

Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match

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

The design of the New York City (NYC) high school match involved trade-offs among efficiency, stability, and strategy-proofness that raise new theoretical questions. We analyze a model with indifferences—ties—in school preferences. Simulations with field data and the theory favor breaking indifferences the same way at every school—single tiebreaking—in a student-proposing deferred acceptance mechanism. Any inefficiency associated with a realized tiebreaking cannot be removed without harming student incentives. Finally, we empirically document the extent of potential efficiency loss associated with strategy-proofness and stability, and direct attention to some open questions. (JEL C78, D82, I21)

In 2003–2004, the authors of this paper assisted the New York City Department of Education (NYCDOE) in redesigning the student assignment mechanism used to match over 90,000 entering students to public high schools each year (Abdulkadiroğlu, Pathak, and Roth 2005). The old system—three rounds of a decentralized process of offers, acceptances, and wait lists with counterproductive incentives and insufficient offers to clear the market—was replaced with a more centralized single-offer system based on a student-proposing deferred acceptance algorithm adapted to satisfy various constraints of the NYCDOE.

The NYCDOE initially contacted us because they were familiar with the design of the medical match (see Roth 1984; Roth and Elliot Peranson 1999). And much of the theory on which the design of the new student assignment mechanism was based involves the same kind of matching theory (see Abdulkadiroğlu and Tayfun Sönmez 2003). But much of the prior theory of two-sided matching concentrates on the case that all parties have strict preferences, mainly because indifferences in preferences were viewed as a “knife-edge” phenomenon in applications like labor markets (cf. Roth and Marilda Sotomayor 1990). In contrast, a primary feature of school choice in NYC and in other cities including Boston (in which a new design was implemented in 2006, see Abdulkadiroğlu et al. 2005, 2006) is that there are indifferences—ties—in how students are ordered by at least some schools. How to break these ties raised some significant design decisions, which have potentially important strategic and welfare consequences (cf. Aytek Erdil and Haluk Ergin 2008). This paper describes those decisions, and evaluates them both empirically and with the aid of some new theory involving the trade-offs among efficiency, stability, and strategy-proofness.

* Abdulkadiroğlu: Department of Economics, Duke University, Durham NC, 27708 (e-mail: atila.abdulkadiroglu@duke.edu); Pathak: Society of Fellows and MIT, Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02142 (e-mail: ppathak@mit.edu); Roth: Department of Economics, Harvard University, Cambridge, MA 02138 and Harvard Business School, Boston, MA 02163 (e-mail: aroth@hbs.edu). We thank Susan Athey, Neil Dorosin, Aytek Erdil, Onur Kesten, Fuhito Kojima, and Tayfun Sönmez for helpful comments, and seminar participants at various universities and conferences. We are especially grateful to the three anonymous referees for helpful comments. For financial support, Abdulkadiroğlu gratefully acknowledges an Alfred P. Sloan Research Fellowship and an NSF-CAREER award, Pathak thanks the Spencer Foundation, and Roth thanks the National Science Foundation.

In two-sided matching models (David E. Gale and Lloyd S. Shapley 1962), there are two disjoint sets of agents, students, and schools. Every agent is an active strategic player with preferences over the agents in the other set.¹ A matching of students to schools (that respects the relevant capacity constraints) is *stable* if it is individually rational and there is no *blocking* pair of a student and a school who each prefer to be assigned to one another than to their mate in the matching.² Stable matchings are Pareto efficient with respect to the set of all agents, and in the core of the market whose rules are that any pair of agents on opposite sides of the market can be matched (subject to capacity constraints) if and only if they both agree. The current system in NYC involves some schools that are active strategic players that rank students in terms of preferences, while others use exogenous criteria—priorities—to rank students passively. When schools actively rank students, if there is a blocking pair, the school has an incentive to circumvent the match to enroll the students it would prefer. This was an important feature of the old system in NYC, when some schools concealed capacity in an effort to be matched later with preferable students.³ Empirical observations suggest that centralized matching mechanisms in labor markets are most often successful if they produce stable matchings (Roth 1991, 2002, 2008). Therefore, the fact that both schools and students are active participants in New York called for a stable matching algorithm for the NYC high school match.

A second desideratum guiding the design was that the school match process should promote student welfare to the extent possible; that is, it should be *efficient* for students. Indeed, the old system was not even capable of generating enough offers to clear the market and “[i]t [was] not unusual for up to 45 percent of students who apply to schools outside their neighborhood to be rejected by all their choices,” after which they would be assigned administratively, without regard for their stated preferences.⁴ When preferences are strict, a class of *deferred acceptance algorithms* identify a stable match that is optimal for one or the other side of the market, in the sense that every agent on one side of the market does at least as well at their optimal stable matching as at any other stable matching (Gale and Shapley 1962).

A third concern was the gaming aspect of the old system. Some schools gave higher priority to those students who ranked them as a first or second choice. This information was made public in the Directory of Public High Schools by NYCDOE. Consequently, students were forced to make a strategic decision while ranking schools. As a result, *strategy-proofness* for students, which requires that it should be a dominant strategy to state true preferences, became the third goal of the design. There do not exist any stable mechanisms that are strategy-proof for all agents. None is strategy-proof for schools matched to more than one student, but the *student-proposing deferred acceptance mechanism*, which produces the student-optimal stable matching, is strategy-proof for the students (Lester E. Dubins and David Freedman 1981; Roth 1982, 1985).

One way to think about these three design concerns is that Pareto efficiency for the students is the primary welfare goal, but that stability of the matching, and strategy-proofness in the elicitation of student preferences, are incentive constraints that likely have to be met for the system to

¹ Gale and Shapley (1962) phrased their discussion in a school choice context as a “college admissions” problem and studied the set of stable matchings. Michel Balinski and Sönmez (1999) considered the theory of two-sided matching in relation to college admissions where ordering of students at colleges is determined via students’ scores at an entrance exam. Abdulkadiroğlu and Sönmez (2003) introduced the problem of student assignment in school choice. Gale and Shapley already considered nonstrict preferences (albeit obliquely), but concentrated on the case of strict preferences.

² Stable matchings may be relevant even when all schools are passive. In this case, stable matchings eliminate “justified envy.” See Abdulkadiroğlu (2005); Abdulkadiroğlu and Sönmez (2003); and Abdulkadiroğlu et al. (2006) for discussion in the context of Boston’s new student assignment mechanism.

³ E.g., the deputy chancellor of schools, quoted in the *New York Times* (November 19, 2004): “Before you might have a situation where a school was going to take 100 new children for ninth grade, they might have declared only 40 seats, and then placed the other 60 outside the process.”

⁴ See “Many Are Shut Out in High School Choice,” *New York Times*, March 11, 2003.

produce substantial welfare gains over the old system.⁵ One goal of the present paper is to assess empirically how much student welfare is lost by meeting these constraints, by comparing the outcomes that result from Pareto improvements among students due to relaxing these constraints (if these improvements could actually be achieved) with the best student outcomes that can be achieved by a stable and strategy-proof mechanism. Another goal is to make clearer theoretically the trade-off between strategy-proofness and student welfare. From this point on, we will mostly be concerned with student welfare and incentives, and we will speak about Pareto efficiency and strategy-proofness with respect to students only.

Student assignment in school choice, introduced by Abdulkadiroğlu and Sönmez (2003) as an application of matching theory, requires the consideration of a model that allows for indifferences in school's preferences. When school preferences are not strict, there will not in general exist a unique stable match for each side of the market that is weakly Pareto optimal for that side; rather, there will be a nonempty set of stable matches that are weakly Pareto optimal for agents on that side. It will be useful in what follows to call any stable matching that is weakly Pareto optimal *a* student-optimal stable matching.⁶

With the student-proposing deferred acceptance algorithm as the basis of the new system, the mechanism must specify how to order equivalent students from the point of view of schools with limited space. For instance, one can assign each student a distinct number, break ties in school preferences according to those assigned numbers, then apply the deferred acceptance algorithm to the strict preferences that result from the tiebreaking. This preserves the stability and strategy-proofness of the student-proposing deferred acceptance algorithm. However, tiebreaking introduces artificial stability constraints (since, after tiebreaking, schools appear to have strict preferences between students for whom they are indifferent), and these constraints can harm student welfare. In other words, when the student-proposing deferred acceptance algorithm is applied to the strict preferences that result from tiebreaking, the outcome it produces may not in fact be a student-optimal stable matching in terms of the original preferences.

Nevertheless, some forms of tiebreaking may be preferable to others. One of the first design decisions we confronted was whether to assign numbers to each student at each school (multiple tiebreaking), or to give each student a single number to be used for tiebreaking at every school (single tiebreaking). Computations with simulated and then actual submitted preferences indicated that single tiebreaking had superior welfare properties.⁷ The computations based on the revealed preferences will be discussed in Section II. Proposition 2 provides some theoretical insight into the difference. We observe that if there is a matching produced by student-proposing deferred acceptance with multiple tiebreaking that cannot be produced by deferred acceptance with single tiebreaking, then the matching is not a student-optimal stable matching.⁸

However, single tiebreaking can also lead to a matching that is not a student-optimal stable matching. Thus there will sometimes be a potential opportunity to improve on the outcome of deferred acceptance with single tiebreaking. Theorem 1, our main theoretical result, implies that there exists no strategy-proof mechanism (stable or not) that Pareto improves on the deferred acceptance algorithm with single tiebreaking. That is, the potential inefficiency of student-proposing deferred acceptance with single tiebreaking is the cost of strategy-proofness.

⁵ Strategy-proofness can also be a goal in itself, as was the case in the design of the Boston school choice system, where it came to be seen as a criterion of equal access for families with different degrees of sophistication about the system.

⁶ See Robert Day and Paul Milgrom (2008) for a similar treatment of the core for auctions with nontransferable utility, and its relation to matching.

⁷ Tayfun Sönmez also played an important role in some of the early discussions of this.

