students from the school’s *walk zone*. Not every residential location in the city has a school for which they obtain walk-zone preference. Students who live in these locations are then given priority for assignment to their first- and second-choice schools. Additional priorities are assigned by random numbers generated once for each student. After the first registration period there is no longer a walk-zone priority.

* Abdulkadiroğlu: Department of Economics, Columbia University, New York, NY 10027; Pathak and Roth: Harvard Business School and Department of Economics, Harvard University, Cambridge, MA 02138; Sönmez: Department of Economics, Koç University, Istanbul, Turkey, and Harvard University.

Within each priority class, students' random numbers determine a strict priority order. Each school has a maximum capacity determined by BPS. The Boston mechanism assigns students as follows:

Step 1.—For each school, consider the students who have listed it as their *first* choice and assign seats to these students in priority order until either no seats remain or no student remains who has listed it as first choice.

Step k.—For each school with seats still available, consider the students who have listed it as their *k*th choice and assign seats to these students in priority order until either no seats remain or no student remains who has listed it as *k*th choice.

The procedure terminates when each student is assigned a seat (or all submitted choices are considered).

If a student does not get her top choice, she may be added to a school's waiting list. Students who get their second choice go on the wait-list for their first choice. Students who get neither their first nor second choice are placed on wait-lists for both. Students who do not get any of their choices go on wait-lists for up to three choices. The priority on the wait-list is based on sibling preference, round of application, and random number. When the school year starts, if a student leaves the public-school system, the student may no longer stay on a wait-list. All wait-lists expire in January of the next school year.

During the 2002–2003 assignment process, about 11 percent of students were on wait-lists. In 2004, two major changes were introduced: caps to the size of the wait-list and active confirmation of interest in a wait-list. This year, students may go on wait-lists only until the wait-list contains 25 percent of the number of seats at the grade level in the school. Also, students already on the number of wait-lists they are entitled to according to the school choice they received must leave one list before being added to another.

At the end of the assignment process, if a student is not given any of his choices, or did not return an application, BPS assigns the student to the school closest to home that has space.

The Boston mechanism is a *priority matching mechanism* (Roth, 1991). Priority mechanisms have been used to match medical graduates to internships in several regions of the United Kingdom, starting in the 1960s. Each of these mechanisms was abandoned after being gamed by the participants. Yan Chen and Sönmez (2005) experimentally examine preference manipulation under the Boston mechanism and observe the associated welfare loss.

Priority mechanisms are common in school choice. The largest district we know of with a priority mechanism is Hillsborough County School District in Tampa-St. Petersburg, the 11th largest school district in the United States, with about 170,000 students.¹ Cambridge, Denver, Minneapolis, and Seattle also have priority mechanisms.

The idea that students and parents should be cautious in choosing their first choice is included in the reference material provided to students and parents. BPS states “for a better chance of getting your ‘first choice’ school ... consider choosing less popular schools” (*Introducing Boston Public Schools*, 2004, p. 3 [quotation marks in original]). In Seattle and Tampa-St. Petersburg, the incentives for such preference manipulation are advocated in the local press (see Haluk Ergin and Sönmez, 2005). Note that when students rank less competitive programs first, many get their stated first choice. Approximately 80 percent of students who submit preferences in the first registration period get their stated first choice in Boston. Of course, this is not necessarily their most preferred school.

II. Two Alternative Matching Mechanisms

It is costly in the Boston mechanism to list a first-choice that you do not succeed in getting because, once other students are assigned their first-choice places, they cannot be displaced even by a student with higher priority. A class of mechanisms that avoid this are *deferred-acceptance algorithms* (David Gale and Lloyd Shapley, 1962) of the kind adopted by New

¹ Often, the precise allocation rules are not publicly specified by the school districts.

York City high schools (Abdulkadiroğlu et al., 2005) and elsewhere (Roth, 2002):

Step 1.—Each student “proposes” to her first choice. Each school tentatively assigns its seats to its proposers one at a time in their priority order. Any remaining proposers are rejected.

Step k.—Each student who was rejected in the previous step proposes to her next choice if one remains. Each school considers the students it has been holding together with its new proposers and tentatively assigns its seats to these students one at a time in priority order. Any remaining proposers are rejected.

The algorithm terminates when no student proposal is rejected, and each student is assigned her final tentative assignment.

In contrast with the Boston algorithm, the deferred-acceptance algorithm assigns seats only tentatively at each step, so students with higher priorities may be considered in subsequent steps. Consequently it is stable in the sense that there is no student who loses a seat to a lower-priority student and receives a less-preferred assignment. Moreover all students prefer their outcome to any other stable matching (Gale and Shapley, 1962), and the induced student-optimal stable mechanism is dominant-strategy incentive-compatible (Roth, 1982a). (Unlike in New York City, the schools are not strategic players in Boston, as the priorities are set centrally.) If the intention of the school board is that priorities be “strictly enforced,” this mechanism is a leading candidate.

