American Economic Review 101 (February 2011): 1–14
http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.1.1

AQ1

Resolving Conflicting Preferences in School Choice: The "Boston Mechanism" Reconsidered

By Atila Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda*

Public school choice—the initiative for broadening families' access to schools beyond their residence area—has broad public support and has been increasingly adopted across the US and abroad. Yet, how to operationalize school choice, i.e., what procedure should be used to assign students to schools, remains hotly debated.

An important debate centers around the procedure known as the "Boston mechanism," which was used by Boston Public Schools (BPS) until the 2004–2005 school year to assign K–12 pupils to the city schools. Beginning with the seminal article by Abdulkadiroğlu and Tayfun Sönmez (2003), authors recognized problems with the mechanism, and BPS ultimately decided in 2005 to replace the mechanism with the student-proposing deferred acceptance (henceforth DA) mechanism, originally proposed by David E. Gale and Lloyd S. Shapley (1962). While the switch has received some academic support, it was met with resistance from some parents. Most importantly, the Boston mechanism still remains among the most popular school choice programs. It is thus sensible to gain fuller understanding of the two mechanisms before a similar switch is recommended more widely. In this context, the current paper provides a new perspective on the debate and, in so doing, cautions against hasty rejection of the Boston mechanism in favor of the DA.

The criticisms of the Boston mechanism are multifaceted, but they are traced to its poor incentive property. In the Boston mechanism, the seats of each school are assigned *according to the rank students assign that school*; those who rank it first are accepted first, followed by those who rank it second *only* when seats are available, and so forth. Those who rank a school the same are assigned in the order of their priorities at that school, ties being broken randomly, but students who rank it more highly have strict priority at that school ahead of those who don't. This feature implies that students may not prefer to rank schools truthfully. In particular, they may refrain from top-ranking a popular school: top-ranking such a school will not improve their odds with that school appreciably but may rather jeopardize their shot at their second, or even less preferred school, which could have been available to them had they top-ranked it. That strategic ranking may be beneficial

* Abdulkadiroğlu: Department of Economics, Duke University, Durham, NC 27708 (e-mail: atila.abdulkadiroglu@duke.edu); Che: Department of Economics, Columbia University, New York, NY 10027, and YERI, Yonsei University (e-mail: yeonkooche@gmail.com); Yasuda: National Graduate Institute for Policy Studies, Tokyo 106-8677, Japan (e-mail: yosuke.yasuda@gmail.com). We are grateful to Eric Budish, Caterina Calsamiglia, Jinwoo Kim, Fuhito Kojima, Muriel Niederle, Antonio Romero-Medina, and Jay Sethuraman for their comments. Atila Abdulkadiroğlu gratefully acknowledges an NSF-CAREER award and an Alfred P. Sloan Research Fellowship. Yeon-Koo Che gratefully acknowledges NSF Grant SES 0721053 and KRF's World Class University Grant R32-2008-000-10056-0. Yosuke Yasuda gratefully acknowledges the financial support of the JSPS Grant-in-aid for Science Research Start-Up 20830024.

presents some difficulties. First, it is not clear how families should strategize their rankings of schools. Second, there is a potential issue of equity since participants who act naïvely or honestly may be disadvantaged by those who are strategically sophisticated.

The DA mechanism avoids the incentive problem by making truthful ranking a dominant strategy for the participants, a property known as "strategy-proofness" (Lester E. Dubins and David A. Freedman 1981; Alvin E. Roth 1982). In the first round of the DA, students apply to their top-ranked schools, and each school *tentatively* admits up to its capacity from its applicants according to its rankings of students, ties being broken randomly, and reject the others. In each subsequent round, those rejected in the previous round apply to their next highest-ranked school, and each school reselects, again tentatively, up to its capacities from those held from the previous round and from new applicants, only based on its ranking of the students, and rejects the others. This process continues until no students are rejected, at which point the tentative assignment becomes final. Since schools select the students based *solely on the schools' own priorities*, top-ranking a very popular school under the DA does not sacrifice a student's chances at her less preferred schools in the event she fails to get into her top school.

