Comparing school choice mechanisms by interim and ex-ante welfare

Peter Troyan

Stanford University, Department of Economics, 579 Serra Mall, Stanford, CA 94305, United States

A B S T R A C T

Recent work has highlighted welfare gains from the use of the Boston mechanism over deferred acceptance (DA) in school choice problems, in particular finding that when cardinal utility is taken into account, Boston interim Pareto dominates DA in certain incomplete information environments with no school priorities. We show that these previous interim results are not robust to the introduction of (weak) priorities. However, we partially restore the earlier results by showing that from an ex-ante utility perspective, the Boston mechanism Pareto dominates *any* strategyproof mechanism (including DA), even allowing for arbitrary priority structures. Thus, we suggest ex-ante Pareto dominance as a relevant criterion by which to compare school choice mechanisms. This criterion may be of particular interest to school districts, as they can be thought of as social planners whose goal is to maximize the overall ex-ante welfare of the students.

© 2012 Elsevier Inc. All rights reserved.

1. Introduction

Variants of the Gale–Shapley deferred acceptance (DA) algorithm (Gale and Shapley, 1962) and what has come to be known as the Boston mechanism (which we will often refer to simply as 'Boston') are widely used by school districts throughout the United States to assign K-12 students to schools. Understanding the advantages and disadvantages of these mechanisms is a matter of great practical importance that has been the focus of extensive research, both theoretical and empirical.[1] The two mechanisms have been widely studied, and choosing one over the other involves trade-offs between incentive and welfare properties. While earlier work promoted DA over Boston (e.g., Abdulkadiroğlu and Sönmez, 2003 and Ergin and Sönmez, 2006), several recent papers have re-examined the Boston mechanism in settings in which participants have limited information about the preferences of the other students and the (post-lottery) priority structures at the schools and have found advantages for the Boston mechanism. In particular, in incomplete information environments with common ordinal preferences for the students and no school priorities,[2] Abdulkadiroğlu et al. (2011) (henceforth, ACY) and

---

E-mail address: petetroyan@gmail.com.

[1] Abdulkadiroğlu and Sönmez (2003) began this line of research into school choice. See Abdulkadiroğlu et al. (2005a, 2005b, 2009) for applications of the theory to the redesign of school choice mechanisms in Boston and New York City.

[2] We will often call such priority structures "trivial" priority structures. A nontrivial priority structure is any priority structure that is not trivial.

Miralles (2008) show that the (symmetric) equilibrium outcomes of the Boston mechanism actually interim Pareto dominate[3] that of deferred acceptance.[4]

The contribution of the current work is two-fold. First, we study the robustness of the interim Pareto dominance results mentioned above to the introduction of more realistic, nontrivial priority structures. We give two examples with weak priorities in which some students are strictly (interim) better off under DA, and show in a general model that the same will be true for any priority structure that satisfies a mild condition that is satisfied by many real-world priority structures. These results are similar in flavor to many "impossibility theorems" in the matching literature, which construct preferences to show that certain properties of matching mechanisms will not hold in general (see Roth and Sotomayor, 1990, Chapter 4).

Second, because neither mechanism will interim Pareto dominate the other once we allow for priorities, we must search for alternative criteria by which to compare school choice mechanisms. We thus introduce a new criterion, ex-ante Pareto dominance, to the school choice literature. This criterion examines welfare before students know their own types (cardinal utilities and priorities). From this perspective, we show that we can once again rank mechanisms, and in particular, that Boston ex-ante Pareto dominates *any* strategyproof (and anonymous) mechanism (including DA and the top trading cycles algorithm), even allowing for arbitrary priority structures. Thus, there is an explicit welfare cost associated with the use of strategyproof mechanisms.

