

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

Atila Abdulkadiroğlu[†] Parag A. Pathak[‡] Alvin E. Roth[§]

This version: November 2008

forthcoming, *American Economic Review*

Abstract

The design of the New York City (NYC) High School match involved tradeoffs 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 tie breaking– in a student-proposing deferred acceptance mechanism. Any inefficiency associated with a realized tie breaking 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.

*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 co-editor and 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 NSF.

[†]Duke University, email: atila.abdulkadiroglu@duke.edu, address: Department of Economics, Duke University, Durham NC, 27708.

[‡]Society of Fellows and MIT, email: ppathak@mit.edu, address: Department of Economics, Massachusetts Institute of Technology, Cambridge, MA 02142.

[§]Harvard University, email: aroth@hbs.edu, address: Department of Economics, Harvard University, Cambridge, MA 02138 and Harvard Business School, Boston, MA 02163.

In 2003-04, 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 (Atila Abdulkadiroğlu, Parag A. Pathak and Alvin E. 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, Pathak, Roth, and Sönmez 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 tradeoffs among efficiency, stability, and strategy-proofness.

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

¹Gale 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 are 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 non-strict preferences (albeit obliquely), but concentrated on the case of strict preferences.

²Stable 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, Pathak, Roth and Sönmez (2006) for discussion in the context of Boston’s new student assignment mechanism.

on opposite sides of the market may 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 passively rank students. 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 sides 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 are 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

³E.g. the Deputy Chancellor of Schools, quoted in the New York Times (11/19/04): “Before you might have a situation where a school was going to take 100 new children for 9th grade, they might have declared only 40 seats, and then placed the other 60 outside the process.”

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

have to be met for the system to 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 tradeoff 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 a non-empty set of stable matches that are weakly Pareto optimal for agents on that side. It will be useful in what follows to call any member of that set for students *a* student-optimal stable matching.⁶

With the 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 tie breaking. This preserves the stability and strategy-proofness of the student-proposing deferred acceptance algorithm. However, tie breaking introduces artificial stability constraints (since, after tie breaking, 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 tie breaking, the outcome it produces may not in fact be a student-optimal stable matching in terms of the original preferences.

Nevertheless, some forms of tie breaking 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 tie breaking), or to give each student a single number to be used for tie breaking at every school (single tie breaking). Computations with simulated and then actual submitted preferences indicated that single tie breaking had superior welfare prop-

⁵Strategy-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.

⁶See Bob Day and Paul Milgrom (2007) for a similar treatment of the core for auctions with non-transferable utility, and its relation to matching.

erties.⁷ The computations based on the revealed preferences will be discussed in Section 2. 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 tie breaking that cannot be produced by deferred acceptance with single tie breaking, then the matching is not a student-optimal stable matching.⁸

However, single tie breaking can also lead to a matching which 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 tie breaking. 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 tie breaking. That is, the potential inefficiency of student-proposing deferred acceptance with single tie breaking is the cost of strategy-proofness.

This theorem does not speak to the magnitude of the tradeoff to expect in real school choice plans. We take a step towards 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% of 8th graders in our sample could be matched to schools they prefer over their assignment from deferred acceptance with single tie breaking without harming any others, if the same preference information could be elicited by a non-strategy-proof mechanism. A further (approximately) 4,300 students or 5.5% of 8th 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 2 describes the model. Section 3 presents our main theoretical and empirical results. The last section concludes.

1 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 herself,

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

⁸That the manner of tie breaking 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 the strategy-proof top trading cycles mechanism that produces efficient (but not stable) matchings is unaffected by the choice of multiple versus single tie breaking, when all schools are indifferent between students.

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 $sR_i s'$ if $sP_i s'$ 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 , $iR_s s$. 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 $sP_i \mu(i)$, and either $|\mu(s)| < q_s$ and $i \succ_s s$ or $[i \succ_s i'$ 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) which 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 ,

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

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

⁹When 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 (forthcoming) 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.

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 tie breaking procedure to convert the school preferences into strict preferences, and then applying the algorithm above. A tie-breaker, 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 tie breaking and multiple tie breaking rules. A single tie breaking rule uses the same tie-breaker at each school, while a multiple tie breaking rule may use a different tie breaker at each school. For a particular set of tie breakers $\tau = (r_s)_{s \in S}$, 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 tie breaking rule, and DA-MTB where τ is a multiple tie breaking 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 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 v (stable or not) such that $v(i)P_i\mu(i)$ for all $i \in I$.*

1.1 The Model and the New York City HS Match

The model defined above captures the main features of the New York City HS 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 HS 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 only rank up to 12 schools of their choice. Across years of the new system, between 72%-80% of applicants rank fewer than 12 schools. Finally, students in the top 2% 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 affect 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 non-Specialized 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 a obtaining a welfare improvement. We do not include students who are in the top

¹⁰Proofs not in the text are in the appendix.