However, if welfare considerations apply only to students, there is tension between stability and Pareto optimality (Roth, 1982a). If priorities are merely a device for allocating scarce spaces, it might be possible to assign students to schools they prefer by allowing them to trade their priority at one school with a student who has priority at a school they prefer. The following *top trading cycles* (TTC) mechanism creates a virtual exchange for priorities:

Step 1.—Assign counters for each school to track how many seats remain available. Each student points to her favorite school, and each school points to the student with the highest priority. There must be at least one cycle. (A

cycle is an ordered list of distinct schools and students (student 1 - school 1 - student 2 - ... - student k - school k) with student 1 pointing to school 1, school 1 to student 2, ... , student k to school k , and school k pointing to student 1.) Each student is part of at most one cycle. Every student in a cycle is assigned a seat at the school she points to and is removed. The counter of each school is reduced by 1, and if it reaches zero, the school is removed.

Step k.—Each remaining student points to her favorite school among the remaining schools, and each remaining school points to the student with highest priority among the remaining students. There is at least one cycle. Every student in a cycle is assigned a seat at the school she points to and is removed. The counter of each school in a cycle is reduced by 1, and if it reaches zero, the school is removed.

The procedure terminates when each student is assigned a seat (or all submitted choices are considered).

This version of the TTC mechanism was introduced by Abdulkadiroğlu and Sönmez (2003) and is an extension of Gale’s “top trading cycles mechanism” described in Shapley and Herbert Scarf (1974). Many properties of TTC carry over to school choice, including Pareto efficiency (Shapley and Scarf, 1974) and dominant-strategy incentive compatibility (Roth, 1982b). Variations of this procedure can also be considered which may reduce instability (e.g., Onur Kesten, 2005). See also the recent design of a kidney exchange clearinghouse (Roth et al., 2004, 2005).

III. Design Considerations

Unlike in New York (Abdulkadiroğlu et al., 2005), students’ priorities at schools are not set by schools, but by the central administration. There does not seem to be any issue of individual schools gaming the system in Boston. Therefore, it is natural to ask whether the benefits that “stable” matching produces in New York have parallel benefits in the different situation in Boston, and if not, whether the welfare improvements that might be available from a TTC-like mechanism should be considered.

At a public meeting of the Boston School

Committee in October 2004 we were asked for advice about how to think about this. We replied with the question “Would anyone mind if two students who each preferred the school in the other student’s walk zone were to trade their priorities and enroll in those schools?” If this is not desirable (e.g., because of transportation costs, or because walk-zone priorities reflect a public good that results when parents walk children to school, or because lawsuits might follow if a child is excluded from a school while another with lower priority is admitted), then stable matchings would efficiently combine student preferences with priorities. But if helping the students this way is worth whatever transportation and other costs might be incurred, then only the students’ preferences need to be taken into account and a TTC-like mechanism might be more appropriate.

IV. Recent Developments

In December 2003, the Boston School Committee initiated an evaluation of all aspects of student assignment. The final task-force report recommends changing the student assignment algorithm. The task force observed that, even though students can select three schools, many children do not get any of their picks because, if a parent and student choose three popular schools and do not get their first choice, they may also miss their second and third choice.

A memorandum from Superintendent Payzant in December 2004 states that BPS plans to change the computerized process used to assign students to schools. Although the task-force report recommended that BPS adopt the TTC assignment algorithm, the School Committee is interested in simulations of both mechanisms and in understanding the extent of preference manipulation under the Boston mechanism. They are also thinking through their philosophical position on the trade-off between stability and efficiency.

REFERENCES

Abdulkadiroğlu, Atila; Pathak, Parag A. and Roth, Alvin E. “The New York City High

School Match.” *American Economic Review*, 2005 (*Papers and Proceedings*), 95(2), pp. 364–67.

Abdulkadiroğlu, Atila and Sönmez, Tayfun. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 2003, 93(3), pp. 729–47.

Chen, Yan and Sönmez, Tayfun. “School Choice: An Experimental Study.” *Journal of Economic Theory*, 2005 (forthcoming).

Ergin, Haluk and Sönmez, Tayfun. “Games of School Choice under the Boston Mechanism.” *Journal of Public Economics*, 2005 (forthcoming).

Gale, David and Shapley, Lloyd. “College Admissions and the Stability of Marriage.” *American Mathematical Monthly*, 1962, 69(1), pp. 9–15.

Kesten, Onur. “Student Placement to Public Schools in the US: Two New Solutions.” Mimeo, University of Rochester, 2005.

Roth, Alvin E. “The Economics of Matching: Stability and Incentives.” *Mathematics of Operations Research*, 1982, 7(4), pp. 617–28.

Roth, Alvin E. “Incentive Compatibility in a Market with Indivisible Goods.” *Economics Letters*, 1982b, 9(2), pp. 127–32.

