

The New York City High School Match

By Atila Abdulkadiroğlu, 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 2003-04, 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% 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 thirty community school districts. In the 1990s, the city began to take more centralized control, and in 2002, a newly reorganized NYCDOE began to reform many aspects of the school system.

In May 2003, Dr. 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 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) Dr. Lack, his colleagues (particularly Elizabeth Sciabarra and Neil Dorosin), and the DOE's software vendor, about the design of the match.

I. The prior (2002-03) NYC matching procedure

There are seven specialized high schools in NYC whose places are allocated by entrance exam (one by auditions). Rising high school students (mostly 8th, 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% 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% be allocated to students who were rated top performers in a standardized English Language Arts exam, 68% to middle performers, and 16% 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 a top 2% ELA score 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, waitlist, 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 waitlist. 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 waitlisted 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 5 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 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. And 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 2002-03 *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 and

Abdulkadiroğlu, Pathak, Roth, and Sönmez 2005.) Third, schools were also strategic: a substantial number 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 NYC 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, an efficient assignment such that no school and student not matched to one another would both prefer to be. So the question was, are the students the only real players in the system, with choices by schools merely a device for allocating scarce spaces? In the latter 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 NYC 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. And the fact that school administrators gamed the system indicated they were strategic players.

The medical match employs an applicant-proposing deferred-acceptance algorithm (Gale and Shapley, 1962, Roth and Peranson 1999.) Ignoring for the moment the details of NYC 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:

- 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.
- 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.
- 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 Sotomayor 1990, Sönmez 1997).

Of course, adapting the mechanism to the regulations and customs of NYC 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% 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 to also be given an offer from a non-specialized school. Therefore students who applied to specialized schools submit a preference list of non-specialized 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 non-specialized school. After they respond, capacities are adjusted and 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 non-specialized 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 NYC 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 5 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 this year than last year.

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

NYC 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/waitlist process, and provides more straightforward incentives to applicants.

References

- Abdulkadiroğlu, Atila and Sönmez, Tayfun. "School Choice: A Mechanism Design Approach." American Economic Review, June 2003, 93-3: 729-747.
- Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. "The Boston Public School Match," American Economic Review, Papers and Proceedings, 2005 (this issue).
- Gale, David and Shapley, Lloyd. "College Admissions and the Stability of Marriage," American Mathematical Monthly, January 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, vol. 92, 991-1016.
- Roth, Alvin E. "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics," Econometrica, July 2002, 70(4), 1341-1378.
- Roth, A.E. and E. Peranson. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," American Economic Review, September 1999, 89(4), 748-780.
- Roth, A.E. and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monograph Series, Cambridge University Press, 1990.
- Roth, Alvin E. and Xiaolin Xing. "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," Journal of Political Economy, April 1997, 105, 284-329.
- Sönmez, Tayfun. "Manipulation via Capacities in Two-Sided Matching Markets," *Journal of Economic Theory*, November 1997, 77(1), 197-204.

* Abdulkadiroğlu: Columbia University. Pathak and Roth: Harvard University.