Strategy-proofness is an important property to have, but that property alone would not be sufficient. For instance, a pure lottery assignment is also strategy-proof for a trivial reason but would not be considered desirable. The DA scores well on the welfare ground as well, *so long as schools have strict rankings over all students* (in addition to students having strict preferences over schools). In that case, the DA produces the so-called *student optimal stable matching*—a matching that is most preferred by every student among all stable matchings (Gale and Shapley 1962).[1] By contrast, the Boston mechanism may produce *any* stable matching in a Nash equilibrium *if all participants know all other participants' preferences as well as their priorities at all schools* (Haluk Ergin and Sönmez 2006).

In reality, however, schools do not have strict priorities over all students. For instance, BPS gives each student priorities based on whether she has a sibling enrolled at a school or whether she lives within the walk zone of a school. This leaves many students in the same priority class. In the DA, any tie among such students must be broken randomly. This makes the assumption of full information particularly problematic. Not only is it unlikely for students to know others' preferences, but it is simply impossible for them to know others' and even their own priorities at schools if they are determined randomly after students submit applications.

More important, coarse priorities alter the nature of welfare consideration itself. Families tend to value similar qualities about schools (e.g., safety, academic reputation, etc.), which causes them to have similar ordinal preferences. Indeed, the BPS data exhibits strong correlation across students' preferences over schools. In 2007–2008, only 8 out of 26 schools (at grade level 9) were overdemanded—that is, top-ranked by more participants than the seats available—whereas an average of 22.21 (s.d. 0.62) schools should have been overdemanded if students' preferences

[1] A matching is *stable* if no student or school can do strictly better by breaking off the current matching either unilaterally or by rematching with some other partner without making it worse off.

had been uncorrelated.[2] Correlated ordinal preferences entail conflicts among participants, and these conflicts cannot be resolved by the schools' priorities if they are coarse. Standard welfare concepts such as Pareto efficiency and student optimal stable matching then lose their relevance. For instance, if all students have the common ordinal preferences and schools have no priorities, then any arbitrary assignment meets these efficiency standards, and mechanisms become indistinguishable on these criteria. Yet this does not mean that all assignments or all mechanisms are equally desirable. Participants may still differ in their relative preference intensities over alternative schools, so it is sensible to resolve conflicts based on these intensities (henceforth called cardinal utilities). For instance, if a seat is desired by two students, it seems sensible to assign that seat to the one who would gain more from that seat relative to her next alternative.

The Boston mechanism and the DA differ in the way they resolve conflicts. The DA makes truthful ranking a dominant strategy and resolves any conflict purely by random lotteries regardless of the underlying cardinal utilities. Therefore, the outcome of the DA is completely insensitive to *cardinal* preferences (although it is sensitive to students' true ordinal preference rankings). In contrast, the Boston mechanism allows participants to influence how ties are broken, so it has the potential to resolve conflicts based on students' cardinal utilities. In fact, the feature of the Boston mechanism often vilified as engendering "gaming" or "strategizing" may be useful for efficient resolution of conflicting interests. These subtleties didn't go unnoticed by the parents. In the wake of the BPS school redesign, parents noted:

> ... if I understand the impact of Gale Shapley, and I've tried to study it and I've met with BPS staff... I understood that in fact the random number... [has] preference over your choices... (Recording from the BPS Public Hearing, 6-8-05).

> I'm troubled that you're considering a system that takes away the little power that parents have to prioritize... what you call this strategizing as if strategizing is a dirty word... (Recording from Public Hearing by the School Committee, 05-11-04).

We argue that the participants' cardinal welfare can be captured well by ex ante Pareto efficiency[3]—this is useful since the welfare evaluation need not involve interpersonal utility comparison—and that, from that perspective, the DA entails a clear welfare loss relative to the Boston mechanism, given common ordinal preferences and coarse priorities. To illustrate, suppose three students, $\{1, 2, 3\}$, must be assigned to three schools, $\{s_1, s_2, s_3\}$, each with one seat. Schools have no intrinsic priorities over students, and student i has a von-Neumann Morgenstern (henceforth, vNM) utility value of v_j^i when she is assigned to school j:

[2] This comparison is based on submitted preferences under the DA introduced in 2005. Since the DA is strategy-proof and BPS paid significant attention in communicating that feature of the DA to the public, we assume that those submitted preferences are a good approximation of the underlying true preferences. For the counter-factual, we generated 100 different preference profiles by drawing a school as first choice for each student uniformly randomly from the set of schools and computed the number of overdemanded schools given school capacities.