Roth, Alvin E. “A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom.” *American Economic Review*, 1991, 81(3), pp. 414–40.

Roth, Alvin E. “The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics.” *Econometrica*, 2002, 70(4), pp. 1341–78.

Roth, Alvin E.; Sönmez, Tayfun and Ünver, M. Utku. “Kidney Exchange.” *Quarterly Journal of Economics*, 2004, 119(2), pp. 457–88.

Roth, Alvin E.; Sönmez, Tayfun and Ünver, M. Utku. “A Kidney Exchange Clearinghouse in New England.” *American Economic Review*, 2005 (*Papers and Proceedings*), 95(2), pp. 376–80.

Shapley, Lloyd and Scarf, Herbert. “On Cores and Indivisibility.” *Journal of Mathematical Economics*, 1974, 1(1), pp. 23–28.

The Gastroenterology Fellowship Market: Should There Be a Match?

By MURIEL NIEDERLE AND ALVIN E. ROTH*

We are helping a task force of the American Gastroenterology Association to evaluate the current state of the (decentralized) market for gastroenterology fellows, and to assess the prospects of reorganizing it via a suitably designed centralized clearinghouse, a “match.” This market used a match from 1986 until the late 1990s. Starting in 1996, participation in the match declined precipitously, and it was formally abandoned after 1999. Consequently, the experience of this market when the match was in place, in comparison to the periods before and since, allows an assessment of the effects of the match. An analysis of how the match failed in the 1990s yields insights into the prospects for success of a new match. These events offer economists a rare window on how decentralized labor markets clear, and on how market clearinghouses succeed and fail.

I. The Rise and Fall of the Gastroenterology Match

A gastroenterologist, after graduating from medical school, completes three years as an internal medicine resident, and then a gastroenterology fellowship. Like many other entry-level labor markets, gastroenterology experienced “unraveling” prior to 1986, as offers were made earlier from year to year, at dispersed times, well over a year before fellowships began. Such early offers are typically also “exploding,” they do not leave candidates time to consider many other offers (Roth and Xiaolin Xing, 1994; Niederle and Roth, 2004b).

In 1986, after other attempts to halt unraveling and create a thicker and more orderly market, gastroenterology, and a number of other

specialties, were successfully organized through a centralized match, the Medical Specialties Matching Program (MSMP), which operates along the lines of the larger resident match for first-year doctors (cf. Roth, 1984; Roth and Elliott Peranson, 1999). After a period of interviewing, medical residents and gastroenterology program directors ranked each other and submitted these lists to the match. A version of a deferred acceptance algorithm (David Gale and Lloyd Shapley, 1962) produced a *stable* matching (i.e., one in which no resident and program who are not matched together would both prefer to be). But in the late 1990’s, the match itself unraveled, as positions were filled before the match was conducted.

Up to 1995, well over 300 fellowship positions were advertised annually through the match, which attracted at least 1.3 applicants per position, with a fill rate of 88 percent and higher. A planned reduction of 25–50 percent in fellowship positions over five years began in 1996, when about 300 positions were advertised. Unexpectedly, there was an even larger reduction in the number of applicants, and in 1996 only 0.9 applicants per position participated in the match, and only about 75 percent of positions were filled through the match. While the number of applicants quickly returned to excess supply, a perceived shortage of “high quality” applicants remained, and it seems that many fellowship programs had lost the confidence to wait for the match and preferred to make offers to candidates when they interviewed them. The next year, 16 percent of the positions initially advertised through the match were withdrawn, leaving only 213 positions in the match. In 1998, 60 percent of advertised positions were withdrawn, leaving only 99 positions in the match, and in 1999, the last year the match was formally conducted, only 14 positions participated.

While we know of about a hundred markets that have been organized by a stable matching mechanism, we know of only a handful that have failed (Niederle and Roth, 2004a), and so

* Niederle: Department of Economics, Stanford University, Stanford, CA 94305, and NBER; Roth: Harvard Business School and Department of Economics, Harvard University, Cambridge, MA 02138. We gratefully acknowledge the support of the NSF and helpful conversations and collaboration with Deborah Proctor.

the cause of the failure of the gastroenterology match is worth investigating. For this, historical field data can only take us so far. However, when we reproduce this market on a small scale in the laboratory (C. Nicholas McKinney et al., 2005) we can subject the market to different kinds of supply and demand shocks, under different information conditions.

Our experimental results confirm that it is hard to unravel a match even through a shock that reverses which side of the market is short. In the lab, when applicants were on the long side of the market, they eagerly accepted early offers, but programs had little incentive to make them. When applicants were on the short side and this was common knowledge, programs made early offers, but applicants preferred to wait for the outcome of the match. The feature of the 1996 market that makes a big difference in the lab is that the sudden shortage of applicants was unanticipated, and hence applicants went into the market thinking that positions would be scarce. So did fellowship programs, but they were more quickly able to discern the true state of affairs, when they did not get their expected number of applications. In this case, in the experiment, programs made early offers, and applicants accepted them. And, of course, once many programs are making early offers, and having them accepted, then many positions are withdrawn, and the attraction of waiting for the match diminishes.