¹¹Guillaume Haeringer and Flip Klijn (2006) 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. Pathak and Sönmez (2008) show how a student-proposing deferred acceptance mechanism where participants can rank at most k schools is manipulable at each preference profile that the mechanism where participants can rank at most l schools is manipulable, when $k < l$.

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

2% 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

1.2 Motivating Examples

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

Example 1. (Tie-breaking does not always yield student-optimal stable matchings.)

Tie-breaking 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:

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

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 tie breaking 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 tie breaking always yields a student-optimal stable matching. On the contrary, the student-proposing deferred acceptance algorithm with multiple tie breaking 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 tie breaking, but that cannot be produced by deferred acceptance with any single tie breaking is not a student-optimal stable matching.

¹³Specifically, we consider that preference lists are capped at 12, that some schools give priority to students with top 2% 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.

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 :

$$\begin{aligned} s_2 P'_{i_1} s_3 P'_{i_1} s_1 \\ s_1 P_{i_2} s_2 P_{i_2} s_3 \\ s_1 P_{i_3} s_2 P_{i_3} s_3. \end{aligned}$$

Under P' , DA^τ produces

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

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.

2 Welfare Consequences of Tie Breaking

In the course of designing the New York City High School Match, policymakers from the Department of Education were concerned with the fairness of tie breaking. 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 tie breaking might have superior welfare properties to multiple tie breaking. The DOE remained unconvinced until student preferences had already been submitted, and computational experiments could be conducted comparing single and multiple tie breaking using actual data from the first round in 2003-04.

Table 1 presents a comparison of the distribution of school assignments based on the preferences submitted by Grade 8 applicants in New York City 2006-07. The column labeled Deferred Acceptance Single Tie Breaking (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-tie breaking rule, where tie-breakers 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-tie breaking rule (DA-MTB), where tie-breakers 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 7th choice or better is greater under DA-STB, while the number receiving their 8th choice or better is greater under DA-MTB. However, note that more than 89% of applicants receive their 7th 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 tie-breaker 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_I, R_S) , any matching that can be produced by deferred acceptance with some multiple tie breaking, but that cannot be produced by deferred acceptance with any single tie breaking 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,

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

¹⁵A 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 tie breaking 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

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, 2006-07, out of the 250 draws of DA-STB, none are 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 tie breaking rule is not strategy-proof.

Theorem 1. For any tie breaking 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.

Claim: Suppose that ν dominates $\mu = DA^\tau(P_I; R_S)$ for a given tie-breaking 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 tie-breaking rule r such that φ dominates DA^τ . There exists a profile P_I such that

$$\begin{aligned} &\varphi_i(P_I; R_S)R_iDA_i^\tau(P_I; R_S) \text{ for all } i \in I, \text{ and} \\ &\varphi_i(P_I; R_S)P_iDA_i^\tau(P_I; R_S) \text{ for some } i \in I. \end{aligned}$$

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 $\hat{s}_iP_i s_i$.

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, $s_i = DA_i^\tau(P_I; R_S)R_iDA_i^\tau(P'_I; R_S)$, 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 $\varphi_i(P'_I; R_S) = i$.

the $n!$ single tie breaking rules. This shows that the set of single tie breaking rules is the smallest such set.

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. \diamond

In other words, any inefficiency associated with a realized tie breaking 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 tie breaking 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 tie breaking. 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 tie breaking, with or without strict preferences. Finally, neither serial dictatorship nor top trading cycles dominates deferred acceptance with single tie breaking (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 tradeoff between strategy-proofness and efficiency, a tradeoff that hinges on stability. Deferred acceptance with tie breaking, 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,

¹⁶The exact details on the implementation of stable improvement cycles are contained in the Online Appendix.

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

2.1 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 Supplementary Appendix. 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 is the student-optimal stable 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 2006-07, 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% 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 have incentives to circumvent the match. If the blocking pair involves a student and a school that only uses 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

¹⁷In an earlier version of this paper, we computed the same empirical exercise for 2003-04, 2004-05, and 2005-06 in New York. The quantitative results are of similar magnitudes.

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% of the blocking pairs involve a student and an “Educational option” program, while about 7% 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% are involved in a blocking pair at one school. 6,707.8 or 19.2% 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 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

2.2 Comparison to Boston

In 2005-06, Boston Public Schools changed their assignment system to one based on the student-proposing deferred acceptance algorithm with a single tie breaker (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.

Table 3 reports the average choice received by elementary school applicants in Boston

¹⁸Appendix A2 contains a description of the admissions criteria employed at both Educational Option and Screened programs.

¹⁹In an earlier version of this paper, we computed the same empirical exercise for 2003-04, 2004-05, and 2005-06 in New York. The quantitative results are of similar magnitudes.

²⁰The 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.

in 2006-07.²¹ 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 the student-optimal stable 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 are student-optimal stable matchings, in Boston we find that 6% 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 2006-07, 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 pairs in the efficient matching for Boston's K2 2006-07 is 1.07 averaging over 250 draws, and only 12% of matchings (30 out of 250 draws) have any blocking pairs. That is, deferred acceptance with single tie breaking 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.