[3] An assignment is ex ante Pareto efficient if it is Pareto efficient prior to the realization of any random lotteries necessary to break ties; namely, it is impossible to reallocate probability shares of different schools in a Pareto improving fashion.

	v_j^1	v_j^2	v_j^3
$j = s_1$	0.8	0.8	0.6
$j = s_2$	0.2	0.2	0.4
$j = s_3$	0	0	0

Every feasible matching is stable due to schools' indifferences. More important, any such assignment is ex post Pareto efficient, hence *student optimal stable*, since students have the same ordinal preferences. Yet, their ex ante welfare depends crucially on how the students' conflicting interests are resolved.

To see this, first consider the DA mechanism with random tie breaking. All three students submit true (ordinal) preferences, and they are assigned to the schools with equal probabilities. Hence, they obtain expected utilities of $EU_1^{DA} = EU_2^{DA} = EU_3^{DA} = 1/3$.

This assignment is ex ante Pareto-dominated by the following assignment: *Assign student* 3 *to* s_2 *and students* 1 *and* 2 *randomly between* s_1 *and* s_3, which yields expected utilities of $EU_1^B = EU_2^B = EU_3^B = 0.4 > 1/3$. Surprisingly, this latter, Pareto-dominating assignment arises as the unique equilibrium of the Boston mechanism.[4] Students 1 and 2 have a dominant strategy of ranking the schools truthfully, and student 3 has a best response of (strategically) ranking s_2 as her first choice.

This example has assumed, for ease of illustration, that participants have complete information about their preferences, but as will be seen, the underlying insight holds much more generally. In our baseline model, we consider a general school choice setting in which participants have common ordinal preferences and schools have no priorities. These latter two assumptions are needed to generate a clear result for the Boston mechanism; it is difficult to analyze the strategic interaction of players in a fully general setting. However, these two assumptions also reflect the salient features of school choice—correlated preferences and coarse school priorities—and serve to isolate their effects in the most transparent form. Some real world problems in fact involve no priorities on the school side. The Supplementary round of the New York City high school match and the choice procedure of Seoul set to begin in 2010 are two such examples.

Other than these two features, we make no further assumptions. Importantly, we consider a realistic setting in which participants have incomplete information about others' preferences. We then focus on Bayesian Nash equilibria in symmetric strategies—those that specify the same (possibly mixed) action for students with the same von-Neumann Morgenstern (vNM) utilities. The symmetry restriction seems well justified, especially when no particular pattern of asymmetry is known a priori. Our results are summarized as follows:

[4] This does not contradict Ergin and Sönmez (2006) since they assume strict school priorities.

First, in keeping with the example, every participant is at least weakly better off in any symmetric equilibrium of the Boston mechanism than in the dominant strategy equilibrium of the DA. This result rests on the intuition that the Boston mechanism allows the participants to communicate their cardinal utilities and resolve their conflicting interests in a more efficient way than does the DA.

We next address an important concern that the Boston mechanism may harm those participants who are not strategically sophisticated. While strategically sophisticated players generally do better than naïve ones with the same vNM values (almost by definition), naïve players may also benefit from the presence of strategic players. The latter participants avoid ranking popular schools highly, and this raises the naïve participants' odds of getting into those schools. We show that naïve participants have a higher chance of attending a popular school under the Boston mechanism than under the DA, and some of them are better off under the Boston mechanism.

Finally, we study how the alternative mechanisms interact with a neighborhood priority policy. An important goal of school choice is to provide students in poor neighborhoods with opportunities to attend good schools. This goal is served best by guaranteeing equal access to all schools regardless of where a child lives. Since neighborhood priorities favor children living in proximate neighborhoods of schools, they may interfere with the equal access goal. The extent to which this interference occurs differs between the two mechanisms. In the DA, a student need not give up her neighborhood priority to be considered for other (good) schools, whereas the Boston mechanism forces the participants to give up their neighborhood priorities when ranking other schools highly. In other words, the inhibitive power of the neighborhood priority is diminished in the Boston mechanism, and this increases access to good schools for those who do not have priority at those schools.

One may take away several broad implications from the current paper. First, we offer a new welfare perspective on school choice—the importance of resolving conflicting interests based on participants' cardinal utilities. This perspective has been missing in the prior school choice debate because authors have largely focused on "ordinal" notions of welfare such as ex post Pareto efficiency and student optimal stable matching. However, we believe the current "cardinal welfare" perspective is very important in settings such as school choice where participants have similar ordinal preferences.