This rare failure of a stable clearinghouse, following a disruption in supply and demand, also gives us an unusually clear way to assess what the clearinghouse accomplished while it was in use.

II. The Effects of a Match

A. *Timing and Market Thickness*

With the demise of the match, the market unraveled once again, and interviews for gastroenterology fellowships moved steadily earlier (Niederle and Roth, 2004a). Compared to internal-medicine subspecialties that continued to use the MSMP, the bulk of gastroenterology interviews had moved two months earlier for positions starting in 2003, and three months earlier for positions starting in 2005 (and 20 months before employment would begin). Interviews also became more dispersed. For exam-

ple, there are never as many as 70 percent of the gastroenterology programs interviewing at the same time, while the comparably large internal medicine subspecialties that continue to use the MSMP have over 70 percent of programs interviewing at the same time for several months. For gastroenterology, by the time 80 percent of programs have started interviews, more than 50 percent have already finished. These differences between gastroenterology and the subspecialties that continue to use the match are even more consequential than they appear, because, for specialties that use a match, offers do not immediately follow interviews, and candidates can consider in the match all programs for which they have interviewed. We conducted a survey of gastroenterology program directors about the timing of offers, and the replies confirm that offers closely follow interviews.

B. *The Effects on Mobility: Who Matches to Whom?*

When hiring moves increasingly far in advance of employment, it may become more difficult to gather information on candidates, or to secure reliable commitments from them. For these reasons, we suspected that unraveling would be associated with increased reliance on local networks. And when offers are exploding, candidates may be able to more readily secure prompt counteroffers from local programs than from those that would require distant interviews. We therefore examined the mobility of gastroenterologists, as they moved from their internal-medicine residency to a fellowship.

In Niederle and Roth (2003b) we tracked the 9,180 fellows who completed both a residency and a gastroenterology fellowship in the United States after 1977. Before the match, and since its demise, fellows were much more likely to stay at the hospital at which they did their residency, to remain in the same city, and in the same state, than during the match. The fact that mobility declines after the breakdown of the match makes us more confident that the increase in mobility during the match is due to the match and is not simply an increase in mobility over time. The effect of the match is bigger for large (and presumably more prestigious) hospitals, which employ more fellows from a different hospital, city, and state.

The use of a centralized match therefore affects not only the *timing* of the market, but also the outcome, who matches to whom.

III. Does a Match Affect the Terms of Employment?

In 2002, 16 law firms filed a class action lawsuit, on behalf of three former residents, seeking to represent the class of all residents and fellows, arguing that the NRMP (the match for medical residents) violated antitrust laws and was a conspiracy to depress wages. The lawsuit was against a class of defendants including the NRMP (which also operates the MSMP), other medical organizations, and the class of all hospitals that employ residents.

One way to investigate whether a match affects wages of medical fellows is to examine comparable medical subspecialties, only some of which use a match. Niederle and Roth (2003a) compare wages of the 1,148 nonmilitary U.S. fellowship programs in all internal-medicine subspecialties that require three years of prior residency. Controlling for the hospital, we find that specialties that use a match have no lower wages than those that do not. Thus it appears that, in these medical labor markets, wages are determined by factors other than whether a centralized clearinghouse is used.

One by-product of the suit is that it brought renewed attention to the fact that many entry-level labor markets have *impersonal* wages that are part of the job description, so that people hired at the same time for the same kind of position by the same firm may all begin at the same salary. Jeremy Bulow and Jonathan Levin (2003) observe that a centralized clearinghouse may promote this tendency, since positions have to be offered in the match to all desirable candidates (i.e., without knowing in advance who will fill them). They note that many labor markets that do not use a match also often have impersonal wages: they mention law, investment banking, and academia. Bulow and Levin (2003) show that a market with nonpersonalized wages tends to lower the average wage and compress the wages of applicants compared to a competitive market (see also Ulrich Kamecke, 1998).

The evidence from medical subspecialties suggests that the absence of a match may pro-

mote neither more personalized wages nor a more competitive market. The gastroenterology market became thinner after the demise of the match, since dispersed exploding offers do not allow applicants to compare multiple offers.

Reflecting these considerations, President George W. Bush signed into law, as an addendum to the Pension Funding Equity Act of 2004, legislation that included a Congressional finding that “*Antitrust lawsuits challenging the matching process, regardless of their merit or lack thereof, have the potential to undermine this highly efficient, pro-competitive, and long-standing process ...*.” The legislation goes on to “*confirm that the antitrust laws do not prohibit sponsoring, conducting, or participating in a graduate medical education residency matching program, or agreeing to do so ...*.” Following this legislation, the antitrust suit was dismissed (although legal skirmishing remains).