3 Discussion

3.1 What happened in NYC after the new system was adopted?

As of 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 stu-

²¹In an earlier version of this paper, we computed the same empirical exercise for elementary, middle and high school applicants in 2005-06 and 2006-07 in Boston. The quantitative results are of similar magnitudes.

dents after the Main Round for the first four years of the new mechanism.²² In 2003-04, 31,021 students received their first choice. In 2006-07, 6,249 more students received their first choice. Moreover, the cumulative distribution in 2006-07 has a higher percentage of students receiving one of their top k choices or better (for $k = 1, \dots, 12$) than in 2003-04. In 2005-06, 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%, which is almost the same percentage as in 2006-07. 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 2003-04 and 346 in 2006-07. 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 only considered 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 have 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.

3.2 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,

²²The data includes both 8th and 9th 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.

and also about 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 what preferences to submit. As we discuss in Appendix A3, this model is a fairly close approximation—although only an approximation—to 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 revealed by submitted preferences to be unpopular have been closed.²³) Market designers and school policy-makers 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 (NY Times 10/24/03) 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 (NY Times 11/3/03): “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 2003-04, 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 only identifiable 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

²³See 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.

such a mechanism, or what their welfare consequences would be. Consequently, there is room for more work to further illuminate the tradeoff 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, Che and Yasuda 2008, Erdil and Ergin 2008, and Kesten and Ünver 2008)²⁴

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 A2). 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, forthcoming).²⁵

On the other hand, for students, a feature of the information environment in NYC and Boston that market designers can not 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 manipulations may depend on what information is in fact available to participants, and what kinds of manipulations are observed in the field.²⁷

²⁴Erdil and Ergin (2008) take some preliminary steps in this direction by analyzing strategic behavior in their stable improvement cycles mechanism when students have symmetric beliefs. In contrast, Abdulkadiroğlu, Che and Yasuda (2008) and Kesten and Ünver (2008) study efficiency from an ex-ante point of view. In particular, Abdulkadiroğlu, Che and Yasuda (2008) introduce signaling of preference intensities in tie-breaking to improve efficiency from an ex-ante point of view 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 tradeoffs between incentives and efficiency.

²⁵Under the old NYC system, which produced unstable outcomes, schools had an incentive not to reveal their full capacity so that they could match afterwards 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 (forthcoming) show that these incentives become small as the market becomes large in an appropriate way.

²⁶Budish and 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.

²⁷In 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

3.3 Conclusions

This paper fills in 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.

with the new multi-hospital 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).

References

- [1] Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth (2005), “The New York City High School Match,” *American Economic Review*, Papers and Proceedings, May 2005.
- [2] Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez (2005), “The Boston Public School Match,” *American Economic Review*, Papers and Proceedings, May 2005.
- [3] Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez (2006), “Changing the Boston Mechanism: Strategyproofness as Equal Access,” Unpublished mimeo, Harvard University.
- [4] Abdulkadiroğlu, Atila and Tayfun Sönmez (2003), “School Choice: A Mechanism Design Approach,” *American Economic Review*, 93: 729-747.
- [5] Abdulkadiroğlu, Atila (2005), “College Admissions with Affirmative Action,” *International Journal of Game Theory*, 33: 535-549.
- [6] Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda (2008), “Expanding “Choice” in School Choice,” Unpublished mimeo, Duke University.
- [7] Balinski, Michel and Tayfun Sönmez (1999), “A Tale of Two Mechanisms: Student Placement,” *Journal of Economic Theory*, 84: 73-94.
- [8] Budish, Eric and Estelle Cantillon (2008), “Strategic Behavior in Multi-Unit Assignment Problems: Theory and Evidence from Course Allocation,” Unpublished mimeo, Harvard.
- [9] Day, Robert and Paul Milgrom (2007), “Incentives in Core-Selecting Auctions,” *International Journal of Game Theory*, 1432-1470.
- [10] Dubins, Lester E. and David A. Freedman (1981), “Machiavelli and the Gale-Shapley Algorithm,” *American Mathematical Monthly*, 88: 485-494.
- [11] Ehlers, Lars (2006), “Respecting Priorities when Assigning Students to Schools,” CIREQ, February 2006.
- [12] Erdil, Aytok (2006), “Two-sided Matching with Ties.” University of Chicago, PhD Dissertation submitted to the Department of Mathematics.
- [13] Erdil, Aytok and Haluk Ergin (2008), “What’s the Matter with Tie-breaking? Improving Efficiency in School Choice,” *American Economic Review*, 98: 669-689.
- [14] Featherstone, Clayton and Muriel Niederle (2008), “Manipulation in School Choice Mechanisms,” working paper, Stanford.
- [15] Gale, David and Lloyd Shapley (1962), “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69: 9-15.