Our ex-ante Pareto dominance concept is similar to the ex-ante viewpoint often used in the mechanism design literature, and one of our main contributions is to show how it can be useful in school choice settings. Since it seems reasonable to assume that students do in fact know their own preferences and priorities at the time they submit their rankings, our ex-ante viewpoint is more usefully interpreted as a normative concept, rather than a positive one, and can be given several normative justifications. First, we can argue from behind a Rawlsian "veil of ignorance" and say that a student who was asked to choose between mechanisms before she knew her place in society (i.e., her preferences and priorities) would pick the ex-ante Pareto dominant mechanism. Additionally, from a policy perspective, we argue that this criterion is especially relevant for school districts, as they can be thought of as social planners who do not know the realized preferences of any individual and whose goal is to maximize overall ex-ante welfare. Finally, since policymakers must decide on mechanisms that will be used for several years, they may be more concerned with how the mechanism performs relative to the underlying distribution of preferences and priorities rather than optimizing performance for one specific realization.[5]

Thus, while introducing nontrivial priorities causes difficulties in ranking mechanisms from an interim perspective, our results provide some justification for the use of Boston mechanism from an ex-ante perspective, even when schools have priorities. This is not to say that DA (or strategyproof mechanisms in general) should be rejected in favor of the Boston mechanism; strategyproofness certainly has its own advantages. However, we are able to generalize the results of Abdulkadiroğlu et al. (2011) and Miralles (2008) to situations with arbitrary priority structures and point out clear welfare losses due to strategyproofness. Which of these issues is more important empirically, or whether there exist mechanisms that will perform better than both Boston and DA, is an important open question.

1.1. Related literature

This paper is related to the large number of works that have aided in the design of real-world institutions by examining the incentive and welfare properties of centralized matching mechanisms in general, and school choice mechanisms in particular. On the incentives side, Roth (1982) and Dubins and Freedman (1981) show that deferred acceptance is strategyproof, while Abdulkadiroğlu and Sönmez (2003) point out that the Boston mechanism requires students (or parents) to play a complicated strategic game and may harm naive students who fail to strategize. In fact, it is this feature that was important in the city of Boston's decision to abandon its namesake mechanism for a deferred acceptance procedure. On the efficiency side, Ergin (2002) and Ergin and Sönmez (2006) were the first to discuss possible ex-post Pareto inefficiencies of the two mechanisms in a school choice context. However, here we will be concerned with interim and ex-ante efficiency losses as a result of the tie-breaking necessary to construct schools' strict priority orderings over students, issues first raised by Abdulkadiroğlu et al. (2008), Erdil and Ergin (2008), and Abdulkadiroğlu et al. (2009).[6]

The papers most closely related to this one are Abdulkadiroğlu et al. (2011) and Miralles (2008), both of which investigate interim efficiency and show that Boston may actually interim Pareto dominate DA in situations with common ordinal preferences and no school priorities. The intuition for these results is that Boston allows students to indicate a relatively high cardinal utility for a school by promoting it above its true ordinal rank. However, the assumption of no school priorities

---

[3] By interim utility, we mean a situation in which students know their own types but only the distribution of the types of other students. Much of the previous work on this topic calls this "ex-ante" utility, but, in this paper we will also examine welfare from the perspective before students know even their own types, and we reserve the term "ex-ante" for this situation.

[4] See also Featherstone and Niederle (2008) who find gains to the Boston mechanism over DA in an experimental setting. Pais and Pintér (2008) is another experimental study that examines the top trading cycles (TTC) algorithm in addition to Boston and DA, finding that limited information may actually improve efficiency. Özek (2008) provides simple examples of problems in which the Boston mechanism may Pareto dominate DA.

[5] A similar argument is made by Featherstone and Niederle (2008).