IV. Reconstituting a Gastroenterology Match

What issues must the American Gastroenterology Association consider, as it contemplates a new match? Whether a match is desirable has the potential to be contentious, because a move to a later, thicker, more competitive market may not be a Pareto improvement. Less competitive programs may, in the present unraveled market, be able to retain their hospital’s best medical residents, who would be more mobile in a match (cf. M. N. Ehrinpreis, 2004). However, evidence from the early years of the MSMP suggests that to start a match successfully requires substantial rates of initial participation by programs.

There are also several kinds of gastroenterology fellows, not only clinical fellows, but also basic science research and clinical research fellows. It appears that some programs may wish to hire a few research fellows (but not exclusively research fellows), but would like to fill those research positions with clinical fellows if they cannot. If so, it may be desirable to design the match to allow unfilled research positions to “revert” to clinical positions (Roth and Peranson, 1999; Roth, 2002). For the gastroenterology market, an alternative might be to have the research market operate before the clinical match.

V. Concluding Remarks

To facilitate efficiency, markets need to be thick, and many markets achieve efficiency by aggregating buyers and sellers in time (and sometimes in space). Unraveling works against this: dispersed and exploding offers make the market more like a series of bilateral encounters.

To realize the efficiencies that a thick market allows, the market needs to overcome congestion: having lots of applicants available does not help if employers only have time to consider a few of them. Prior to the start of the gastroenterology match in 1986, attempts were made to organize the market simply via a system of rules about when offers can be made, how long they must remain open, and so forth. Many markets have tried and failed to organize themselves by such rules: the problem is that they experience congestion, so that not enough offers can be processed in the available time. (By the time an offer is rejected, other candidates may no longer be available, and so employers have incentives to start making offers earlier, and to leave them open for less time, which makes the market unravel.)

Clearinghouses solve both problems: they bring participants to the market at the same time, and they overcome congestion.

To more fully understand how a wide variety of labor markets clear, we need to better understand how, and how well, other decentralized as well as centralized market institutions perform these tasks.

Added in proof:

The American Gastroenterology Association announced in June 2005 that it will reinstate a match starting in 2006.

REFERENCES

- Bulow, Jeremy and Levin, Jonathan.** "Matching and Price Competition." Working paper, Stanford University, December 2003.
- Ehrinpreis, M. N.** "Con: The Gastroenterology Fellowship Match: R.I.P." *American Journal of Gastroenterology*, 2004, 99(1), p. 7.
- Gale, David and Shapley, Lloyd.** "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 1962, 69(1), pp. 9–15.
- Kamecke, Ulrich.** "Wage Formation in a Centralized Matching Market." *International Economic Review*, 1998, 39(1), pp. 33–53.
- McKinney, C. Nicholas; Muriel, Niederle and Roth, Alvin E.** "The Collapse of a Medical Labor Clearinghouse (and Why Such Failures Are Rare)." *American Economic Review*, 2005, 95(3), pp. 878–89.
- Niederle, Muriel and Roth, Alvin E.** "Relationship Between Wages and Presence of a Match in Medical Fellowships." *Journal of the American Medical Association*, 2003a, 290(9), pp. 1153–54.
- Niederle, Muriel and Roth, Alvin E.** "Unraveling Reduces Mobility in a Labor Market: Gastroenterology with and without a Centralized Match." *Journal of Political Economy*, 2003b, 111(6), pp. 1342–52.
- Niederle, Muriel and Roth, Alvin E.** "The Gastroenterology Fellowship Match: How It Failed, and Why It Could Succeed Once Again." *Gastroenterology*, 2004a, 127(2), pp. 658–66.
- Niederle, Muriel and Roth, Alvin E.** "Market Culture: How Norms Governing Exploding Offers Affect Market Performance." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 10256, 2004b.
- Roth, Alvin E.** "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory." *Journal of Political Economy*, 1984, 92(6), pp. 991–1016.
- Roth, Alvin E.** "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." *Econometrica*, 2002, 70(4), pp. 1341–78.
- Roth, Alvin E. and Peranson, Elliott.** "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 1999, 89(4), pp. 748–80.
- Roth, Alvin E. and Xing, Xiaolin.** "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review*, 1994, 84(4), pp. 992–1044.

A Kidney Exchange Clearinghouse in New England

By ALVIN E. ROTH, TAYFUN SÖNMEZ, AND M. UTKU ÜNVER*

In September, 2004, the Renal Transplant Oversight Committee of New England approved the establishment of a clearinghouse for kidney exchange, proposed by Francis Delmonico, Susan Saidman, and the three authors of this paper. We outline here the potential gains from kidney exchange and discuss practical constraints encountered as we begin designing and implementing a matching mechanism.