- [16] Haeringer, Guillaume and Flip Klijn (2006), “Constrained School Choice,” Unpublished mimeo, Universitat Autònoma de Barcelona.
- [17] Kesten, Onur (2004), “Student Placement to Public Schools in the US: Two New Solutions,” Unpublished mimeo, University of Rochester.
- [18] Kesten, Onur and Ünver, M. Utku (2008), “A Theory of School Choice Lotteries: Why Ties Should Not Be Broken Randomly,” Unpublished mimeo, Carnegie Mellon University and Boston College
- [19] Kojima, Fuhito and Parag A. Pathak (forthcoming), “Incentives and Stability in Large Two-Sided Matching Markets,” *American Economic Review*.
- [20] Pathak, Parag A. (2006), “Lotteries in Student Assignment,” Unpublished mimeo, Harvard University.
- [21] Pathak, Parag A. and Tayfun Sönmez (2008), “Comparing Mechanisms by their Vulnerability to Manipulation,” Unpublished mimeo, MIT.
- [22] Roth, Alvin E. (1982), “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 7: 617-628.
- [23] 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: 991-1016.
- [24] Roth, Alvin E. (1985), “The College Admissions Problem is Not Equivalent to the Marriage Problem,” *Journal of Economic Theory*, 36: 277-288.
- [25] Roth, Alvin E. (1991), “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.” *American Economic Review*, 81: 415-440.
- [26] Roth, Alvin E. (2002), “The Economist as Engineer: Game Theory, Experimental Economics, and Computation as Tools of Design Economics.” *Econometrica*, 70: 1341-1378.
- [27] Roth, Alvin E. (2008), “What Have We Learned from Market Design?” *Economic Journal*, 118: 285-310.
- [28] Roth, Alvin E. and Elliot Peranson (1999), “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design.” *American Economic Review*, 89: 748-780.
- [29] Roth, Alvin E., Tayfun Sönmez and M. Utku Ünver (2004), “Kidney Exchange.” *Quarterly Journal of Economics*, 119: 457-488.
- [30] Roth, Alvin E., Tayfun Sönmez and M. Utku Ünver (2005), “Pairwise Kidney Exchange,” *Journal of Economic Theory*, 125: 151-188.

- [31] Roth, Alvin E. and Marilda Sotomayor (1990), *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge: Cambridge University Press.
- [32] Sönmez, Tayfun (1997), “Manipulation via Capacities in Two-Sided Matching Markets,” *Journal of Economic Theory*, 77: 197-204.
- [33] Sönmez, Tayfun (1999), “Can Pre-Arranged Matches be Avoided in Two-Sided Matching Markets?” *Journal of Economic Theory*, 86: 148-156.

Appendices

A1. 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 $iP_{\mu(i)}j$. The tie breaking 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) . \diamond

Proof of Proposition 2. The result will directly follow from showing 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,

$$B(s) = \{i : sP_i\mu(i) \quad \text{and} \quad iR_sj \quad \text{for every } j \text{ such that } sP_j\mu(j)\}.$$

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

Let A be the set of students in $B(S)$ who are assigned to a school under μ : $A = \{i \in B(S) : \mu(i) \in S\}$. 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 $jR_{\mu(j)}i$; 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 are 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, then 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 only be displaced 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 $i_kR_{\mu(i_k)}i_{k+1}$ 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 that there is some $i \in C$ and $j \in I - C$ such that $\mu(i)P_j\nu(j)$ and $jR_{\mu(i)}i$. 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 $i \in I$. \diamond

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

A2. 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% of students who participate in the admissions process are 8th graders, and the rest are 9th 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.

Besides the specialized schools in New York City, there are three types of schools, known as mainstream schools: 1) schools that actively evaluate applicants and submit a ranking to the mechanism, 2) 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 3) 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 7th 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

²⁹Specialized High Schools include Brooklyn Technical, High School (HS) of Math and Science at City College, Stuyvesant, Queens HS for Science at York College, HS of American Studies at Lehman College, and the Bronx HS of Science. Fiorello H. LaGuardia HS of Music & 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 2003-04 school year, Staten Island Technical High School and Brooklyn Latin School also became Specialized High Schools.

³⁰For 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 90th percentile on Math and Reading standardized tests as well as a minimum grade point average of 90 in June of 7th grade when being considered for a 9th grade seat.

³¹For example, the academic comprehensive program at Forest Hills, an unscreened school in Queens places

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 2003-04, 352 (10.7%) 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 7th grade standardized reading test: top 16% (High), middle 68% (Middle), and bottom 16% (Low). In the algorithm, each Educational Option programs is split into six separate programs corresponding to **L**ow, **M**iddle and **H**igh categories, and the **S**elect and **R**andom 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 two percent on the 7th grade standardized reading test will automatically be matched to an Educational Option program if she lists it as her first choice.

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 Schools admissions are

students who live in an attendance zone near the school in a higher priority class than students from outside the priority zone.