Second, from this new welfare perspective, there is a clear welfare loss associated with the DA relative to the Boston mechanism. Although the past literature has recognized the possibility that some students may be better off (while others are worse off) under the Boston mechanism relative to the DA in the Bayesian setting, it does not point to a clear welfare loss associated with the DA.[5] By contrast, the welfare loss identified here is systematic in nature and occurs in plausible school choice scenarios.

Third, this welfare loss can be seen as the "price" paid for achieving strategy-proofness. For instance, in the above example, the very feature of the Boston mechanism responsible for strategic behavior (i.e., student 3 lying about her preference) leads to efficient resolution of conflicts in that case. More formally, it is not possible

[5] Ergin and Sönmez (2006) provide an example with an ambiguous welfare comparison between the two mechanisms. However, truth telling is the unique equilibrium in their example, so they are silent on the issue we raise here.

for (symmetric) mechanisms to have both strategy-proofness and ex ante Pareto efficiency in general circumstances (Lin Zhou 1990).

Fourth, the trade-off between incentive and cardinal welfare (or ex ante Pareto efficiency) has a policy implication for the design of desirable school choice procedure. As is much emphasized in the prior literature, strategy-proofness is an important property. Somewhat less appreciated, however, is what we highlight: strategy-proofness has its own cost that seems important particularly in school choice.[6] This is not to argue that the DA should be rejected in favor of, say, the Boston mechanism (or "the clock should be turned back" in the case of BPS).[7] Such a conclusion is unwarranted, just as it would be unwarranted to reject the Boston mechanism on account of what we know so far. What is crucial however is that the school choice debate must be informed on both sides of the trade-off. More importantly, further work is needed to quantify the benefits and costs associated with strategy-proofness, particularly on the empirical and experimental fronts.

I. DA versus Boston in the Baseline Model

We first consider the Bayesian model in which each student (family) knows her own cardinal preferences about the schools but does not know about the others', except for the underlying probability distribution. Such a model is realistic, more so than the complete information model in which agents are assumed to know all other players' preferences.[8] We show that if the students share the same ordinal preferences but may differ in their preference intensities and the schools have no priorities, the Boston mechanism Pareto dominates the DA.[9]

There are $m \geq 2$ schools, $S = \{s_1, \dots, s_m\}$ with the index set $A := \{1, \dots, m\}$. School $s_a \in S$ has capacity q_a. There are $n \geq 2$ students, each of whom draws vNM utility values $\mathbf{v} = (v_1, \dots, v_m)$ about the schools from a finite set $\mathcal{V} = \{(v_1, \dots, v_m) \in [0,1]^m \,|\, v_1 > v_2 \dots > v_m\}$ with probability $f(\mathbf{v})$.[10] The students all have the same ordinal preferences preferring school s_a to school s_b if $a < b$. Importantly, though,

[6] Exceptions are Erdil and Ergin (2008) and Abdulkadiroğlu, Pathak, and Roth (2009), who focus on ex post inefficiencies and do not deal with the Boston mechanism.

[7] Incidentally, the clock did turn back in the case of Seattle Public Schools (SPS), which has recently switched from a version of the DA to a new assignment plan, in which open enrollment applications for attendance at area schools are processed in the order of choice as in the Boston mechanism.

See http://www.seattleschools.org/area/newassign/final_assign_plan_June17_Cleveland.pdf for a more detailed description, document retrieved on 12/07/09.

[8] For instance, it is conceivable to find a geographical area with several schools with a clear quality ranking, which coincides with parents' preference ordering of schools. Yet it is impossible to find a case in which parents have heterogeneous preference orderings and every parent knows every other parent's ordinal preference relation and every other parent's random number so that she knows schools' strict priority rankings of students.

[9] The result that the Boston mechanism may Pareto dominate the DA from an ex ante efficiency standpoint was first brought to the debate by Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda (2008). Subsequently, Antonio Miralles (2008) and Clayton Featherstone and Muriel Niederlee (2008) revisited the Boston mechanism. Miralles (2008) adopts the framework and efficiency notions of Abdulkadiroğlu, Che, and Yasuda (2008) and studies a variant of the Boston mechanism with roundwise tie breakers. Featherstone and Niederlee (2008) introduce incomplete information to a model with completely symmetric students and schools and show that the Boston mechanism induces truth telling in equilibrium; therefore it assigns more students to their first choices.