⁸ That the manner of tiebreaking has important consequences for the outcome of deferred acceptance algorithms (when stability is an issue) contrasts in a surprising way with the case of one-sided matching. Pathak (2006) shows

This theorem does not speak to the magnitude of the trade-off to expect in real school choice plans. We take a step toward investigating this question using student preference data from New York City. Our simulations reveal significant costs of imposing strategy-proofness: about 1,500 students, or 1.9 percent of eighth graders in our sample, could be matched to schools they prefer over their assignment from deferred acceptance with single tiebreaking without harming any others, if the same preference information could be elicited by a nonstrategy-proof mechanism. A further (approximately) 4,300 students, or 5.5 percent of eighth graders in our sample, could improve their assignment if we relaxed the stability constraint, if the same rates of participation could be elicited by an unstable mechanism. In contrast, in Boston, the welfare cost of strategy-proofness and stability are negligible. These results raise new questions concerning when potential efficiency gains exist, and whether they can be realized.

The rest of the paper is organized as follows. Section I describes the model. Section II presents our main theoretical and empirical results. The last section concludes.

I. Model

A matching problem consists of a finite set of students I (for “individuals”) and a finite set of schools S . Each school $s \in S$ has q_s available seats. A **matching** is a correspondence $\mu : I \cup S \rightarrow S \cup I$ such that each student is assigned either to only one school or to herself, and the number of students matched at each school is less than or equal to its capacity. If $\mu(i) = i$, i remains unmatched. If $|\mu(s)| < q_s$, some seats at school s remain unfilled.

Every student i has a strict preference relation P_i over schools and being unmatched, $S \cup \{i\}$. A school is acceptable to a student if she strictly prefers to be matched to that school over being unmatched. Let $s R_i s'$ denote that $s P_i s'$ or $s = s'$. Every school s has a weak preference relation R_s over students and keeping unfilled seats, $I \cup \{s\}$. A student i is acceptable to a school s if s prefers to be matched to i , $i R_s s$. Let \succ_s and \sim_s represent the asymmetric and symmetric parts of R_s , respectively. To simplify, we assume that either $i \succ_s s$ or $s \succ_s i$, but not $i \sim_s s$. Throughout the paper, we will assume that the preference relation of a school s over subsets of students is **responsive** to R_s . This means that a school's preferences over groups of students is such that, for any group of students I' with $|I'| < q_s$, the school prefers $I' \cup \{i\}$ to $I' \cup \{j\}$ if and only if $i \succ_s j$, and prefers $I' \cup \{i\}$ to I' if and only if $i \succ_s s$ (Roth 1985). For $I' \subset I$, let $P_{I'} = (P_i)_{i \in I'}$ and let $P_{-i} = (P_j)_{j \neq i}$. We define $R_{S'}$ and R_{-s} similarly. We fix I, S, R_S , and q throughout the paper.

When schools actively rank all students, this model reduces to the college admissions model. When every school is indifferent between all students, it reduces to the house allocation model where there are potentially multiple places in each house. More importantly, the model allows both types of schools, as well as schools with several indifference classes simultaneously.

A matching μ is **individually rational** if it matches every $x \in I \cup S$ with agent(s) that is(are) acceptable for x . A matching μ is **blocked by** (i, s) if $s P_i \mu(i)$, and either $[|\mu(s)| < q_s \text{ and } i \succ_s s]$ or $[i \succ_s i' \text{ for some } i' \in \mu(s)]$. μ is **stable** if μ is individually rational and not blocked by any student-school pair (i, s) . A matching μ **dominates** matching ν if $\mu(i) R_i \nu(i)$ for all $i \in I$, and $\mu(i) P_i \nu(i)$ for some $i \in I$. A stable matching μ is a **student-optimal stable matching** if it is not dominated by any other stable matching. A matching μ is **efficient** if there is no other matching (stable or not) that dominates it.

A direct **mechanism** φ is a function that maps every (P_i, R_S) to a matching. For $x \in I \cup S$, let $\varphi_x(P_i, R_S)$ denote the set of agents that are matched to x by φ . A mechanism φ is **dominant strategy incentive compatible (DSIC)** for $i \in I$ if for every (P_i, R_S) and every P_i' ,

the strategy-proof top trading cycles mechanism that produces efficient (but not stable) matchings is unaffected by the choice of multiple versus single tiebreaking, when all schools are indifferent between students.

$$\varphi_i(P_i; R_S) R_i \varphi_i(P_i', P_{-i}; R_S).$$

DSIC for schools is defined similarly. A mechanism is **strategy-proof** if it is DSIC for all students.⁹

We say that a mechanism φ **dominates** ψ if

for all $P_i : \varphi_i(P_i; R_S) R_i \psi_i(P_i; R_S)$ for all $i \in I$, and

for some $P_i : \varphi_i(P_i; R_S) P_i \psi_i(P_i; R_S)$ for some $i \in I$.

The mechanisms we will consider are based on the student-proposing deferred acceptance algorithm which, for a given profile of *strict* preferences, produces the unique student-optimal stable matching. The algorithm is defined as follows:

Step 1: Each student proposes to her most preferred school. Each school tentatively assigns its seats to its proposers one at a time in the order of its preference. When all of its seats are tentatively assigned, it rejects all the proposers who remain unassigned.

In general, at

Step k : Each student who was rejected in the previous step proposes to her next preferred school. Each school considers the set of students it has been holding and its new proposers. It tentatively assigns its seats to these students one at a time in the order of its preference. When all of its seats are tentatively assigned, it rejects all the proposers who remain unassigned.

The algorithm terminates when no student proposal is rejected. Each student who was tentatively assigned to some school when the algorithm ended is now assigned to that school.

Note that the student-proposing deferred acceptance algorithm can be adapted to our economy with indifferences, by using any tiebreaking procedure to convert the school preferences into strict preferences, and then applying the algorithm above. A tiebreaker, which is a bijection $r : I \rightarrow \mathbb{N}$, breaks ties at school s by associating R_s with a strict preference relation P_s as follows: $i P_s j \Leftrightarrow [(i \succ_s j) \text{ or } (i \sim_s j \text{ and } r(i) < r(j))]$.

We focus on single tiebreaking and multiple tiebreaking rules. A single tiebreaking rule uses the same tiebreaker at each school, while a multiple tiebreaking rule may use a different tiebreaker at each school. For a particular set of tiebreakers $\tau = (r_s)_{s \in S}$, let the mechanism DA^τ be the student-proposing deferred acceptance algorithm acting on the preferences (P_i, P_s) , where P_s is obtained from R_s by breaking ties using r_s , for all s . Define DA-STB to be the mechanism DA^τ , where τ is a single tiebreaking rule, and DA-MTB where τ is a multiple tiebreaking rule. The dominant strategy incentive compatibility of the student-proposing deferred acceptance mechanism for every student implies that DA^τ is strategy-proof for any τ .

When preferences are strict, there is no individually rational matching (stable or not) that is preferred to the unique student-optimal stable matching by every student, but even when preferences are strict, there may be unstable matchings at which some students do better than the

⁹ When preferences are strict, there is no mechanism that is DSIC for students and schools (Roth 1982), and there is no mechanism that is DSIC for schools (Roth 1985). These negative results generalize directly to our model. Fuhito Kojima and Pathak (2009) show that, as markets get large with bounded preference lists, the ability of schools to manipulate through misstating either their preferences or their capacities gets small.

student-optimal stable match, and no student does worse (Roth 1982). In our model, there may be multiple student-optimal stable matchings. This weak Pareto efficiency result generalizes for every student-optimal stable matching.¹⁰

PROPOSITION 1: *If μ is a student-optimal stable matching, there is no individually rational matching ν (stable or not) such that $\nu(i) P_i \mu(i)$ for all $i \in I$.*

A. The Model and the New York City High School Match

The model defined above captures the main features of the New York City high school match. However, there are some differences between the model and the actual New York City system. We briefly discuss three differences here (and in more detail in the Appendix). First, the New York City high school match consists of multiple rounds. Multiple rounds are used to accommodate the placement of students into specialized high schools, which account for about 4,000–5,000 seats a year, as well as to place unassigned and appealing students at the conclusion of the main process. Second, in the main round of the mechanism, students can rank only up to 12 schools of their choice. Across years of the new system, between 72 and 80 percent of applicants rank fewer than 12 schools. Finally, students in the top 2 percent of the grade 7 English language arts test score distribution are guaranteed their assignment if they rank a particular type of program, known as "educational option," as their first choice.

Each of these differences affects the dominant strategy properties of the mechanism.¹¹ However, as detailed in the Appendix, the consequences of these constraints are relatively small. Moreover, the issues we consider in the rest of the paper are present even in the absence of these constraints.¹² For our empirical analysis, we will focus on the initial rank order lists of nonspecialized schools. These rank order lists are used to assign the vast majority of applicants. We will use all rank order lists, including those that rank 12 schools, treating those stated preferences as true preferences. The Appendix shows that the students who rank 12 schools do not have an incentive to misreport their relative ordering of schools, so we can in fact treat a student who receives a higher ranked school as obtaining a welfare improvement. We do not include students who are in the top 2 percent and who rank an educational option program as their first choice in our calculations, because these students receive their stated top choice. The Appendix contains a thorough discussion of these constraints and their magnitudes.¹³

B. Motivating Examples

In this section, we develop the main ideas via two examples.

¹⁰ Proofs not in the text are in the Appendix.

¹¹ Guillaume Haeringer and Flip Klijn (2009) analyze the school choice problem when there is a constraint on the number of schools a student may rank. They identify necessary and sufficient conditions on the priorities needed to guarantee stability. See also Caterina Calsamiglia, Haeringer, and Klijn (forthcoming) for a related experiment. Pathak and Sönmez (2008) show that a student-proposing deferred acceptance mechanism where participants can rank at most k schools is manipulable whenever the mechanism where participants can rank at most l schools is manipulable, when $k < l$.

¹² For readers concerned with these idiosyncratic constraints, we report additional empirical analysis from Boston's school choice system, which does not have these idiosyncratic features.

¹³ Specifically, we consider that preference lists are capped at 12, that some schools give priority to students with top 2 percent reading scores if those students rank them first, and that the matching algorithm is conducted in multiple rounds. Because market design involves the application of theory to complex environments (which are known to the designers in great detail), it is useful to make clear how closely a simple model corresponds to the actual environment, and what consequences follow from the differences.