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% of these students receive their top choice mainstream school, and between 75-85% are placed to one of their top three choices. Of those receiving an offer, between 60 and 78% of these students take their Specialized High School placement, and of those who do not, 60-70% take their top choice assignment.

Once the students who obtain a Specialized offers 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 2003-04, 40.2% of students were assigned to their first choice at the end of the Main round. This number improved to 46.5% in the 2006-07 assignment. 1.1% of students in 2003-04 were assigned to their 12th and last choice and this number fell to 0.2% in 2006-07. The distribution of assignments from 2003-04 to 2006-07 is presented in the Online 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

³²Placement at Specialized schools is determined by a serial dictatorship in which students are ordered based on their score on the SHSAT, and an even tie breaker is used to order students with the same score. That is, the student with the highest test score is placed to 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 only obtain 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 are placed at the conclusion of the Supplementary Round.

Finally, the office of high school enrollment runs an appeals process each year. In 2003-04, 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. Students must also explain 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 a student appeals, 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 2006-07, the NYC DOE evaluated appeals and classified them into either valid or invalid appeals, and processed only those they considered valid.

Further details are contained in the Online Appendix.

A3. 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 Online 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.

A3.1. Students may rank no more than 12 choices

The first issue that we focus on is that students can only rank 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% 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

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

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, Lemma 8.1.): In the student-proposing deferred acceptance mechanism where a student may only rank k schools,

- if a student prefers fewer than 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 which selects k schools among the set of schools she prefers to being unassigned and ranking them according to her true preference ordering.

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

A3.2. Top 2% Priority at Educational Option Programs

Students who score in the top two 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 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 submit 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 Online Appendix, shows that while the student does not have a dominant strategy, she does have a strategy which 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 8th 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 these students are competitive for a spot at a Specialized High School, and between 30-35% of applicants accept a Specialized High School offer. Investigating the student rank order lists, between 22-36% of applicants exercise their ed-opt 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 33-41% of the population who are not in the top 2% category rank an Educational Option program as their top choice.

A3.3. 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 to not 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% of these students receive their top choice mainstream school. Of the remaining fraction of students, between 17-34% of students take a mainstream offer which is not their top choice over their specialized offer. This corresponds to between 365-816 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, then 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 12th choice is less than 1% of all applicants, and students who rank 12 schools are almost

never unassigned.

The last potential strategic issue facing students is whether they may 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.

Table 1— Tie-breaking for Grade 8 Applicants in NYC in 2006-07

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	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 2006-07 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 tie-breaking (DA-STB). Column (2) reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with school-specific tie breaking. 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 the matching produced by DA-STB. Columns (1), (2), (3) and (4) are based on 250 random draws. Simulation standard errors are reported in parenthesis.

Table 2— Welfare Consequences of Stability for Grade 8 Applicants in in 2006-07

Choice	Student-Optimal Stable Matching (1)	Efficient Matching (2)	Improvement from Student-Optimal Stable Matching	Number (3)	k	Count of Students with k 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