[10] The finiteness is assumed to simplify the existence of the Bayesian equilibrium of the Boston mechanism. The argument for the comparison works for any arbitrary distribution.

the students may differ in their relative preference intensities. For convenience, we assume that $\sum_{a \in A} q_a = n$.[11]

Gale-Shapley's Deferred Acceptance Algorithm.—It is a dominant strategy for each student to report truthfully, so we focus on such an equilibrium. Each student is then assigned to school s_a whenever she fails to win a seat at schools that are preferred to s_a but wins a seat at s_a. Since ties are broken in an ex ante symmetric way for students, the probability of that event is given by

$$(1) \qquad \hat{P}_a := \left(\frac{n - q_1}{n}\right)\left(\frac{n - q_1 - q_2}{n - q_1}\right)\cdots\left(\frac{q_a}{n - q_1 - \cdots - q_{a-1}}\right) = \frac{q_a}{n}.$$

The Boston Mechanism.—Let Π be the set of all rank-order lists of S, and $\Delta(\Pi)$ the set of probability distributions over Π. A Bayesian strategy is a mapping $\sigma : \mathcal{V} \to \Delta(\Pi)$. We focus on a symmetric strategy where every agent follows the same Bayesian strategy, meaning that they play the same mixed strategy for each realized $\mathbf{v} \in \mathcal{V}$. Such an equilibrium can be shown to exist by the standard Nash existence argument. Fix *any* such equilibrium $(\sigma^*, \ldots, \sigma^*)$.

AQ 2

For any mixed strategy $\sigma \in \{\sigma^*(\mathbf{v})\}_{\mathbf{v} \in \mathcal{V}}$ used in equilibrium, let $P_a(\sigma)$ be the probability that a student is assigned to school s_a if that student employs the strategy σ and all other students play the symmetric equilibrium strategy σ^*. For each $a \in A$, we must have

$$(2) \qquad \sum_{\mathbf{v} \in \mathcal{V}} n P_a\big(\sigma^*(\mathbf{v})\big) f(\mathbf{v}) = q_a.$$

To see this, note first that the LHS is the total expected number of students who are assigned to school s_a. Each of n students realizes preference $\mathbf{v}$ with probability $f(\mathbf{v})$ and plays $\sigma^*(\mathbf{v})$. She then gets assigned to school s_a with probability $P_a\big(\sigma^*(\mathbf{v})\big)$ under the Boston mechanism equilibrium. Summing over possible types gives the expected number of students assigned to school s_a. The RHS represents the total number of seats at school s_a that are assigned in equilibrium. Clearly, equation (2) must hold for each $a \in A$.

Fix any type $\tilde{\mathbf{v}} \in \mathcal{V}$ of student. Suppose that student $\tilde{\mathbf{v}}$ picks the following strategy: $\tilde{\sigma} := \sum_{\mathbf{v} \in \mathcal{V}} \sigma^*(\mathbf{v}) f(\mathbf{v})$. That is, $\tilde{\sigma}$ involves playing $\sigma^*(\mathbf{v})$ with probability $f(\mathbf{v})$, i.e., according to the probability distribution of types that play that strategy. Then, that student will be assigned to school $s_a \in S$ with probability

$$(3) \qquad P_a(\tilde{\sigma}) \equiv \sum_{\mathbf{v} \in \mathcal{V}} P_a\big(\sigma^*(\mathbf{v})\big) f(\mathbf{v}) = \frac{q_a}{n} = \hat{P}_a,$$

where the first equality follows from (2) and the second follows from (1).

Since $\tilde{\sigma}$ need not be an equilibrium strategy, we must have

$$\sum_{a \in A} \tilde{v}_a P_a\big(\sigma^*(\tilde{\mathbf{v}})\big) \geq \sum_{a \in A} \tilde{v}_a P_a(\tilde{\sigma}) = \sum_{a \in A} \tilde{v}_a \hat{P}_a.$$

[11] Every student is entitled to public education by law so that $\sum_{a \in A} q_a \geq n$. The results hold for that more general case; see Abdulkadiroğlu, Che, and Yasuda (2009).