Example 1. (Tiebreaking does not always yield student-optimal stable matchings.) Tiebreaking has important welfare consequences. Suppose that school s_1 is indifferent among students, students i_1, i_2, i_3 , and schools s_2 and s_3 have the following strict preferences:

Student preferences	School preferences
$s_2 P_{i_1} s_1 P_{i_1} s_3$	$i_1 \sim_{s_1} i_2 \sim_{s_1} i_3$
$s_1 P_{i_2} s_2 P_{i_2} s_3$	$i_2 \succ_{s_2} i_1 \succ_{s_2} i_3$
$s_1 P_{i_3} s_2 P_{i_3} s_3$	$i_3 \succ_{s_3} i_1 \succ_{s_3} i_2$

The stable matchings are the following:

$$\mu_1 = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}, \quad \mu_2 = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}, \quad \mu_3 = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}.$$

Note that μ_1, μ_2 , and μ_3 are produced by the student proposing deferred acceptance algorithm (DA) when the indifference in s_1 's preferences is broken as $i_1 \succ_{s_1} i_3 \succ_{s_1} i_2, i_2 \succ_{s_1} i_x \succ_{s_1} i_y$ and $i_3 \succ_{s_1} i_x \succ_{s_1} i_y$, respectively. However, μ_2 dominates μ_1 despite μ_1 being stable. That is, DA need not produce a student-optimal stable matching even if ties at schools are broken the same way.

An interesting comparison of single versus multiple tiebreaking can be obtained via the following variation of Example 1: if none of the schools strictly ranks students, then the student-proposing deferred acceptance algorithm with single tiebreaking always yields a student-optimal stable matching. On the contrary, the student-proposing deferred acceptance algorithm with multiple tiebreaking yields the suboptimal matching μ_1 when ties at schools are broken as $i_1 \succ_{s_1} i_3 \succ_{s_1} i_2, i_2 \succ_{s_2} i_1 \succ_{s_2} i_3$, and $i_3 \succ_{s_3} i_1 \succ_{s_3} i_2$. In fact, Proposition 2 below states that given any problem, any matching that can be produced by deferred acceptance with some multiple tiebreaking, but that cannot be produced by deferred acceptance with any single tiebreaking, is not a student-optimal stable matching.

In Example 1, what happens if one attempts to improve student welfare when DA yields the suboptimal matching μ_1 ? We demonstrate in the following example that such welfare improvement yields loss of strategy-proofness.

Example 2. (Welfare enhancement involves losing strategy-proofness.) Consider Example 1 and the student proposing deferred acceptance mechanism when the ties at s_1 are broken as $i_1 \succ_{s_1} i_3 \succ_{s_1} i_2$. We will refer to that mechanism as DA^τ . When the students' preference profile is given as in Example 1, which we refer as P , DA^τ produces μ_1 . Suppose that there is a strategy-proof mechanism φ that dominates DA^τ , and in particular produces μ_2 under P .

Consider the following preference profile P' , which we obtain by changing i_1 's preferences in P :

$s_2 P'_{i_1} s_3 P'_{i_1} s_1$
$s_1 P_{i_2} s_2 P_{i_2} s_3$
$s_1 P_{i_3} s_2 P_{i_3} s_3$

Under P', DA^τ produces

$$\mu' = \begin{pmatrix} i_1 & i_2 & i_3 \\ s_3 & s_2 & s_1 \end{pmatrix}.$$

TABLE 1—TIEBREAKING FOR GRADE 8 APPLICANTS IN NYC IN 2006–2007

Choice	Deferred acceptance single tiebreaking DA-STB (1)	Deferred acceptance multiple tiebreaking DA-MTB (2)	Student-optimal stable matching (3)	Improvement from DA-STB to student-optimal	Number of students (4)
1	32,105.3 (62.2)	29,849.9 (67.7)	32,701.5 (58.4)	+1	633.2 (32.1)
2	14,296.0 (53.2)	14,562.3 (59.0)	14,382.6 (50.9)	+2	338.6 (22.0)
3	9,279.4 (47.4)	9,859.7 (52.5)	9,208.6 (46.0)	+3	198.3 (15.5)
4	6,112.8 (43.5)	6,653.3 (47.5)	5,999.8 (41.4)	+4	125.6 (11.0)
5	3,988.2 (34.4)	4,386.8 (39.4)	3,883.4 (33.8)	+5	79.4 (8.9)
6	2,628.8 (29.6)	2,910.1 (33.5)	2,519.5 (28.4)	+6	51.7 (6.9)
7	1,732.7 (26.0)	1,919.1 (28.0)	1,654.6 (24.1)	+7	26.9 (5.1)
8	1,099.1 (23.3)	1,212.2 (26.8)	1,034.8 (22.1)	+8	17.0 (4.1)
9	761.9 (17.8)	817.1 (21.7)	716.7 (17.4)	+9	10.2 (3.1)
10	526.4 (15.4)	548.4 (19.4)	485.6 (15.1)	+10	4.7 (2.0)
11	348.0 (13.2)	353.2 (12.8)	316.3 (12.3)	+11	2.0 (1.1)
12	236.0 (10.9)	229.3 (10.5)	211.2 (10.4)		
Unassigned	5,613.4 (26.5)	5,426.7 (21.4)	5,613.4 (26.5)	Total:	1,487.5

Notes: Data from the main round of the New York City high school admissions process in 2006–2007 for students requesting an assignment for grade 9 (high school). Column 1 reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with single tiebreaking (DA-STB). Column 2 reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with school-specific tiebreaking. Column 3 reports the average choice received distribution of applicants in a student-optimal stable matching, which is computed from DA-STB followed by stable improvement cycles. Column 4 reports the average number of students and how many places on their rank order list students improve in the student-optimal stable matching relative to the matching produced by DA-STB. Columns 1, 2, 3, and 4 are based on 250 random draws. Simulation standard errors are reported in parentheses.

Since φ dominates DA^τ , φ must also produce μ' under $(P'; R_s)$. Note that i_1 is matched with s_3 , her second choice, at μ' . But then i_1 can manipulate φ under P' by misrepresenting her preferences as $s_2 P_{i_1} s_1 P_{i_1} s_3$, because then φ matches her with s_2 , which she prefers to s_3 , the school she is matched with by reporting truthfully. Therefore no such strategy-proof mechanism exists for this problem.

Theorem 1, our main result, generalizes this insight.

II. Welfare Consequences of Tiebreaking

In the course of designing the New York City high school match, policymakers from the Department of Education were concerned with the fairness of tiebreaking. Even after the decision to use a student-proposing deferred acceptance algorithm, they believed that each student should receive a different random number at each program they applied to and this number should be used to construct strict preferences of schools for students. Through simple examples and simulations, we suggested that single tiebreaking might have superior welfare properties to multiple tiebreaking. The DOE remained unconvinced until student preferences had already been submitted, and computational experiments could be conducted comparing single and multiple tiebreaking using actual data from the first round in 2003–2004.

Table 1 presents a comparison of the distribution of school assignments based on the preferences submitted by grade 8 applicants in New York City 2006–2007. The column labeled Deferred Acceptance Single Tiebreaking (DA-STB) reports the average number of students receiving their top choice, second choice, and so on under the student-proposing deferred acceptance algorithm

with a single-tiebreaking rule, where tiebreakers are drawn independently from a uniform distribution a total of 250 times. For instance, on average 32,105.3 students are matched to their first choice, while 5,613.4 are unassigned. Simulation standard deviations are reported in parentheses.

The second column reports the average number of students receiving their top choice, second choice, and so on under the student-proposing deferred acceptance algorithm with a multiple-tiebreaking rule (DA-MTB), where tiebreakers are drawn independently from a uniform distribution across each school. The column reports the average from 250 independent draws from a uniform distribution. In contrast to DA-STB, on average 29,849.9 students, or 2,255.4 fewer students, receive their top choice under DA-MTB. The distributions produced by the two mechanisms are statistically different (two-sided, Wilcoxon $p < 0.001$), with the mean rank received under DA-STB being less than DA-MTB.¹⁴ While the mean and the number of students receiving their top choices are higher under DA-STB than DA-MTB, the simulations reveal that there is no stochastic dominance ordering (weaker than Pareto dominance) between the two mechanisms. For example, the number of students receiving their seventh choice or better is greater under DA-STB, while the number receiving their eighth choice or better is greater under DA-MTB. Note, however, that more than 89 percent of applicants receive their seventh choice or better.

The greater number of students obtaining one of their top choices in a similar simulation and in the first year of submitted preference data convinced New York City to employ a single tiebreaker in their assignment system. The following result provides some insight into the difference in the distribution of matchings produced by DA-STB and DA-MTB.

PROPOSITION 2: *For any (P, R_S) , any matching that can be produced by deferred acceptance with some multiple tiebreaking, but that cannot be produced by deferred acceptance with any single tiebreaking, is not a student-optimal stable matching.*

In other words, the set of student-optimal stable matchings is included in the set of stable matchings generated by a DA-STB, which is a subset of the set of stable matchings generated by a DA-MTB.¹⁵ This result has no direct implication for the distribution of matchings produced by both mechanisms ex ante. As we observed in the simulations, DA-STB and DA-MTB cannot be compared by stochastic dominance, a welfare criterion that relies only on ordinal information about preferences. In New York City, 2006–2007, out of the 250 draws of DA-STB, none is student-optimal. This is also the case for DA-MTB.

In case of a matching that is not a student-optimal stable matching, one can easily imagine Pareto improving the matching. Our next result states that such improvements harm incentives: any mechanism that can Pareto improve DA with some tiebreaking rule is not strategy-proof.

THEOREM 1: *For any tiebreaking rule τ , there is no mechanism that is strategy-proof and dominates DA^τ .*

PROOF:

We begin by establishing the following property of a matching that dominates a stable matching.

¹⁴ All statistical tests in this paper are based on the Mann-Whitney-Wilcoxon rank-sum test.

¹⁵ A similar version of this result, which states that every student-optimal stable matching can be produced by some DA-STB, was discovered independently by both Erdil (2006) and Lars Ehlers (2006). A natural question is if there is a smaller set of tiebreaking rules that will yield any student-optimal stable matching. To see that there is not, consider a problem in which all n students have the same preferences over all schools and all schools are indifferent between students. There are $n!$ student-optimal stable matchings, which correspond exactly to the $n!$ single tiebreaking rules. This shows that the set of single tiebreaking rules is the smallest such set.