I. Background

In 2003 there were 8,665 transplants of deceased donor kidneys for the approximately 60,000 patients waiting for such transplants in the United States. While waiting, 3,436 patients died. There were also 6,464 kidney transplants from living donors (Scientific Registry of Transplant Recipients web site).¹ Live donation is an option for kidneys, since healthy people have two and can remain healthy with one. While it is illegal to buy or sell organs, there have started to be kidney *exchanges* involving two donor–patient pairs such that each (living) donor cannot give a kidney to the intended recipient because of blood type or immunological incompatibility, but each patient can receive a kidney from the other donor. So far these have been rare: as of December 2004, only five exchanges had been performed in the 14 transplant centers in New England. One reason there have been so few kidney exchanges is that there have not been databases of incompatible patient–donor pairs. Incompatible donors were simply sent home. (Databases are now being assembled not only in New England, but also in Ohio and Baltimore.)

* Roth: Harvard Business School and Department of Economics, Harvard University, Cambridge, MA 02138 (e-mail: aroth@hbs.edu); Sönmez: Koç University, Istanbul, Turkey, and Harvard University (e-mail: tsonmez@ku.edu.tr); Ünver: Koç University (e-mail: uunver@ku.edu.tr). We are grateful to Francis Delmonico and Susan Saidman for comments, and to NSF for support.

¹ <<http://www.ustransplant.org/srtr.php>>

Lainie Friedman Ross et al. (1997) discussed the possibility of exchange between incompatible patient–donor pairs. Not only have a few such two-way exchanges been performed, but two three-way exchanges (in which the donor kidney from one pair is transplanted into the patient in a second pair, whose donor kidney goes to a third pair, whose donor kidney goes to the first pair) have been performed at Johns Hopkins. There have also been a number of “list exchanges” in which an incompatible patient–donor pair makes a donation to someone on the waiting list for a cadaver kidney, in return for the patient in the pair receiving high priority for a cadaver kidney when one becomes available.

II. Scope and Design of a Kidney Clearinghouse:

In Roth et al. (2004a), we considered how to organize all these kinds of exchanges efficiently, in a way that would give patients and their doctors straightforward incentives. (Because medical information is decentralized, some of the procedures for allocating cadaver organs have experienced incentive problems.) We modeled patients as having strict preferences over compatible kidneys, and we allowed exchanges among any number of patient–donor pairs (including not only incompatible pairs, but also compatible pairs who might nevertheless be able, through exchange, to obtain a preferred kidney). We allowed list exchanges to be integrated with live exchanges, so a patient–donor pair who decided to exchange their kidney for priority on the deceased donor list would not necessarily donate their kidney to someone on the list, but might instead donate their kidney to another patient–donor pair who would in turn donate a kidney to the list (or to another pair who would in turn donate a kidney to the list, etc.).

In our model each agent is a patient and her donor(s). Agents have strict preferences over other agents (based on compatibility, closeness

of tissue match, and age of donor), and over priority on the cadaver wait-list.

If we exclude list exchange, this is the “housing market” of Lloyd Shapley and Herbert Scarf (1974), and David Gale’s method of top trading cycles (TTC) produces efficient, core allocations. There is a unique such allocation (Roth and Andrew Postlewaite, 1977), and the mechanism that selects it is dominant-strategy incentive-compatible (Roth, 1982).

TTC works as follows: Each agent points to her most preferred agent (the patient with the agent’s favorite donor). There is at least one *cycle* [an ordered list of agents (a_1, a_2, \dots, a_n) in which each agent points to the next, and agent a_n points to a_1], and no agent can be part of more than one cycle. The implied exchange in each cycle is carried out, and the procedure continues with each remaining agent pointing to her favorite among the remaining agents.

When list exchange is included the model is close to the “room assignment” of Atila Abdulkadiroğlu and Sönmez (1999). At some point of the TTC procedure there may be no cycles, but only “*w*-chains” in which a_n is pointing to the waiting list. An agent may be part of several *w*-chains and therefore the procedure needs a selection rule for *w*-chains. In Roth et al. (2004a) we called this class of procedures *top trading cycles and chains* (TTCC) and identified a version that is Pareto efficient and dominant-strategy incentive-compatible.

To solve one aspect of the incentive problem, all surgeries in a live-donor exchange are conducted simultaneously. Thus a two-way exchange (involving just two patient–donor pairs) involves four simultaneous surgeries, a three-way exchange involves six, and so on.

III. Logistical Constraints:

Our medical colleagues worried that, at least initially, they could not manage exchanges larger than two-way. They were also inclined to exclude list exchanges, and to allow only incompatible patient–donor pairs to participate. As a first approximation, their feeling was that a patient should be indifferent between any compatible exchanges.