In other words, the following is true:

THEOREM 1: *In any symmetric equilibrium of the Boston mechanism, each type of student is weakly better off than she is under the DA with any symmetric tie breaking.*

II. Does the Boston Mechanism Harm Naïve Players?

The appeal of strategy-proof mechanisms such as the DA and others (e.g., the top trading cycles mechanism) is that participants' strategic sophistication becomes irrelevant since ranking schools according to their true preferences is their dominant strategy. By contrast, the Boston mechanism may expose strategically naïve participants, who always submit their true rankings. Indeed, Abdulkadiroğlu et al. (2006) provide potential evidence that some players may have behaved naïvely and suffered as a consequence under the Boston mechanism. They find that as many as 20 percent of the applicants ranked two overdemanded schools as their first and second choices. These applicants could never get admitted by their second choice schools, so they would have done better by using their second rank for some other school. Of course, their ex post suboptimal behavior does not mean that their behavior was necessarily suboptimal ex ante. Their behavior may as well have been optimal if they put sufficiently high chance, quite possibly rationally, to the event that these schools were not overdemanded. Nevertheless, the concern about the potential strategic exploitation of strategically naïve participants was an important consideration in the redesign of the BPS program.[12]

Consistent with this concern, Parag A. Pathak and Sönmez (2008) argue that strategically sophisticated participants exploit naïve ones in the Boston mechanism, to such an extent that *the sophisticated effectively enjoy a higher priority over the naïve at every school except for the latter's most preferred.* While naïve players are generally expected to do worse, the particular sense and extent to which they are exploited is striking. A closer look reveals, however, that this characterization rests crucially on the two modeling features: *strict school priorities* and *complete information by strategic players*. Given these assumptions, each strategic player knows exactly who her competitors are and what their priorities are at each school. So, if a strategic player stands no chance at getting admitted to her favorite school, but her competitor at the next best school is a naïve player and that school is the naïve player's second most preferred, say, then the former will exploit the latter by top-ranking that school under the Boston mechanism. Indeed, one can show that a naïve player is (at least weakly) worse off from the presence of a strategic player in such an environment.[13]

Clearly, the preceding argument relies crucially on complete information. Absent complete information, a strategic player cannot be sure whom she will face as competitors, so she cannot target naïve players for manipulation. Hence, a naïve player need not be the victim of the strategic behavior. On the contrary, a naïve player may

[12] BPS Superintendent Payzant noted: "A strategy-proof algorithm levels the playing field by diminishing the harm done to parents who do not strategize or do not strategize well." (Memorandum, May 25, 2005)

[13] This result holds under common ordinal preferences, complete information, and strict priorities by schools. See Proposition A of our online Appendix.

actually benefit from a strategic play. Given nonstrict school priorities, ties are broken randomly, so it is impossible for the strategic player to know the priorities of her competitors. Hence, a strategic player may end up forgoing a spot at her favorite school even though she would have gotten it had she ranked it truthfully. That spot will then go to another participant; a naïve player may as well be the beneficiary. In fact, there is a clear sense in which naïve players benefit from the presence of strategic behavior when schools have coarse priorities and participants have similar ordinal preferences. In that case, strategic players tend to avoid popular schools, and this increases the chance for naïve players to get admitted by their favorite schools (likely to be the popular schools given correlated ordinal preferences), which they will rank truthfully as first choice.

To illustrate, consider our example in the introduction, except now each school has a quota of two, and there are two students of each type, one naïve and one strategically sophisticated. That is, there are a total of six seats and six students. Just as before, under the DA, every student ranks truthfully, and each student has a one third chance of getting admitted to each school. In the Boston mechanism, naïve students (there are three, one for each type) and type 1 and 2 strategic students rank schools truthfully, namely $s_1 - s_2 - s_3$ in that order. Just as before, the strategic type 3 student submits the ranking of $s_2 - s_1 - s_3$. Consequently, the type 3 strategic student gets assigned to school s_2. All others, strategic and naïve, are assigned to the schools with probabilities $(P_{s_1}, P_{s_2}, P_{s_3}) = (0.4, 0.2, 0.4)$. Naïve students lose priority at school s_2 to the strategic type 3 student. Yet, they enjoy a higher probability of getting assigned to school s_1 due to that strategic player. In this example, the two naïve type 1 and 2 players are better off, whereas the naïve type 3 student is worse off. The next proposition formalizes this observation in our Bayesian model where each type $\mathbf{v} \in \mathcal{V}$ of student is naïve with probability $x \in (0,1)$.