CLAIM: Suppose that ν dominates $\mu = DA^\tau(P_I; R_S)$ for a given tiebreaking rule τ . Then the same set of students are matched in both ν and μ .

If there exists a student who is assigned under μ and unassigned under ν , then $\nu(i) = iP_i\mu(i)$, which implies that μ is not individually rational, a contradiction. So every student assigned under μ is also assigned under ν . Therefore $|\nu(S)| \geq |\mu(S)|$. If $|\nu(S)| > |\mu(S)|$, then there exists some $s \in S$ and $i \in I$ such that $|\nu(s)| > |\mu(s)|$ and $\nu(i) = s \neq \mu(i)$. This implies that there is a vacancy at s under μ and i is acceptable for s . Furthermore, $sP_i\mu(i)$ since ν dominates μ . These together imply that μ is not stable, a contradiction. So $|\nu(S)| = |\mu(S)|$. Then the same set of students are matched in both ν and μ since $|\nu(S)| = |\mu(S)|$ and every student assigned under μ is also assigned under ν . This completes the proof of the claim.

Fix R_S . Suppose that there exists a strategy-proof mechanism φ and tiebreaking rule r such that φ dominates DA^τ . There exists a profile P_I such that

$$\varphi_i(P_I; R_S) R_i DA_i^\tau(P_I; R_S) \text{ for all } i \in I, \text{ and}$$

$$\varphi_i(P_I; R_S) P_i DA_i^\tau(P_I; R_S) \text{ for some } i \in I.$$

We will say that the matching $\varphi(P_I; R_S)$ dominates the matching $DA^\tau(P_I; R_S)$, where $DA^\tau(P_I; R_S)$ denotes the student-optimal stable matching for $(P_I; P_S^\tau)$.

Let $s_i = DA_i^\tau(P_I; R_S)$ and $\hat{s}_i = \varphi_i(P_I; R_S)$ denote i 's assignment under $DA^\tau(P_I; R_S)$ and $\varphi(P_I; R_S)$, respectively, where $\hat{s}_i P_i s_i$.

Consider profile $P_I' = (P_i', P_{-i})$, where P_i' ranks \hat{s}_i as the only acceptable school. Since DA^τ is strategy-proof, $s_i = DA_i^\tau(P_I; R_S) R_i DA_i^\tau(P_I'; R_S)$, and since $DA_i^\tau(P_I'; R_S)$ is either \hat{s}_i or i , we conclude that $DA_i^\tau(P_I'; R_S) = i$. Then the claim proved above implies that $\varphi_i(P_I'; R_S) = i$.

Now consider the economy where $(P_I'; R_S)$ are the actual preferences. In this case, i could state P_i' and be matched to $\varphi_i(P_I'; R_S) = \hat{s}_i$, which under P_i' she prefers to $\varphi(P_I'; R_S) = i$. This shows that φ is not strategy-proof.

In other words, any inefficiency associated with a realized tiebreaking cannot be removed without harming student incentives. This result has close connections to three other results. First, Erdil and Ergin's (2008) new algorithm, *stable improvement cycles*, can be used to find a student-optimal stable matching that Pareto dominates the outcome of student-proposing deferred acceptance with any tiebreaking rule. They show that no student-optimal stable mechanism, such as their stable improvement cycles mechanism, is strategy-proof. Theorem 1 implies a stronger result, namely that it not possible for any strategy-proof mechanism to Pareto improve upon deferred acceptance with tiebreaking. Onur Kesten (2004) shows that when preferences are strict, there is no strategy-proof and Pareto efficient mechanism that Pareto dominates the deferred acceptance mechanism. Again, Theorem 1 gives a tighter bound, i.e., even when the efficiency requirement is relaxed, no strategy-proof mechanism can Pareto improve upon deferred acceptance with some tiebreaking, with or without strict preferences. Finally, neither serial dictatorship nor top trading cycles dominates deferred acceptance with single tiebreaking (since they are both strategy-proof), although neither one is dominated by deferred acceptance, since both are Pareto efficient for students. Thus, there is a trade-off between strategy-proofness and efficiency, a trade-off that hinges on stability. Deferred acceptance with tiebreaking, a stable mechanism, is on the efficient frontier of all strategy-proof mechanisms. Given the importance of stability and strategy-proofness, this theorem, combined with the evidence from the simulations with field data and Proposition 2, supports NYC's selection of DA-STB, even though it does not always yield a student-optimal stable matching.

This result also allows us to interpret the lack of student-optimality associated with a tie-breaking rule as the cost of providing straightforward incentives to students. In Table 1, the third column labeled student-optimal stable matching reports the average number of students receiving their first choice, second choice, and so on in a student-optimal stable matching. This matching was computed by starting from a particular DA-STB in column 1, and employing the stable improvement cycles procedure of Erdil and Ergin (2008).¹⁶ The third column shows that 32,701.5 students receive their top choice in a student-optimal stable matching, which is an improvement of 596.2 students on average. In column 4, we report the average improvement in the ranking of the match received in the student-optimal stable matching from DA-STB. For instance, 633.2 students on average receive a school that is ranked one place higher on their rank order list under the student-optimal stable matching i.e., a student who received her n th choice under DA-STB, receives her $(n - 1)$ th choice in a student-optimal stable matching. Similarly, 338.6 students on average improve on their rank order list by two places, i.e., go from their n th choice under DA-STB, to their $(n - 2)$ th choice in a student-optimal stable matching. In total, 1,487.5 students on average receive a more preferred matching in the student-optimal stable matching. This represents the cost of strategy-proofness associated with DA-STB. The distribution of the student-optimal stable matching is statistically different from DA-STB (one-sided, $p < 0.001$).¹⁷

A. Welfare Cost of Stability

The student-optimal stable matching is not dominated by any other stable matching. However, it is not necessarily an efficient matching and thus can be dominated by an efficient matching. In this section, we compute the cost of stability by finding an efficient matching that dominates the student-optimal stable matching.

The empirical exercise starts with a student-optimal stable matching. We next find all welfare-enhancing swaps among assigned students across schools. These swaps are computed by Gale's top trading cycles algorithm, which is explained in detail in the Web Appendix (available at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.5.1954>). After students trade among themselves, the resulting efficient matching will dominate the student-optimal stable matching. Table 2 reports the results of this exercise. Column 1 is the same column as Table 1's column 3, and reports the distribution of assignments for a student-optimal matching. The next column, labeled efficient matching, reports the average number of students across the 250 independent draws of the student-optimal stable matchings from column 1 receiving their top choice, second choice, and so on in the efficient matching. In 2006–2007, in grade 8, 34,707.8 students on average receive their first choice, which is an improvement of 2,602.6 students. The next columns report the magnitude of the improvement students obtain in the efficient matching. Column 4 shows that a total of 4,296.6 students obtain a school that is better than the school that they obtained in a student-optimal stable matching. Of these students, 1,819.7, or 42.4 percent, go from their n th choice to their $(n - 1)$ th choice. The distribution of efficient matching is statistically different from the student-optimal matching (one-sided, $p < 0.001$).

Each of the welfare improving swaps among the students will lead to the formation of a blocking pair of a student and a school. If the blocking pair involves a student and a school that actively ranks applicants, then both the student and school will prefer to be matched with each other, and

¹⁶ The exact details on the implementation of stable improvement cycles are contained in the Web Appendix (available at <http://www.aeaweb.org/articles.php?doi=10.1257/AER.99.5.1954>).

¹⁷ In an earlier version of this paper, we computed the same empirical exercise for 2003–2004, 2004–2005, and 2005–2006 in New York. The quantitative results are of similar magnitudes.

TABLE 2—WELFARE CONSEQUENCES OF STABILITY FOR GRADE 8 APPLICANTS IN 2006–2007

Choice	Student-optimal stable matching (1)	Efficient matching (2)	Improvement from student-optimal stable matching	Number (3)	<i>k</i>	Count of students with <i>k</i> blocking pairs (4)
1	32,701.5 (58.4)	34,707.8 (50.5)	+1	1,819.7 (41.3)	1	22,287.5 (205.1)
2	14,382.6 (50.9)	14,511.4 (51.1)	+2	1,012.8 (26.4)	2	6,707.8 (117.9)
3	9,208.6 (46.0)	8,894.4 (41.2)	+3	592.0 (19.5)	3	2,991.0 (79.6)
4	5,999.8 (41.4)	5,582.1 (40.3)	+4	369.6 (16.0)	4	1,485.4 (56.5)
5	3,883.4 (33.8)	3,492.7 (31.4)	+5	212.5 (12.0)	5	716.6 (32.5)
6	2,519.5 (28.4)	2,222.9 (24.3)	+6	132.1 (9.1)	6	364.6 (22.9)
7	1,654.6 (24.1)	1,430.3 (22.4)	+7	77.0 (7.1)	7	183.1 (13.6)
8	1,034.8 (22.1)	860.5 (20.0)	+8	43.0 (5.6)	8	85.6 (10.9)
9	716.7 (17.4)	592.6 (16.0)	+9	26.3 (4.5)	9	44.7 (6.4)
10	485.6 (15.1)	395.6 (13.7)	+10	11.6 (2.8)	10	22.6 (4.9)
11	316.3 (12.3)	255.0 (10.8)	+11	4.8 (2.0)	11	9.9 (3.0)
12	211.2 (10.4)	169.2 (9.3)			12	3.2 (1.6)
Unassigned	5,613.4 (26.5)	5,613.4 (26.5)	Total:	4,296.6		34,898.8

Notes: Data from the main round of the New York City high school admissions process in in 2006–2007 for students requesting an assignment for grade 9 (high school). Column 1 reports the average choice received distribution of applicants in a student-optimal stable matching, which is computed from DA-STB followed by stable improvement cycles. Column 2 reports the average choice received distribution of applicants in a Pareto efficient matching, computed with Gale's top trading cycles, which dominates the matching in column 1. Column 3 reports the average number of students and how many places on their rank order list students improve in the efficient matching relative to a student-optimal stable matching. Column 4 reports the average number of students who form *k* different blocking pairs in the efficient matching. Columns 1, 2, 3, and 4 are based on 250 random draws. Simulation standard errors are reported in parentheses.

have incentives to circumvent the match. If the blocking pair involves a student and a school that uses only various fixed priorities (i.e., passively ranks applicants), then the student involved in the blocking pair will have higher priority than the student who receives the school in the efficient matching.