Therefore, in Roth et al. (2004b), each agent is a patient with incompatible donors and is

indifferent between all donors compatible with her. No exchange larger than two-way is feasible. Building on well-known results in graph theory we showed that there are constrained-efficient dominant-strategy incentive-compatible mechanisms. These include deterministic “priority” mechanisms like those organ banks use to allocate cadaver organs, and stochastic mechanisms that address equity considerations.

The gains from kidney exchange will depend on several factors including:

- (i) the size of the patient–donor database;
- (ii) whether list exchanges are included (while list exchanges have distributional implications for the deceased donor wait-list, their inclusion increases the potential gains from exchange);
- (iii) the maximum number of transplants that can be simultaneously carried out (equivalently, the size of largest feasible cycle and/or *w*-chain); and
- (iv) whether compatible patient–donor pairs can participate in exchange.

Consider pairs A and B: donor A is compatible with both patients, but donor B is compatible only with patient A. While donor A can directly give her kidney to patient A, both patients receive a kidney if pairs A and B exchange. Such an exchange is called an *altruistically unbalanced exchange* (E. Steve Woodle and Ross, 1998) and is unlikely to be recommended to couple A as long as such exchanges are unusual. But if patients have strict preferences over donors, it could be that both pairs obtain a preferred kidney from such an exchange. (Consideration of compatible pairs, and altruistically unbalanced exchanges, will help us estimate an upper bound on the gains that can be achieved.)

We turn to simulations to estimate the impact of each of these factors on the number of patients who can benefit from exchange.

IV. Simulations

For simplicity we consider non-blood-related patient–donor pairs. Distributions of blood types (48 percent O, 34 percent A, 14 percent B,

4 percent AB), PRA levels (discussed below), and gender of the patients (41 percent female), and percentage of spouses among the unrelated donors (49 percent) are from the UNOS/OPTN data.²

Tissue-type incompatibility (a *positive cross-match*) arises when a patient has antibodies against a donor protein. (The positive cross-match probability between female patients and their husbands is approximately 33 percent, compared to approximately 11 percent between random pairs (Stefanos Zenios et al., 2000), because antibodies can develop during childbirth.) Patients in the UNOS database are divided into three groups based on the odds that they have a crossmatch with a random donor. For simplicity we simulate patients in discrete PRA (*percent reactive antibody*) levels:

- (i) 70 percent low-PRA patients, each of whom has a positive crossmatch probability of 5 percent with a random donor,
- (ii) 20 percent medium-PRA patients, each of whom has a positive crossmatch probability of 45 percent, and
- (iii) 10 percent high-PRA patients, each of whom has a positive crossmatch probability of 90 percent.

We randomly simulate patient–donor pairs using Monte-Carlo simulation size of 100 random population constructions for each of the 16 scenarios described below:

- (1) We consider two population sizes: 25 and 100.
- (2) We consider including compatible pairs in exchange as well as excluding them. (For example, in a population of 25 patient–donor pairs, if compatible pairs are excluded from exchange, only the smaller number of incompatible pairs will be available for exchange, and these will have a different distribution of characteristics than the general

population; O donors will be rare, and high PRA patients will be more common.)

- (3) Either
 - (a) list exchanges are unavailable; or
 - (b) list exchanges are available but only 40 percent of incompatible pairs consider a transplant from a deceased donor and only if a live donor is unavailable.
- (4) The largest feasible cycle/*w*-chain is either 2 or unbounded.

These possibilities yield $2 \times 2 \times 2 \times 2 = 16$ scenarios, and for each realization we search for a feasible exchange that includes the maximum number of patients.

For simplicity we assume that patients are indifferent between compatible live donors but prefer any such donor to priority on the deceased-donor wait-list. We use versions of Jack Edmond's (1965) algorithm to find a maximal exchange when the largest feasible cycle/*w*-chain is 2. We know of no efficient algorithm to determine a maximal exchange when cycle/*w*-chain size is unbounded. In these scenarios we search for a maximal exchange among efficient matchings via the TTCC algorithm.

Table 1 makes clear that the gains from all kinds of exchange increase as the population *n* of patient–donor pairs grows. The exchanges that are initially likely to be achievable are those involving no list exchange (0 percent wait-list), and only incompatible patient–donor pairs. When only two-way exchanges are feasible, exchange yields on average an additional 3.96 such transplants when $n = 25$ (16 percent of the patient population), but 23.04 when $n = 100$. Allowing list exchange, or allowing larger than two-way exchanges each gives a comparable increase in the number of transplants that can be achieved.

The largest gains in the table come from including compatible pairs in the population eligible for exchange. As the bottom of Table 1 indicates, it is at least conceivable that in a large population in which all patient–donor pairs could participate in exchange, virtually every patient (98.83 percent) with a willing donor would be able to receive a kidney. But we emphasize that this is an upper bound, since for many *compatible* pairs, exchange will not be desirable.

² UNOS/OPTN 2003 Annual Report, 1993–2002 (<http://www.optn.org>). Patient characteristics are from new waiting-list registrations, living donor relational type is from living-donor transplants data.