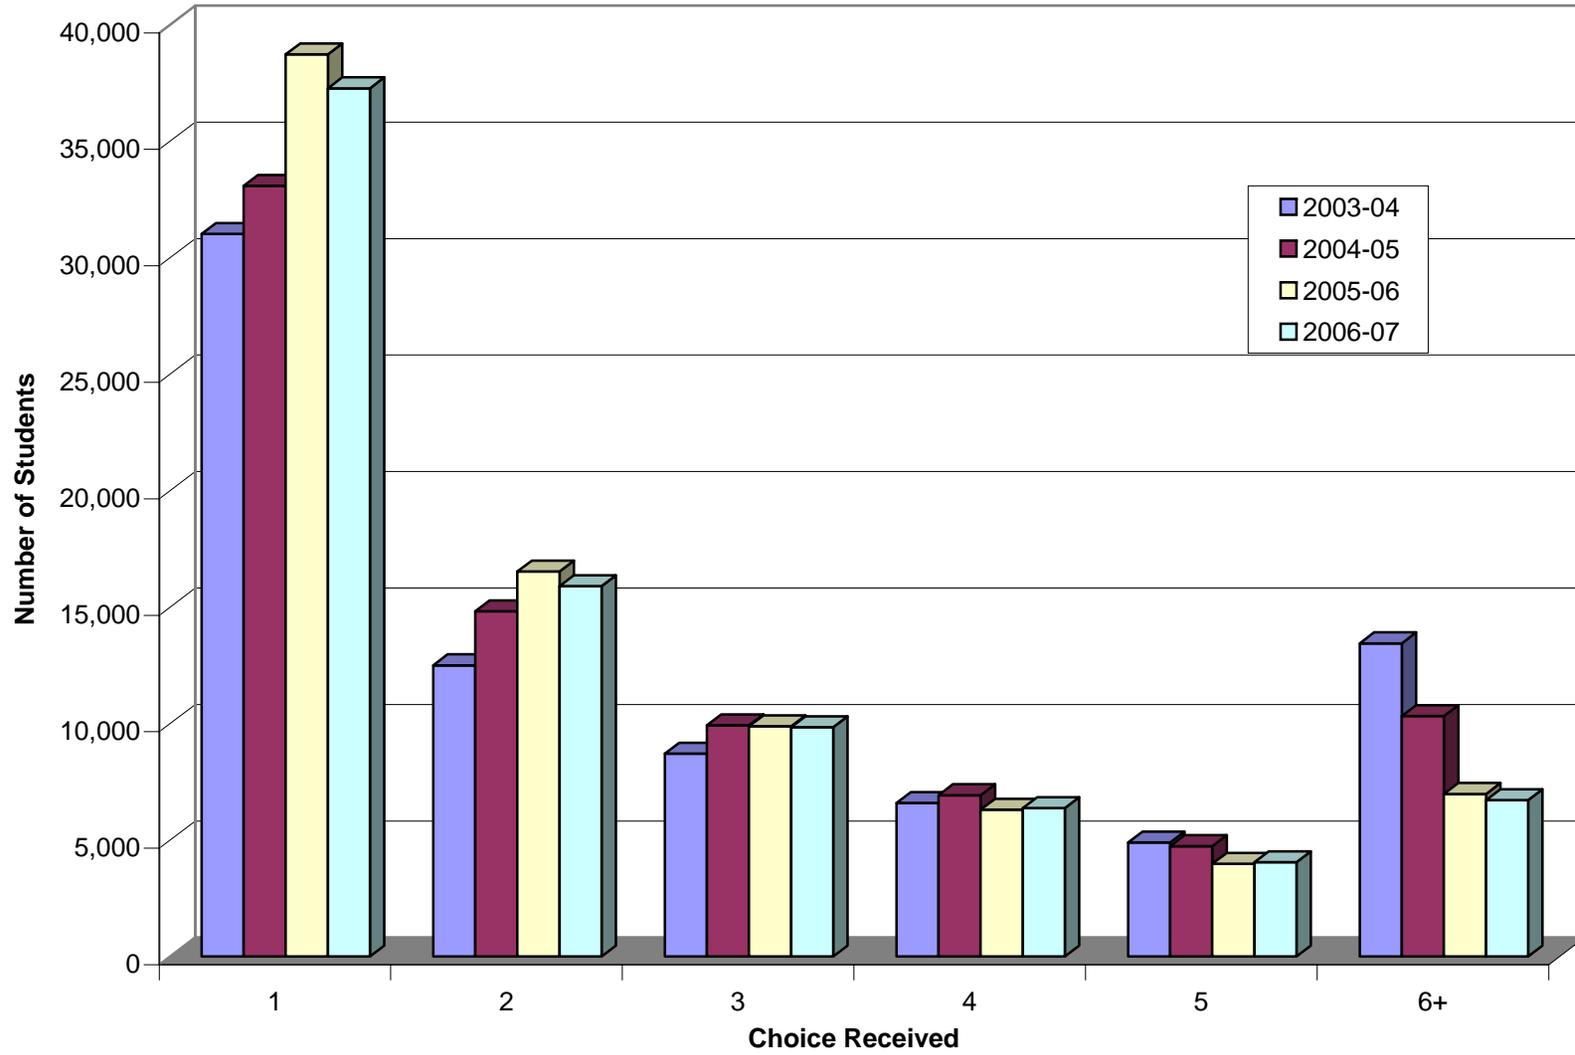
Notes: Data from the main round of the New York City High School Admissions Process in in 2006-07 for students requesting an assignment for Grade 9 (high school). Column (1) reports the average choice received distribution of applicants in a student-optimal stable matching, which is computed from DA-STB 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 parenthesis.

Table 3— Tie-breaking for Elementary School Applicants in Boston in 2006-07

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 School’s student assignment process in Round 1 in 2006-07 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 tie-breaking (DA-STB). Column (2) reports the average choice received distribution of applicants from the student-proposing deferred acceptance algorithm with school-specific tie breaking. 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 parenthesis.

Figure 1: Distribution of Choices Received in Round 1 or 2 by Year



Online Appendix

This appendix provides supplementary material for “Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match.” The numbering of sections parallels portions of the Appendix (i.e. A2’ corresponds to Appendix A2).

A2’: Description of New York City High School Admissions

The following table summarizes the distribution of assignments from students in Round 1 and Round 2.