The last two columns of Table 2 report both types of blocking pairs that result from the efficient matching. The total number of students involved in at least one blocking pair is 34,898.8 on average, which corresponds to a total of 57,046 blocking pairs involving a student and school on average across the simulations. More than 90 percent of the blocking pairs involve a student and an educational option program, while about 7 percent of the remaining blocking pairs involve a “screened” program.¹⁸ Since both of these program types submit a rank order list on students, this type of violation of stability may lead programs to circumvent their assignment by matching with different students after the placements have been announced.

Of the students involved in a blocking pair, 22,287.5, or 63.9 percent, are involved in a blocking pair at one school; 6,707.8, or 19.2 percent, of students are involved in a blocking pair at two schools. There are even 3.2 students on average who are involved in a blocking pair at 12 schools. These students are unassigned, but in the efficient matching, there is a student who is assigned with either lower priority or lower ranking at each of their 12 choices. Taken together, this table shows that, while 4,296.6 students on average could potentially benefit from moving to

¹⁸ Appendix AII contains a description of the admissions criteria employed at both educational option and screened programs.

TABLE 3—TIEBREAKING FOR ELEMENTARY SCHOOL APPLICANTS IN BOSTON IN 2006–2007

Choice	Deferred acceptance single tie-breaking DA-STB (1)	Deferred acceptance multiple tie-breaking DA-MTB (2)	Student-optimal stable matching (3)	Improvement from DA-STB to student- optimal	Number of students (4)
1	2,251.8 (8.4)	2,157.3 (13.4)	2,256.6 (8.2)	+1	4.6 (2.6)
2	309.8 (10.3)	355.5 (12.0)	307.4 (10.0)	+2	1.2 (1.1)
3	154.9 (7.9)	189.3 (10.1)	154.0 (7.7)	+3	0.5 (0.7)
4	59.7 (5.5)	76.1 (7.0)	58.7 (5.5)	+4	0.3 (0.5)
5	27.4 (4.5)	34.1 (4.8)	27.0 (4.4)	+5	0.0 (0.1)
6	4.9 (1.9)	6.0 (2.5)	4.9 (1.9)	+6	0.0 (0.1)
7	2.6 (1.4)	2.8 (1.6)	2.5 (1.4)	+7	0.0 (0.1)
8	1.9 (1.2)	0.9 (0.9)	1.9 (1.2)	+8	0.0 (0.1)
9	1.2 (1.1)	0.4 (0.6)	1.2 (1.0)	+9	0.0 (0.0)
10	0.3 (0.6)	0.1 (0.2)	0.3 (0.5)		
Unassigned	112.4 (4.6)	104.6 (4.5)	112.4 (4.6)	Total:	6.5

Notes: Data from Boston Public Schools' student assignment process in round 1 in 2006–2007 for students requesting an assignment for grade K2 (elementary school). Column 1 reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with single tiebreaking (DA-STB). Column 2 reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with school-specific tiebreaking. Column 3 reports the average choice received distribution of applicants in a student-optimal stable matching, which is computed from DA-STB followed by stable improvement cycles. Column 4 reports the average number of students and how many places on their rank order list students improve in the student-optimal stable matching relative to the matching produced by DA-STB. Columns 1, 2, 3, and 4 are based on 250 random draws. Simulation standard errors are reported in parentheses.

an efficient matching, the consequence would be blocking pairs involving 34,898.8 students and the majority of these blocking pairs could potentially undermine pairwise incentives.¹⁹

B. Comparison to Boston

In 2005–2006, Boston Public Schools changed their assignment system to one based on the student-proposing deferred acceptance algorithm with a single tiebreaker (see Abdulkadiroğlu et al. 2005, 2006). This gives us an opportunity to investigate the quantitative magnitudes of the constraints imposed by strategy-proofness and stability in a different market, in which the idiosyncratic features of NYC are absent.²⁰

Boston differs from NYC in at least two important ways. First, school choice in Boston takes place at three entry points, elementary, middle, and high school, and involves approximately 3,000–4,000 students at each of these levels. Second, the schools in Boston do not actively rank applicants, and instead employ a priority structure based on location and where siblings are enrolled.

¹⁹ In an earlier version of this paper, we computed the same empirical exercise for 2003–2004, 2004–2005, and 2005–2006 in New York. The quantitative results are of similar magnitudes.

²⁰ The situation in Boston is different from that in NYC in a number of important respects, which we discuss elsewhere. For our present purpose, the main difference, we think, will be that the pattern of preferences in Boston is different than in NYC, due in large part to different geographic and transportation situations, and to the fact that in Boston the preferences are for younger children.




FIGURE 1. DISTRIBUTION OF CHOICES RECEIVED IN ROUND 1 OR 2 BY YEAR

Table 3 reports the average choice received by elementary school applicants in Boston in 2006–2007.²¹ The layout of the table mirrors Table 1. The number of students who receive their top choice under DA-MTB (column 2) is 2,157.3, while 2,251.8 students receive their top choice under DA-STB (column 1), a difference of 94.5 students on average. The overall distribution of DA-MTB and DA-STB is statistically different (two-sided, $p < 0.001$). There is a small difference between DA-STB and a student-optimal matching in Boston. Column 3 reports that 2,256.6 students obtain their first choice in the student-optimal matching, a difference of 4.8 from DA-STB. The difference for this grade and year is statistically significant at conventional levels (one-sided, $p = 0.01$), but for other grades, the difference is no longer significant.

In light of Proposition 2, we also computed how many of the stable matching produced by DA-STB are in fact student-optimal. In contrast to NYC, where none of the matchings produced by DA-STB is a student-optimal stable matching, in Boston we find that 6 percent of the matchings produced by DA-STB are in fact student-optimal stable matchings. However, the small difference between the student-optimal stable matching in Boston and the outcome of DA-STB suggests that DA-STB in Boston is very close to producing student-optimal stable matchings. Therefore, the costs of obtaining strategy-proofness by using the DA-STB instead of the student-optimal stable matching are small.

We also computed the difference between the student-optimal stable matching and an efficient matching in Boston. For grade K2, in 2006–2007, we found on average 0.26 students received a better school. The distribution of student-optimal stable matchings was not statistically different from the distribution of efficient matchings (one-sided, $p = 0.45$). The number of blocking

²¹ In an earlier version of this paper, we computed the same empirical exercise for elementary, middle, and high school applicants in 2005–2006 and 2006–2007 in Boston. The quantitative results are of similar magnitudes.

pairs in the efficient matching for Boston's K2 2006–2007 is 1.07 averaging over 250 draws, and only 12 percent of matchings (30 out of 250 draws) have any blocking pairs. That is, deferred acceptance with single tiebreaking in Boston produces a match that is very close to being a student-optimal stable match, and this is in turn very close to being an efficient match. The costs of strategy-proofness and stability in Boston are therefore low.

III. Discussion

A. What Happened in NYC after the New System Was Adopted?

Since 2008, the new matching system in New York City has assigned nearly half a million students to high school. Figure 1 shows the distribution of assignments received by students after the main round for the first four years of the new mechanism.²² In 2003–2004, 31,021 students received their first choice. In 2006–2007, 6,249 more students received their first choice. Moreover, the cumulative distribution in 2006–2007 has a higher percentage of students receiving one of their top k choices or better (for $k = 1, \dots, 12$) than in 2003–2004. In 2005–2006, 38,727 students receive their top choice, the largest of the first four years, but this is partially due to a larger applicant pool. The percent of applicants is about 43 percent, which is almost the same percentage as in 2006–2007. Note that these improvements took place without any further change in the matching algorithm.

There are many factors that might contribute to the improving performance of the match over time. For instance, guidance counselors and school officials have become more familiar with how the system works. Another change is that the city introduced extensive high school fairs, which help families determine what programs are good matches for their students. While a complete study of the factors that have led to changes in the performance of the match is beyond the scope of this paper, it is worth mentioning two factors that may have played a role.

First, the total capacity of the programs (measured by how many students were assigned) was higher in later years than the first year. While there is considerable variation across programs, many of the highly desirable programs declared more seats in later years. For instance, Townsend Harris High School, one of the most popular screened schools, assigned 216 applicants in 2003–2004 and 346 in 2006–2007. This type of change might be expected from a mechanism that produces a stable outcome, since stability removes most of the incentive to withhold positions to fill later, one of the features of the system before 2003.

The second factor is that many programs are ranking more applicants. From the first to the second year of the match, the number of students who were ranked by programs (counting a given student as many times as he is listed by some program) went from about 121,000 to about 213,000. In the old mechanism, many programs considered only students who listed them as their top choice. In the new mechanism, schools do not see how they were ranked by applicants, so some schools have to consider a larger set of applicants. Staff members at the NYC DOE have told us that after the first year, they strongly encouraged programs to rank more applicants, so that they avoid having vacancies after the match. Many schools seem to have adapted to the new system in these two ways.

²² The data include both eighth and ninth grade applicants, as well as students who rank educational option programs as their top choice, so the numbers are not directly comparable to Table 1.

B. Formal Theory and Practical Design

This paper reports both formal theory and practical design, so it is useful to take a moment to think about the close connections between the two, the gaps between them, and also how these gaps might be reduced.

We study the interaction between strategy-proofness, stability, and student welfare by considering a strategic model in which the choices facing the students and schools are only which preferences to submit. As we discuss in Appendix AIII, this model is a fairly close approximation—although only an approximation—of the situation facing students in NYC when they contemplate how to rank schools. The requirement that matchings be stable implicitly recognizes that students and schools may in reality have wider strategic options that allow them to match with each other outside the matching system, as they did under the prior, unstable mechanism.

Since no stable mechanism is strategy-proof for schools, but there are strategy-proof and stable mechanisms for students, there is no way to completely satisfy the two competing design objectives of stability and strategy-proofness for all agents. The argument for focusing on strategy-proofness for students rather than schools in NYC and in other school choice environments is not only because it lifts a burden from students and families. Economists and social planners like strategy-proofness because it yields valuable student preference data. (And in New York City, schools that were revealed by submitted preferences to be unpopular have been closed.)²³ Market designers and school policymakers like it because it allows simple advice to be given to families about how to participate in the matching system.