THEOREM 2: *Consider a symmetric Bayesian equilibrium with naïve players.* (*i*) *If a strategic player manipulates with positive probability*, *then every naïve player is assigned to each of the top j schools*, $\{s_1, \ldots, s_j\}$, *for some* $j \in A$, *with weakly higher probability and to some school in that set with strictly higher probability under the Boston mechanism than under the DA.* (*ii*) *If strategic students with type* $\mathbf{v}$ *rank the schools truthfully in equilibrium*, *then naïve students with the same* $\mathbf{v}$ *are* (*at least weakly*) *better off under the Boston mechanism than under the DA.*

PROOF:

To prove (i), let j be the smallest index in A such that there exists some type of a strategic player who does not rank school s_j as jth. (Call the type "manipulating" type.) By definition, the manipulation involves ranking j lower than the jth position (i.e., ranking it lth for some $l > j$). Since each player, both strategic and naïve, ranks school $s_{j'}$ at the j'th position for $j' < j$, she is assigned to $s_{j'}$, for $j' < j$, with the same probability under the Boston mechanism as under the DA. When the manipulating type player is rejected by all schools $s_1, \ldots, s_{j'}$ (which occurs with positive probability), a naïve player will have a higher priority at s_j than such a player and the same priority as the other strategic player. Hence, a naïve player will have higher probability of assignment to school s_j under the Boston mechanism than under the DA. The second statement holds since a strategic player with $\mathbf{v}$ ranks the schools truthfully in equilibrium and is weakly better off from Boston than from the DA (by Theorem 1); the naïve students with the same $\mathbf{v}$ must be also weakly better off from the Boston.

Remark 1: One can also show as in Theorem 1 that all strategic participants are at least weakly better off under the Boston mechanism than under the DA. Pathak and Sönmez (2008) obtain a similar result with complete information under general preferences.

III. Neighborhood Priority and Access to Good Schools

Neighborhood priority is a common feature of many school choice programs. For instance, BPS gives priority to students who live within one mile of an elementary school, within 1.5 miles of a middle school, and within two miles of a high school in attending those schools. At the same time, one of the major goals of public school choice is to provide equal access to good schools for every student, especially for those in poor neighborhoods with failing schools. This goal is compromised by neighborhood priority.

The extent to which the neighborhood priority inhibits access to good schools for students in failing school districts differs between the two mechanisms. Under the DA, one does not need to give up her neighborhood priority when applying for other (better) schools. This is in sharp contrast to what happens under the Boston mechanism. When a student does not rank her neighborhood school as her first choice under the Boston, she loses her neighborhood priority at that school to those who rank it higher in their choice list. Similarly, if she ranks her neighborhood school as first choice, then she gives up priority at the other schools. In either case, another student would be able to improve her odds at that school or some other school. This feature of the Boston mechanism provides a greater access to good schools for students without neighborhood priority at those schools.

To formalize this insight, consider our general Bayesian model in which each of n students draws her vNM values from $\mathcal{V}$ according to probability distribution f. We further assume $n > q_1 + q_2$, that is, the two most popular schools have excess demand. Suppose that $n_a \geq 0$ students are given neighborhood priority at school $s_a \in \{s_1, \ldots, s_m\}$. We assume $n_a < q_a$ for each a, i.e., the number of students who receive the neighborhood priority is less than the quota at each school.[14] A symmetric Bayesian strategy in this case specifies the same (mixed) action for students with the same vNM value $\mathbf{v} \in \mathcal{V}$ and same priority standing. Then, the following characterizations hold.

THEOREM 3: *Consider any symmetric Bayesian equilibrium of the Boston mechanism. Every student with priority at s_a, $a \geq 3$ or no priority at any school has a strategy that guarantees a strictly higher probability of being assigned to $s_{\tilde{a}}$ for some $\tilde{a} = 1, 2$ compared to the DA.*

[14] A similar characterization holds without the assumption (see Abdulkadiroğlu, Che, and Yasuda 2009).