Table A2: Distribution of Assignments
from Round 1 and Round 2 in New York City

Choice	2003-04	2004-05	2005-06	2006-07
1	31,021 (40.2%)	33,083 (41.4%)	38,727 (47.0%)	37,270 (46.5%)
2	12,504 (16.2%)	14,818 (18.6%)	16,524 (20.1%)	15,898 (19.8%)
3	8,713 (11.3%)	9,929 (12.4%)	9,882 (12.0%)	9,845 (12.3%)
4	6,587 (8.5%)	6,927 (8.7%)	6,308 (7.7%)	6,369 (7.9%)
5	4,893 (6.3%)	4,739 (5.9%)	3,984 (4.8%)	4,051 (5.1%)
6	3,652 (4.7%)	3,415 (4.3%)	2,699 (3.3%)	2,532 (3.2%)
7	2,682 (3.5%)	2,246 (2.8%)	1,603 (1.9%)	1,629 (2.0%)
8	2,160 (2.8%)	1,651 (2.1%)	1,054 (1.3%)	978 (1.2%)
9	1,635 (2.1%)	1,149 (1.4%)	688 (0.8%)	681 (0.8%)
10	1,376 (1.8%)	786 (1.0%)	440 (0.5%)	436 (0.5%)
11	1,063 (1.4%)	600 (0.8%)	291 (0.4%)	275 (0.3%)
12	877 (1.1%)	476 (0.6%)	205 (0.2%)	184 (0.2%)

A3’: Relationship between the Model and Actual NYC System

This section describes the relationship between the model in the main text and the actual New York City high school assignment process.

A3.1': Students may rank no more than 12 choices

The following table shows the distribution of the length of the rank order list in Round 1 across years.

Table A3.1: Length of Applicant ROLs in Round 1

	2003-04	2004-05	2005-06	2006-07
1	7,907 (8.47)	6,123 (6.59)	6,648 (7.18)	6,786 (7.48)
2	4,967 (5.32)	4,369 (4.70)	4,808 (5.20)	4,683 (5.16)
3	6,332 (6.79)	6,048 (6.51)	6,694 (7.23)	6,615 (7.29)
4	6,722 (7.2)	6,697 (7.21)	7,670 (8.29)	7,490 (8.26)
5	6,817 (7.31)	7,159 (7.71)	8,109 (8.76)	8,098 (8.93)
6	6,504 (6.97)	7,480 (8.05)	8,194 (8.86)	8,115 (8.95)
7	5,607 (6.01)	6,320 (6.81)	6,990 (7.55)	7,026 (7.75)
8	5,386 (5.77)	5,798 (6.24)	6,123 (6.62)	6,336 (6.99)
9	4,808 (5.15)	4,841 (5.21)	4,971 (5.37)	5,286 (5.83)
10	5,741 (6.15)	4,952 (5.33)	4,804 (5.19)	5,025 (5.54)
11	8,647 (9.27)	5,561 (5.99)	5,261 (5.69)	5,269 (5.81)
12	23,875 (25.59)	27,524 (29.64)	22,260 (24.06)	19,952 (22.00)

A3.2': Top 2% Priority at Educational Option Programs

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 submit 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.

Proof: Consider a student with a guaranteed placement at a school. Given her preferences, partition her set of strategies into two sets: The first set consists of preference list of at most k schools that rank her guaranteed school as first choice. The second set consists of all other preference lists of at most k schools. We will show that her optimal strategy lies either in the first or the second set.

She is indifferent among all the preference lists in the first set, as she is guaranteed her guaranteed school by submitting any of those preference lists. So, there is no loss of generality in considering a particular strategy from this set, namely the one that ranks the guaranteed school as her first choice, and ranks the rest of her preferences according to her true preference ordering.

By the proposition above, her optimal strategy among the ones in the second set ranks schools in her true preference ordering, yielding the desired conclusion. \diamond

A4': Ex Ante Comparison of DA-STB and DA-MTB

Let p_i^k be the probability that student i receives her k th choice. An allocation is a vector of probabilities $p_i = (p_i^1, \dots, p_i^n)$ for each item on the rank order list P_i such that $\sum_{k=1}^n p_i^k = 1$. We will say that an allocation p_i **ordinally dominates** an allocation q_i for student i , if for all $m = 1, \dots, n$,

$$\sum_{k=1}^m p_i^k \geq \sum_{k=1}^m q_i^k,$$

with strict inequality for some m . An allocation vector $p = (p_i)$ stochastically dominates $q = (q_i)$ if p_i stochastically dominates q_i for some i , and does no worse for all i .

Proposition. There is no ordinal dominance relationship between DA-STB and DA-MTB.

Proof. We present an example where there is no ordinal dominance relationship. Consider an economy with three students i_1, i_2, i_3 and three schools, s_1, s_2, s_3 , each with one seat. Suppose student preferences are:

$$\begin{aligned} i_1 &: s_1 \succ s_2 \succ s_3 \\ i_2 &: s_3 \succ s_1 \succ s_2 \\ i_3 &: s_1 \succ s_3 \succ s_2 \end{aligned}$$

Suppose three schools are indifferent between all applicants. Then DA-STB induces the following distribution over matchings:

$$\frac{1}{3} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_3 & s_2 \end{pmatrix} + \frac{1}{2} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_3 & s_1 \end{pmatrix} + \frac{1}{6} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix}$$