For instance, NYC School Chancellor Joel Klein stated (*New York Times*, October 24, 2003) that the “changes are intended to reduce the strategizing parents have been doing to navigate a system that has a shortage of good high schools.” Furthermore, Peter Kerr, another NYCDOE official, wrote (*New York Times*, November 2, 2003): “The new process is a vast improvement.... For example, for the first time, students will be able to list preferences as true preferences, limiting the need to game the system. This means that students will be able to rank schools without the risk that naming a competitive school as their first choice will adversely affect their ability to get into the school they rank lower.” In every year since 2003–2004, the high school directory makes a point to advise families to express their preferences truthfully. In Boston, too, strategy-proofness was a major factor in deciding to move to a new school choice system (cf. Abdulkadiroğlu et al. 2006).

Moreover, the efficiency losses we calculate for students are identifiable only because we have preferences that we can take as a reasonable approximation of true preferences, since they were elicited from an algorithm that is strategy-proof for the large majority of students. As Theorem 1 makes clear, any algorithm that would improve on DA-STB from an efficiency point of view would not be strategy-proof for students. Nothing is yet known about what kinds of preferences one could expect to be strategically submitted to such a mechanism, or what their welfare consequences would be. Consequently, there is room for more work to further illuminate the trade-off between efficiency and strategy-proofness. In particular, for what kinds of preferences will there be substantial efficiency loss with DA-STB (as in New York but not in Boston)? Can these efficiency losses in fact be reduced by alternative mechanisms (cf. Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda 2008; Erdil and Ergin 2008; Clayton Featherstone and Muriel Niederle 2008; and Kesten and M. Utku Ünver 2008)?²⁴

²³ See e.g., Elissa Gootman, “Lafayette Among 5 High Schools to Close,” *New York Times*, December 14, 2006, for a report that cites demand data in the match as a reason for the closing of South Shore High School.

²⁴ Erdil and Ergin (2008) analyze strategic behavior in their stable improvement cycles mechanism when students have symmetric beliefs. In contrast, Abdulkadiroğlu, Che, and Yasuda (2009) study efficiency from an ex ante point

Likewise, since the mechanism in NYC is not strategy-proof for schools, it is also important to understand under what circumstances we expect schools to be able to act strategically. One design decision that reduced how much NYC schools can manipulate was to stop letting them see students' rank order lists (see Appendix AII). Some of the remaining ways in which the mechanism is not strategy-proof for schools are ameliorated by the size of the system (Kojima and Pathak 2009).²⁵

On the other hand, for students, a feature of the information environment in NYC and Boston that market designers cannot change and simply have to work with is that it is easy for parents and students to observe which schools are popular. So, a mechanism that is theoretically manipulable by moving popular schools higher in submitted preferences may in fact also be practically manipulable in this way in school choice environments.²⁶

This comparison between students and schools suggests that a productive way for theory and practical market design to advance together as we consider mechanisms that are not strategy-proof for all agents will be to include explicit discussion about information environments. This will need to be at least partly an empirical discussion, since the most relevant kinds of manipulation may depend on what information is in fact available to participants, and what kinds of manipulation is observed in the field.²⁷

C. Conclusions

This paper provides some of the new theory demanded by the design of school choice mechanisms, and shows empirically that the efficiency costs of strategy-proofness need not be small. But we are much less sure how to measure the costs of *not* having a strategy-proof mechanism.

When we began the design of the NYC high school match in 2003, we had a lot of highly relevant theory to draw on, but as we looked into the particular requirements of the NYC school match, we found ourselves running into problems beyond the available theory, and using data, simulations, and examples to make design decisions for which no reliable theory yet existed. In the present paper, we develop some of the theory we would have liked to have in 2003, and provide support for some of the design decisions made in a more timely way on the basis of those early simulations and examples. In doing so, we raise some new theoretical questions, to which it would be helpful to have answers before the next major design (or redesign) of school matching systems.

As economists are more often asked to design practical markets and allocation mechanisms, this kind of feedback between theory, data, and design seems likely to become familiar.

of view and show that the outcomes of the Boston mechanism may be more desirable from an ex ante standpoint than the outcomes of DA-STB. Abdulkadiroğlu, Che, and Yasuda (2008) introduce signaling of preference intensities to improve ex ante efficiency without harming students' incentives. Kesten and Ünver (2008) introduce a stable and constrained ordinally efficient lottery mechanism. Featherstone and Niederle (2008) consider another ex ante perspective on the trade-offs between incentives and efficiency.

²⁵ Under the old NYC system, which produced unstable outcomes, schools had an incentive not to reveal their full capacity so that they could match afterward with preferred students. This motivation is addressed by the stability of the current system, but no stable mechanism completely eliminates the possibility of manipulation by withholding capacity (Sönmez 1997, 1999). However Kojima and Pathak (2009) show that these incentives become small as the market becomes large in an appropriate way.

²⁶ Eric Budish and Estelle Cantillon (2008) study the Harvard Business School allocation mechanism for MBA class assignments. They show theoretically that it is manipulable by moving popular classes higher in the submitted rankings, and then observe empirically that it seems to be manipulated this way in practice. This is a direction in which it might also be fruitful to investigate mechanisms like stable improvement cycles.

²⁷ In just this way, initial work on kidney exchange focused on strategy-proofness for patients and surgeons and involving patient/donor information (Roth, Sönmez, and Ünver 2004, 2005). After gaining some experience with the new multihospital clearinghouses that were formed, there are new incentive concerns about whether transplant centers will withhold patient-donor pairs that can be matched internally (Roth 2008).

APPENDICES

AI. Proofs from Main Text

PROOF OF PROPOSITION 1:

Suppose that μ is a student-optimal stable matching for the profile (P_I, R_S) . Construct P_S from R_S as follows: for every i, j such that $\mu(i) \neq \mu(j)$ and $i \sim_{\mu(i)} j$, let $i P_\mu(i) j$. The tiebreaking among other students can be done arbitrarily. Then the student-proposing deferred acceptance algorithm produces μ for (P_I, P_S) . That is, μ is the unique student-optimal stable matching for some (P_I, P_S) such that P_S is consistent with R_S . Then, from Roth (1982), there is no matching ν (stable or not) that is individually rational under (P_I, P_S) and $\nu(i) P_i \mu(i)$ for all $i \in I$.

To the contrary, suppose that there is a matching ν that is individually rational under (P_I, R_S) and $\nu(i) P_i \mu(i)$ for all $i \in I$. Since P_S is consistent with R_S , ν is individually rational under (P_I, P_S) as well. Then $\nu(i) P_i \mu(i)$ for all $i \in I$ contradicts with μ being the student-optimal stable matching for (P_I, P_S) .

PROOF OF PROPOSITION 2:

The result will directly follow from showing that, for every preference profile (P_I, R_S) and every student-optimal stable matching μ for that preference profile, there is a single ordering of students r such that $DA^r(P_I, R_S) = \mu$.

Suppose that μ is a student-optimal stable matching at some (P_I, R_S) . For each school s , consider the students who prefer s to their assignment in μ and are ranked highest in R_s among such students,

$$B(s) = \{i : s P_i \mu(i) \text{ and } i R_s j \text{ for every } j \text{ such that } s P_j \mu(j)\}.$$

Let $B(S) = \cup_s B(s)$.

Let A be the set of students in $B(S)$ who are assigned to a school under $\mu : A = \{i \in B(S) : \mu(i) \in S\}$. A stable improvement cycle consists of students $\{i_1, \dots, i_n\} \subset A$, $n \geq 2$, such that $i_l \in B(\mu(i_{l+1}))$ and $\mu(i_{l+1}) P_{i_l} \mu(i_l)$ for $l = 1, \dots, n$ where $l + 1$ is replaced by 1 when $l = n$ (Erdil and Ergin 2008). Since μ is student-optimal, there does not exist any stable improvement cycle (Corollary 1, Erdil and Ergin 2008).

Construct a directed graph with vertices (nodes) A and a directed edge from node i to node j if $\mu(j) P_i \mu(i)$ and $j R_{\mu(j)} i$; that is, i envies j 's school place and the school j is assigned to is indifferent between i and j . Since there is no stable improvement cycle, the directed graph must be acyclic.

We will construct an ordering based on this directed graph utilizing two properties of the graph. First, there is a node with no incoming edges. To see this, suppose that every node has at least one incoming edge. Then, since there is a finite number of nodes, starting from any node we can always leave a node by an incoming edge until we return to a visited node, which leads to a cycle, and a contradiction. Second, after removing a node, we still have a directed acyclic graph, since if there is a cycle after removing a node, there must be a cycle in the original graph.

Construct an ordering $\rho : A \rightarrow \{1, \dots, |A|\}$ as follows: find a node with no incoming edges. Remove this node and all its outgoing edges. Set the value of ρ of this node to $|A|$. By the two properties above, when we remove this node we still have a directed acyclic graph and there will be a node with no incoming edges. From this graph, we iterate the process and set the value of ρ of the next node to $|A| - 1$, and so on.

Next, construct an ordering $r : I \rightarrow \mathbb{N}$ of students as follows:²⁸ For every $j, k \in A$, set $r(j) < r(k)$ if $\rho(j) < \rho(k)$. For every $i \in I - B(S)$ and $j \in A$, set $r(i) < r(j)$. Finally, for every student $l \in B(S) - A$, set $r(j) < r(l)$ for all $j \in A$.

Let $\nu = DA^r(P_I, R_S)$. We will show that $\nu = \mu$. Suppose to the contrary that there exists $j \in I$ such that $\mu(j) \neq \nu(j)$. Since μ is student-optimal (and since students' preferences are strict), there exists some $i \in I$ such that $\mu(i) P_i \nu(i)$. Let $C = \{i : \mu(i) P_i \nu(i)\}$ be the set of students who prefer μ to ν . For any $i_k \in C$, let $i_{k+1} = \nu(\mu(i_k))$, or $\nu(i_{k+1}) = \mu(i_k)$. Since ν is stable, there is no blocking pair, so $i_{k+1} R_{\mu(i_k)} i_k$.

The proof by contradiction has three steps. First, we will show that for any $i_k \in C$, the student who is matched to $\mu(i_k)$ under ν , i_{k+1} , also prefers her assignment under μ to ν and so is in C . Next, we will show that in the course of $DA^r(P_I, R_S)$, student $i \in C$ can be displaced only by some other student in C . Finally, we argue that i_k could have displaced i_{k+1} , only if i_{k+1} were displaced herself. Therefore, no member of C can be rejected first, and so C must be empty.