PROOF:

If every student ranks s_1 as first choice with probability 1 in equilibrium, then a student with priority at s_a, $a \geq 3$, or no priority at any school can guarantee assignment at s_2 by ranking it first. That probability is smaller than 1 under the DA since $n > q_1 + q_2$. If some type of student ranks s_1 below top with positive probability, then by ranking s_1 first, a student with priority at s_a, $a \geq 3$, or no priority at any school can get assigned to s_1 with a higher probability than under the DA. That follows since every student ranks s_1 first under the DA.

When school priorities are strict and students have the same ordinal preferences, the Nash equilibrium outcome of the Boston mechanism is unique, and it coincides with the unique stable matching of the economy, which in turn implies that there is no randomness or uncertainty in equilibrium. Strategic opportunities characterized in Theorem 3 arise under coarse school priorities and incomplete information. Such opportunities are not present under the DA since submitting true ordinal ranking is a dominant strategy regardless of the priority and information structure.

IV. Conclusion

The Boston mechanism is widely used in school choice programs in the United States and elsewhere.[15] Meanwhile, the matching literature on school choice seems to reject the Boston mechanism. The standard view is that the Boston mechanism has a serious deficiency in both incentives and welfare. Although the incentive problem of the Boston mechanism is well understood, its welfare assessment is not as clear cut. The existing welfare assessment has been shaped largely by models that make unrealistic assumptions such as complete strict priorities on the part of schools and complete information on the part of students. Such evaluation could serve as a reasonable, if not perfect, approximation of truth, either if schools have near-complete priorities over students or if students have divergent preferences. The real-life school choice environment seems far from either stylization, however. In practice, families tend to have similar preferences about schools, and schools have at best coarse priorities. In such an environment, the issue of how a mechanism resolves conflicts based on cardinal welfare—captured by ex ante Pareto efficiency—becomes important. We have shown that, from this perspective, the Boston mechanism possesses several desirable features that other alternatives such as the DA do not.

At the same time, our results should not be seen as an unqualified endorsement of the Boston mechanism. The lack of strategy-proofness remains a significant drawback of the Boston mechanism that may ultimately make it unacceptable. Nevertheless, the current paper has shown a clear trade-off in the choice between the DA and the Boston mechanism. Informing the school choice debate of this trade-off is the most important purpose of this paper. Resolving this trade-off ultimately

[15] In the United States, the Boston mechanism and its variants have been used in Seattle Public Schools, Washington; Cambridge, Massachusetts; Providence, Rhode Island; Fort Collins and Denver, Colorado; Charlotte-Mecklenburg, North Carolina; and Miami-Dade and Tampa-St. Petersburg, Florida. Examples from around the world include the assignment of city schools in Seoul set to begin in 2010, elementary and middle school admissions in Japan, elementary school assignments in Madrid and Barcelona, and college admissions in China and Germany.

necessitates quantifying both sides of the trade-off, which will require much more work on the theoretical, computational, empirical as well as experimental front. Also needed are attempts to explore mechanisms that balance the trade-offs better than the existing mechanisms. They remain topics for ongoing and future research.

REFERENCES

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2008. "Expanding 'Choice' in School Choice." Unpublished.

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2009. "Resolving Conflicting Preferences in School Choice: The 'Boston' Mechanism Reconsidered." Unpublished.

Abdulkadiroğlu, Atila, Parag Pathak, Alvin E. Roth, and Tayfun Sönmez. 2006. "Changing the Boston School Choice Mechanism." National Bureau of Economic Research Working Paper 11965.

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2009. "Strategy-Proofness Versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review*, 99(5): 1954–78.

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

Dubins, Lester E., and David A. Freedman. 1981. "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly*, 88(7): 485–94.

Erdil, Aytek, and Haluk Ergin. 2008. "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice." *American Economic Review*, 98(3): 669–89.

Ergin, Haluk, and Tayfun Sönmez. 2006. "Games of School Choice under the Boston Mechanism." *Journal of Public Economics*, 90(1–2): 215–37.

Featherstone, Clayton, and Muriel Niederle. 2008. "Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation." National Bureau of Economic Research Working Paper 14618.

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

Miralles, Antonio. 2008. "School Choice: The Case for the Boston Mechanism." Unpublished.

Pathak, Parag A., and Tayfun Sönmez. 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *American Economic Review*, 98(4): 1636–52.

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

Zhou, Lin. 1990. "On a Conjecture by Gale About One-Sided Matching Problems." *Journal of Economic Theory*, 52(1): 123–35.