TABLE 1—SIMULATION RESULTS

Comp. pairs	n	WL	Transplants		
			Own	Ex.	w-List
A. Two-Way Exchange:					
Out	25	0	11.56	3.96	0
		40	11.56	5.76	3.71
		100	0	47.49	23.04
In	25	40	47.49	28.79	11.48
		0	1.33	19.00	0
		40	1.33	19.63	2.12
	100	0	1.01	90.14	0
		40	1.01	91.35	4.70
B. Unrestricted Exchange:					
Out	25	0	11.56	5.33	0
		40	11.56	6.32	3.82
		100	0	47.49	28.71
In	25	40	47.49	30.38	11.82
		0	1.44	20.30	0
		40	1.50	20.29	2.20
	100	0	1.67	92.68	0
		40	1.61	92.55	4.67

Notes: The table reports the average number of patients receiving a transplant in each scenario (through own-donor, through an exchange, and through being sent to the top of the waiting list). Key to column headings: Comp. pairs = compatible pairs (excluded ["out"] or included ["in"]); n = population size; WL = percentage of wait-list options; Own = number of patients receiving own donor kidneys; Ex. = number of patients participating in an exchange; w-List = number of patients who get priority in the waiting list through list exchange.

V. Conclusions

Kidney exchange is likely to proceed incrementally, starting with the simplest cases (two-way exchange) and the patients who can benefit most (incompatible pairs). Roth et al. (2005) show that most of the gain from larger than two-way exchange comes from three-way exchange, and so we are hopeful that it will be possible to achieve these gains in the near term. It may also be possible to include list exchanges and nondirected donors (altruistic living donors who do not specify a particular patient). Each of these increases in the scope of exchange will necessitate design changes in the clearinghouse, and there are open theoretical problems remaining for some of them (as is to be expected; cf. the examples in Roth [2002]).

It seems likely that, until exchange becomes well established, only incompatible patient-

donor pairs will be included, as surgeons will be reluctant to advise compatible pairs not to proceed with their own transplant. However, as exchange becomes more routine, there will be opportunities for mutually beneficial exchange between, for example, a 25-year-old patient with a compatible 50-year-old donor and a 50-year-old patient with an incompatible 25-year-old donor.

Fortunately, the gains from even the simplest exchanges are large, and achievable.

REFERENCES

- Abdulkadiroğlu, Atila and Sönmez, Tayfun.** "House Allocation with Existing Tenants." *Journal of Economic Theory*, 1999, 88(2), pp. 233–60.
- Edmonds, Jack.** "Paths, Trees, Flowers." *Canadian Journal of Mathematics*, 1965, 17, pp. 449–67.
- Ross, Lainie Friedman; Rubin, David T.; Siegler, Mark; Josephson, Michelle A.; Thistlethwaite, J. Richard, Jr. and Woodle, E. Steve.** "Ethics of a Paired-Kidney-Exchange Program." *New England Journal of Medicine*, 1997, 336(24), pp. 1752–55.
- Roth, Alvin E.** "Incentive Compatibility in a Market with Indivisibilities." *Economics Letters*, 1982, 9(2), pp. 127–32.
- Roth, Alvin E.** "The Economist as Engineer: Game Theory, Experimental Economics, and Computation as Tools of Design Economics." *Econometrica*, 2002, 70(4), pp. 1341–78.
- Roth, Alvin E. and Postlewaite, Andrew.** "Weak versus Strong Domination in a Market with Indivisible Goods." *Journal of Mathematical Economics*, 1977, 4(2), pp. 131–37.
- Roth, Alvin E.; Sönmez, Tayfun and Ünver, M. Utku.** "Kidney Exchange." *Quarterly Journal of Economics*, 2004a, 119(2), pp. 457–88.
- Roth, Alvin E.; Sönmez, Tayfun and Ünver, M. Utku.** "Pairwise Kidney Exchange." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 10698, 2004b; *Journal of Economic Theory* (forthcoming).
- Roth, Alvin E.; Sönmez, Tayfun and Ünver, M. Utku.** "Efficient Kidney Exchange: Coincidence of Wants in a Structured Market." National Bureau of Economic Research

(Cambridge, MA) Working Paper No. 11402, 2005.

Shapley, Lloyd and Scarf, Herbert. "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1974, 1(1), pp. 23–28.

Woodle, E. Steve and Ross, Lainie Friedman. "Paired Exchanges Should Be Part of the Solution to ABO Incompatibility in Living

Donor Kidney Transplantation." *Transplantation*, 1998, 66(3), pp. 406–7.

Zenios, Stefanos; Woodle, E. Steve and Ross, Lainie Friedman. "Primum Non Nocere: Avoiding Harm to Vulnerable Wait List Candidates in an Indirect Kidney Exchange." *Transplantation*, 2001, 72(4), pp. 648–54.