To show the first step, note that $\mu(i_{k+1}) \neq \nu(i_{k+1})$. Suppose that $\nu(i_{k+1}) = \mu(i_k) P_{i_{k+1}} \mu(i_{k+1})$. Then $i_k R_{\mu(i_k)} i_{k+1}$ by stability of μ so that by construction $i_{k+1} \sim_{\mu(i_k)} i_k$. But then $i_{k+1} \in B(\mu(i_k))$ so that $r(i_k) < r(i_{k+1})$. Then, since $i_k \in C$, $\mu(i_k) = \nu(i_{k+1}) P_{i_k} \nu(i_k)$, which contradicts with stability of $DA^r(P_I, R_S) = \nu$. Therefore $\mu(i_{k+1}) P_{i_{k+1}} \nu(i_{k+1})$, so $i_{k+1} \in C$.

We prove the second step by contradiction. Suppose there is some $i \in C$ and $j \in I - C$ such that $\mu(i) P_j \nu(j)$ and $j R_{\mu(i)} i$. Since $j \in I - C$, we have $\nu(j) R_j \mu(j)$ therefore $\mu(i) P_j \mu(j)$. Then, stability of μ implies that $j \sim_{\mu(i)} i$, which in turn implies that $j \in B(\mu(i))$ so that $r(i) < r(j)$. Therefore, no $i \in C$ is rejected by $\mu(i)$ in $DA^r(P_I, R_S)$ in favor of any $j \in I - C$ such that $\mu(i) P_j \nu(j)$. This implies that every $i \in C$ is rejected by $\mu(i)$ in $DA^r(P_I, R_S)$ in favor of some $i' \in C - \{i\}$.

Finally, in the process of $DA^r(P_I, R_S)$, no $i_k \in C$ will be rejected by $\mu(i_k)$ before $i_{k+1} = \nu(\mu(i_k))$ is rejected by $\mu(i_{k+1})$. Therefore, no $i \in C$ will be rejected by $\mu(i)$ in $DA^r(P_I, R_S)$, so that $C = \emptyset$, i.e., $\nu(i) R_i \mu(i)$. Then optimality of μ implies $\nu(i) = \mu(i)$ for all $i \in I$.

AII. Description of New York City High School Admissions

Each fall, about 90,000 eligible students are asked to work with their families and guidance counselors to fill out a high school application form for the following school year. Nearly 95 percent of students who participate in the admissions process are eighth graders, and the rest are ninth graders, many of whom were enrolled in middle schools that include grade 9. Students receive their application in early October and the rest of the fall is devoted to gathering information about various school options. Typically, in late October or early November, students who are interested in attending one of the Specialized High Schools²⁹ in New York City take the Specialized High School Admissions Test (SHSAT). Between 25,000 and 30,000 students take the SHSAT exam every year. The Specialized High Schools account for between 4,000 and 5,000 school seats a year.

In late November or early December, all students are asked to rank up to twelve schools/programs on the High School Application form. Students who have taken the SHSAT also may submit a separate rank order list expressing their preferences for specialized high schools.

²⁸ Recall that lower numbers means a student is more preferred, i.e., $r(i) = 1$ means that student i is the most preferred student.

²⁹ Specialized High Schools include Brooklyn Technical, High School (HS) of Math and Science at City College, Stuyvesant, Queens High School for Science at York College, High School of American Studies at Lehman College, and the Bronx High School of Science. Fiorello H. LaGuardia High School of Music and Art and Performing Arts is also a Specialized High School where entrance is determined by audition and a review of academic records. After the new matching mechanism was adopted for the 2003–2004 school year, Staten Island Technical High School and Brooklyn Latin School also became Specialized High Schools.

Besides the specialized schools in New York City, there are three types of schools, known as mainstream schools: schools that actively evaluate applicants and submit a ranking to the mechanism; schools that do not evaluate applicants, and instead employ priorities, which are determined not at the school, but by the Department of Education, to order students; and schools at which a fraction of seats are reserved for students who are explicitly ranked by the school, while the rest are automatically categorized into priority groupings set by the DOE.

“Screened” and “audition” schools are examples of the first type of school, at which staff review applicants based on criteria ranging from seventh grade academic performance, attendance, and disciplinary actions to auditions, portfolio submissions, and interviews.³⁰ “Unscreened” schools are examples of the second type of school. Priorities include geographic location, current middle school, or other criteria.³¹ Finally, the third class of schools, “educational option,” are permitted to rank students for half of their positions, and are required to admit students according to priorities for the other half.

Nearly half of all schools are educational option, and more than half of total district capacity is at schools that do not actively rank students. When priorities are used at unscreened and educational option programs, many students fall into the same priority class. For instance, at Forest Hills, an unscreened school where there were 474 seats in 2003–2004, 352 (10.7 percent) of student applicants are from the assignment zone while the remaining 2,937 are from outside.

For each of the school types, neither the students nor the schools see the other party’s rankings. All applicants are processed through a centralized match for mainstream schools, based on the student-proposing deferred acceptance algorithm. First, each student is given a single random number. This number stays with the student across the admissions process and, when necessary, is used to break ties at each program. At each screened and audition programs, the students are ordered based on a strict rank ordering provided by the program. At unscreened programs, the priorities of applicants are determined, and within each priority grouping, students are ordered based on this single random number. Students applying to an educational option program are categorized into one of three categories based on their score on the seventh grade standardized reading test: top 16 percent (high), middle 68 percent (middle), and bottom 16 percent (low). In the algorithm, each educational option programs is split into six separate programs corresponding to **Low**, **Middle**, and **High** categories, and the **Select** and **Random** breakdowns (**LR**, **LS**, **MR**, **MS**, **HR**, **HS**). For the half of seats for which students are ranked by the school, high-level students are ranked above middle- and low-level students at the high program, middle-level students are ranked above high- and low-level students at the middle program, and low-level students are ranked above high- and middle-level students at the low program. For the other half of seats, for which students are not ranked, these same applicants are ordered based on their test score category, and the single random number is used to break ties and construct a strict ordering. The rank order list of a student who applies to an educational option program is modified to rank these six “virtual” subprograms according to the order **HR**, **HS**, **MR**, **MS**, **LR**, **LS**. The preferences are modified in the same way for each student. Finally, any student who scored in the top 2 percent on the seventh grade standardized reading test will automatically be matched to an educational option program if she lists it as her first choice.

³⁰ For example, Townsend Harris in Flushing, NY, one of the most academically rigorous screened schools, evaluates students based on their test scores, attendance, and punctuality. At Townsend Harris, all students are required to have a minimum ninetieth percentile on math and reading standardized tests, as well as a minimum grade point average of 90 in June of seventh grade when being considered for a ninth grade seat.

³¹ For example, the academic comprehensive program at Forest Hills, an unscreened school in Queens, places students who live in an attendance zone near the school in a higher priority class than students from outside the priority zone.

Once applications are submitted, the assignment algorithm is organized into three rounds. The first round exists to accommodate the placement of the students who are given an offer at a Specialized High School. Specialized High School admissions are administered separately, but because of their special position in NYC history, applicants with a specialized school offer also receive an offer from a mainstream school, and have a choice between the two.³² The student-proposing deferred acceptance mechanism produces a mainstream school assignment for each student in the first round. At the end of the first round, only students who have received an assignment at a Specialized High School are informed of their placement from the mainstream schools, together with their specialized assignment, and are asked to pick one of these two offers. The number of such students ranges between 4,000 and 5,000 a year. Since students who receive a placement from a Specialized High School have performed well on the SHSAT, between 50 and 60 percent of these students receive their top choice mainstream school, and between 75 and 85 percent are placed to one of their top three choices. Of those receiving an offer, between 60 and 78 percent of these students take their Specialized High School placement, and of those who do not, 60 to 70 percent take their top choice assignment.

Once the students who obtain a specialized offer are finalized, the capacities at each program are adjusted to remove these placed students. Next, the students who did not receive a specialized offer, and the students who did not apply to Specialized High Schools, are assigned through the student-proposing deferred acceptance mechanism in the second round. The student preferences are the same as those used for the first round. The second round, also known as the main round, produces an assignment for the majority of students. In 2003–2004, 40.2 percent of students were assigned to their first choice at the end of the main round. This number improved to 46.5 percent in the 2006–2007 assignment; 1.1 percent of students in 2003–2004 were assigned to their twelfth and last choice, and this number fell to 0.2 percent in 2006–2007. The distribution of assignments from 2003–2004 to 2006–2007 is presented in the Web Appendix.

If a student is not matched in the main round, she has the opportunity to select up to 12 new schools in the third round, also known as the supplementary round. The schools students may consider are those with remaining capacity, and are almost exclusively unscreened or zoned programs. Students who are not able to be matched to one of their choices will be assigned to an available high school as near to their homes as possible. All students are placed at the conclusion of the supplementary round.

Finally, the office of high school enrollment runs an appeals process each year. In 2003–2004, there were no plans for a formal appeals process. Students who were unhappy with their assignment were invited to write a letter to the Department of Education explaining the reason for their appeal, and a list of programs where they would prefer to be assigned. When DOE received more than 5,000 letters, plans were drawn up for a formal appeals process. While still evolving, the current appeals process invites applicants to file for an appeal after the supplementary round placements, usually in late May. Students are asked to meet with their guidance counselor to obtain an appeals form, which asks them to explain the reason for their appeal and to list up to three programs in order of preference the student wants as a replacement for their assignment.

³² Placement at specialized schools is determined by a serial dictatorship in which students are ordered based on their score on the SHSAT, and an even tiebreaker is used to order students with the same score. That is, the student with the highest test score is placed in her most preferred school, and the student with the next highest test score is assigned to her most preferred school with remaining capacity. A small fraction of seats are assigned at these schools as a result of a Discovery Program, which gives disadvantaged students of demonstrated high potential an opportunity to try the Specialized High School program. These students must be near the SHSAT cutoff, be certified as disadvantaged, and be recommended by their middle schools. At the LaGuardia programs, students are given offers at the programs that accept them after their audition and portfolio submission. For all of the other Specialized High Schools, students can obtain only a single offer. It is possible for a student to receive an offer from both LaGuardia and another Specialized High School, in addition to a mainstream offer.