DA-MTB induces the following distribution over matchings:

$$\frac{1}{4} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_3 & s_2 \end{pmatrix} + \frac{1}{2} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_3 & s_1 \end{pmatrix} + \frac{1}{6} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_1 & s_2 & s_3 \end{pmatrix} + \frac{1}{12} \cdot \begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_3 \end{pmatrix}.$$

Student i_3 is more likely to receive her first or second choice under DA-MTB than DA-STB, while student i_1 is more likely to receive her first or second choice under DA-STB than DA-MTB. Therefore, there is no ordinal dominance relationship between the two mechanisms. \diamond

A5': Implementation of the Stable Improvement Cycles Algorithm

This section describes the Stable Improvement Cycles algorithm of Erdil and Ergin (2008) and explain its implementation. The data we use for New York is for all 8th grade applicants in Round 1 of the New York City High School match. If an applicant is marked as a student who receives top 2% priority at an Educational Option school and ranks the school as their top choice, we do not include the applicant in these tables. The data we use for Boston is all elementary (Grade K2), middle (Grade 6), and high school (Grade 9) applicants in Round 1 for 2005-06 and 2006-07, when Boston employed a student-proposing deferred acceptance algorithm to place students. These students will be receiving their top choice and thus will not be affected by a stable improvement cycle. If an applicant ranked 12 schools, we work with the stated rank order list. Given a stable matching μ , define the following: Let A_s be the set of students assigned to school s under μ ; B_s be the set of students who are ranked highest by s among all who prefer s to their assignment. Formally,

$$A_s = \{i \in I : \mu(i) = s\},$$

$$B_s = \{i \in I : sP_i\mu(i) \text{ and } iR_sj \text{ for all } j \text{ such that } sP_j\mu(j)\}.$$

A stable improvement cycle is a list of distinct students $i_1, \dots, i_n \equiv i_0$, $n \geq 2$, such that $\mu(i_l) \in S$ and $i_l \in B_{\mu(i_{l+1})}$ for $l = 0, \dots, n - 1$. We implement a stable improvement cycle by forming a new matching μ' as

$$\mu'(i) = \begin{cases} \mu(i) & \text{if } i \notin \{i_0, \dots, i_{n-1}\} \\ \mu(i_{l+1}) & \text{if } i = i_l \text{ for some } l = 0, \dots, n - 1 \end{cases}$$

We start with a single tie breaking rule and matching produced by the associated DA-STB.

Given a stable matching, we construct a directed graph as follows: The nodes of the graph are schools. We draw an edge from school s to school s' if there is a student i such that $\mu(i) = s$ and $i \in B_{s'}$. We also associate that edge with the set of all such students, denoted by $E_{ss'}$. Formally,

$$E_{ss'} = \{i \in I : \mu(i) = s \text{ and } i \in B_{s'}\}$$

Students in $E_{ss'}$ are sorted according to the given tie breaking rule. Let E_s be the set of edges originating from s . During a search for a cycle, schools are tried in the alphabetical order. In particular, we start the search for a stable improvement cycle with the first school in the alphabetical order. If we cannot find a cycle after starting the search with a school, we restart the search with the next school in the alphabetical order. When we reach a school s in our search, we continue our search with the schools in E_s in the alphabetical order. When a student is to be moved from s to s' in cycle, the last student i in $E_{ss'}$ is moved from s to s' . Then i is removed from all $E_{ss''}$ for every $s'' \in S \setminus \{s\}$. We find and implement all the cycles in the graph. Then we repeat these steps with the new matching until no cycle is found.

A6': Tradeoff between Stability and Efficiency

As we mention in the text, we take only students' preferences into account for welfare considerations. In order to measure the cost of stability associated with a student-optimal stable matching μ , we find a Pareto efficient matching that Pareto dominates μ .

If a matching is not Pareto efficient, we find a Pareto efficient matching that Pareto dominates it from the perspective of students via Gale's top trading cycle algorithm as follows:

If a matching of students to schools, μ , is not Pareto efficient, then there exists a cycle of students $i_1, i_2, \dots, i_{n+1} \equiv i_1$, $n \geq 2$, such that i_l prefers i_{l+1} 's matched school over her match, that is $\mu(i_{l+1}) P_{i_l} \mu(i_l)$, $l = 1, \dots, n$. A new matching $TTC(\mu)$ can be obtained by picking an arbitrary cycle $i_1, i_2, \dots, i_{n+1} \equiv i_1$, and transferring every i_l to i_{l+1} 's matched school:

$$TTC(\mu)(i) = \begin{cases} \mu(i_{l+1}) & \text{if } i = i_l \text{ for some } l = 1, \dots, n \\ \mu(i) & \text{otherwise} \end{cases}$$

$TTC(\mu)$ Pareto dominates μ . Therefore, a Pareto efficient matching that Pareto dominates μ can be found as the limit of $\mu^{t+1} = TTC(\mu^t)$ where $\mu^0 = \mu$. The limit is obtained in finite steps by finiteness of the model.