Students must also give a reason for their appeal as either (1) Change of Home Address, (2) Change of School Location, (3) Entitled Services Unavailable, (4) New Small High School, and (5) Other, and each appeal requires verification of the reason. When students appeal, they also have the option to retain the right to their assigned school if they cannot be placed to one of the three alternatives. In 2006–2007, the NYCDOE evaluated appeals and classified them into either valid or invalid appeals, and processed only those they considered valid.

Further details are contained in the Web Appendix.

AIII. Relationship between the Model and Actual NYC System

This section describes the differences between the actual NYC high school admissions process and the model analyzed in the main text, and describes some empirical features of these differences. In the Web Appendix, we formally demonstrate that, even though high schools have multiple seats, for the questions concerning the choices facing students when they can submit no more than 12 schools, there is no loss of generality in considering a one-to-one matching model.

A. Students May Rank No More Than 12 Choices

The first issue we focus on is that students can rank only 12 schools. The model in the text analyzes the school choice problem when there is no constraint on the number of schools a student may rank. Between 22 and 30 percent of applicants present in Round 1 rank 12 choices.³³ This represents the maximum number of students who are affected by this constraint who may have wanted to rank more than 12 choices. It is also possible that a fraction of these students actually preferred exactly 12 schools, and the number of students ranking 11 choices is a rough approximation of this number. Students who wish to rank more than 12 choices face an incentive problem. However, the preference revelation problem for these students involves selecting which 12 of their preferred schools they submit to the mechanism, according to the order they want.

PROPOSITION (Haeringer and Klijn 2009, Proposition 4.2): *In the student-proposing deferred acceptance mechanism where a student may rank only k schools:*

- *If a student prefers at most k schools, then she can do no better than submitting her true rank order list;*
- *If a student prefers more than k schools, then she can do no better than employing a strategy that selects k schools among the set of schools she prefers to being unassigned and ranking them according to her true preference ordering.*

Last, it is of note that the number of participants receiving their twelfth choice is less than 1 percent of all applicants. Moreover, students ranking 12 choices are rarely unassigned in Round 1 or Round 2.

B. Top 2 Percent Priority at Educational Option Programs

Students who score in the top 2 percent on the standardized reading exam in grade 7 are guaranteed an educational option program *only* if they rank it as their top choice. A student who

³³ A table that shows the distribution of the length of the rank order list in Round 1 across years is presented in the Web Appendix.

does not prefer an educational option program as her top choice may have an incentive to rank it as her top choice so that she receives it. The following proposition summarizes the incentive problem faced by such applicants.

PROPOSITION: *In the student-proposing deferred acceptance mechanism where a student can rank at most k schools, if a student is guaranteed a placement at a school only if she ranks it first, then she can do no better than:*

- *Either ranking that program as her first choice, and submitting the rest of her preferences according to her true preference ordering; or*
- *Submitting her preferences by selecting at most k schools among the set of schools she prefers to being unassigned and ranking them according to her true preference ordering.*

The proposition, the proof of which we defer to the Web Appendix, shows that while the student does not have a dominant strategy, she does have a strategy that is relatively simple. In particular, it shows that if a student does not rank an educational option program as her top choice, she cannot benefit by reversing the order of any school she prefers. In eighth grade, there are between 1,900 and 3,600 students who receive educational option priority. Since these are students who have performed well in grade 7, many of them are competing for a spot at a Specialized High School, and between 30 and 35 percent of applicants accept a Specialized High School offer. Investigating the student rank order lists, between 22 and 36 percent of applicants exercise their educational option priority and rank an educational option school first. In light of the proposition above, this represents an upper bound on the number of students who may have manipulated and ranked an educational option program above a program they prefer. Some of these students may actually prefer an educational option program as their top choice even in the absence of the priority they receive by ranking the school. For a comparison, a larger fraction of between 33 and 41 percent of the population who are not in the top 2 percent category rank an educational option program as their top choice.

C. Multiple Rounds

The last distinction between the model and the actual NYC high school admissions process is due to multiple rounds. First, a student who takes the SHSAT in Round 1, receives a specialized offer, but prefers her mainstream offer may face a strategic disadvantage for having applied in Round 1 at her mainstream school. In this round, the student will be competing against a larger set of students and this may cause her not to receive her top choice mainstream school. Had she, instead, not received an offer from a Specialized High School, she would be processed through the main round after the students who received a specialized offer were removed from the match. Since students who receive a placement from a Specialized High School have performed well on the SHSAT, between 50 and 60 percent of these students receive their top choice mainstream school. Of the remaining fraction of students, between 17 and 34 percent of students take a mainstream offer that is not their top choice over their specialized offer. This corresponds to between 365 and 816 students who might have fared better by not applying to a Specialized High School. Note, it is also possible that students who accepted their specialized offer and who did not receive their top choice mainstream school would have taken their mainstream school if they received their top choice mainstream school. But since these students accepted their specialized offer, we do not know if they preferred their specialized offer to their top choice mainstream school.

The set of schools that are available in Round 3 are those with remaining capacity after the main round. Thus, if a student had ranked any program that is available to applicants in Round 3,

they would have received a placement there. The only students who could be affected by this calculation are those ranking 12 choices. These students might have been better ranking a school with available seats in Round 3 on their rank order list. However, as we have discussed, the fraction of students receiving their twelfth choice is less than 1 percent of all applicants, and students who rank 12 schools are almost never unassigned.

The last potential strategic issue facing students is whether they might obtain a placement for which they can file an appeal. Since under the current system, each appeal is first deemed valid or not, we suspect this is not a major issue.

REFERENCES

- ▶ **Abdulkadiroğlu, Atila.** 2005. "College Admissions with Affirmative Action." *International Journal of Game Theory*, 33(4): 535–49.
- Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda.** 2008. "Expanding 'Choice' in School Choice." Unpublished.
- Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda.** 2009. "Resolving Conflicting Preferences in School Choice: The 'Boston' Mechanism Reconsidered." Available at SSRN: <http://ssrn.com/abstract=1465293>.
- ▶ **Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth.** 2005. "The New York City High School Match." *American Economic Review*, 95(2): 364–67.
- ▶ **Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez.** 2005. "The Boston Public School Match." *American Economic Review*, 95(2): 368–71.
- Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez.** 2006. "Changing the Boston Mechanism: Strategyproofness as Equal Access." Unpublished.
- ▶ **Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003. "School Choice: A Mechanism Design Approach." *American Economic Review*, 93(3): 729–47.
- ▶ **Balinski, Michel, and Tayfun Sönmez.** 1999. "A Tale of Two Mechanisms: Student Placement." *Journal of Economic Theory*, 84(1): 73–94.
- Budish, Eric, and Estelle Cantillon.** 2008. "Strategic Behavior in Multi-Unit Assignment Problems: Theory and Evidence from Course Allocation." Unpublished.
- Calsamiglia, Caterina, Guillaume Haeringer, and Flip Klijn.** Forthcoming. "Constrained School Choice: An Experimental Study." *American Economic Review*.
- ▶ **Day, Robert, and Paul Milgrom.** 2008. "Core-Selecting Package Auctions." *International Journal of Game Theory*, 36(3–4): 393–407.
- ▶ **Dubins, Lester E., and David A. Freedman.** 1981. "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly*, 88: 485–94.
- Ehlers, Lars.** 2006. "Respecting Priorities when Assigning Students to Schools." Unpublished.
- Erdil, Aytek.** 2006. "Two-sided Matching with Ties." PhD Diss. University of Chicago.
- ▶ **Erdil, Aytek, and Haluk Ergin.** 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice." *American Economic Review*, 98(3): 669–89.
- Featherstone, Clayton, and Muriel Niederle.** 2008. "Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation." Unpublished.
- ▶ **Gale, David, and Lloyd Shapley.** 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69: 9–15.
- ▶ **Haeringer, Guillaume, and Flip Klijn.** 2009. "Constrained School Choice." *Journal of Economic Theory*, 144(5): 1921–47.
- Kesten, Onur.** 2004. "Student Placement to Public Schools in the US: Two New Solutions." Unpublished.
- Kesten, Onur, and M. Utku Ünver.** 2008. "A Theory of School Choice Lotteries: Why Ties Should Not Be Broken Randomly." Unpublished.
- ▶ **Kojima, Fuhito, and Parag A. Pathak.** 2009. "Incentives and Stability in Large Two-Sided Matching Markets." *American Economic Review*, 99(3): 608–627.
- Pathak, Parag A.** 2006. "Lotteries in Student Assignment: The Equivalence of Queuing and a Market-Based Approach." Unpublished.
- Pathak, Parag A., and Tayfun Sönmez.** 2008. "Comparing Mechanisms by their Vulnerability to Manipulation." Unpublished.

- ▶ **Roth, Alvin E.** 1982. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research*, 7(4): 617–28.
- ▶ **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.
- ▶ **Roth, Alvin E.** 1985. "The College Admissions Problem Is Not Equivalent to the Marriage Problem." *Journal of Economic Theory*, 36(2): 277–88.
Roth, Alvin E. 1991. "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*, 81(3): 415–40.
- ▶ **Roth, Alvin E.** 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica*, 70(4): 1341–78.
- ▶ **Roth, Alvin E.** 2008. "What Have We Learned from Market Design?" *Economic Journal*, 118(527): 285–310.
Roth, Alvin E., and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review*, 89(4): 748–80.
- ▶ **Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver.** 2004. "Kidney Exchange." *Quarterly Journal of Economics*, 119(2): 457–88.
- ▶ **Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver.** 2005. "Pairwise Kidney Exchange." *Journal of Economic Theory*, 125(2): 151–88.
Roth, Alvin E., and Marilda A. Oliveira Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. New York: Cambridge University Press.
- ▶ **Sönmez, Tayfun.** 1997. "Manipulation Via Capacities in Two-Sided Matching Markets." *Journal of Economic Theory*, 77(1): 197–204.
- ▶ **Sönmez, Tayfun.** 1999. "Can Pre-Arranged Matches Be Avoided in Two-Sided Matching Markets?" *Journal of Economic Theory*, 86(1): 148–56.