[6] While most of the above works focus on specific aspects of the mechanisms, Kojima and Ünver (2010) take a general axiomatic approach to understanding the Boston mechanism, while Kojima and Manea (2007) do the same for deferred acceptance.

may not apply in many contexts. Many cities classify students into several priority levels at each school,[7] with a student in a higher priority level being admitted before a student in a lower priority level under Boston, if they rank the school the same. To highlight the role of priorities, we keep the common ordinal preferences assumptions found in the prior work for most of the analysis, but allow for arbitrary priority structures. As we show, when this is done, the interim welfare comparison between the two mechanisms is no longer clear cut. However, the ex-ante criterion we propose allows us to rank mechanisms in a wider range of scenarios, providing guidance in mechanism selection for school districts that may have complicated priority structures.

The remainder of the paper is organized as follows. In Section 2, we give two examples of school choice problems with nontrivial priority structures in which the Boston mechanism no longer interim Pareto dominates DA. Section 3 extends these examples to a general model, and identifies a sufficient condition on the priority structure under which Boston and DA cannot be interim Pareto ranked. Section 4 examines welfare from an ex-ante perspective, showing that from this viewpoint, Boston Pareto dominates any strategyproof and anonymous mechanism, even with priorities. Section 5 concludes. The proofs of the main propositions are relegated to Appendix A.

2. Two examples with weak priorities

We first briefly described DA and the Boston mechanism. Since the workings of both are well-known in the literature, we give only an informal description.[8]

2.1. Deferred acceptance

In round 1, each student applies to her most preferred school. Each school *tentatively* accepts the highest ranked (according to its priority ordering) students who apply to it, up to its capacity, rejecting all others. In round $t > 1$, any student who was rejected in round $t - 1$ applies to their most preferred school which has not yet rejected them. The schools consider all applicants held from round $t - 1$ and the new applicants in round t, and once again keep the highest ranked set. The algorithm finishes when every student either has a tentative acceptance or has applied to all acceptable schools.

2.2. Boston mechanism

In round 1 of the Boston mechanism, all students apply to their most preferred school on their list, and again, each school accepts the highest ranked set of students up to its capacity, rejecting all others. Unlike DA, however, these acceptances are not tentative, but permanent. All schools' capacities are then decreased by the number of students accepted. In round t, each student who has yet to be accepted applies to the t-th ranked school on her list, and each school accepts the best set of applicants up to its capacity. The algorithm finishes when every student is either accepted or has exhausted all schools on her list.

A crucial difference between the two mechanisms is that DA is strategyproof, while Boston is not. Under Boston, if a student ranks a school s second, she loses her priority to all those students who rank s first. Under DA, the acceptances in each round are tentative, and so a student who ranks a school s second is still able to apply to and receive this school in later rounds if she happens to be rejected from her first-choice school.

Since Boston is not strategyproof, it allows students to express a relatively high cardinal utility for a school by promoting it over its true ordinal rank. This feature was first identified by ACY and Miralles (2008), who show that with no school priorities, Boston may interim Pareto dominate DA. However, as the two examples below show, when schools have nontrivial priority structures, some students may be strictly better off under DA.

An important distinction is whether priorities are public or private information. If priorities are determined by such things as distance from a school, then they are in principle public. However, it can also be argued that parents may not know the exact priority level of every other student, and instead might have only an estimate of the number of students in a given priority level at a school (formally, we assume that they know the underlying distribution by which the priorities are distributed). Treating priorities as private information keeps the symmetry that is the main driving force for the interim Pareto dominance of Boston over DA, but, as we will show below, we will not be able to interim Pareto rank the mechanisms with nontrivial priorities, regardless of whether they are public or private. Example 1 treats priorities as public information, while Example 2 draws each student's priority level at a school i.i.d. across students and schools.

2.3. Example 1

Let there be three schools A, B, and C with 1 seat each and 3 students labeled 1, 2 and 3. The students all strictly prefer A to B to C, but may have different cardinal utilities for each school (discussed below). Each school has two priority levels

---

Table 1
Priority structure for Example 1.

	School *A*	School *B*	School *C*
Students with high priority	None	1	None
Students with low priority	1, 2, 3	2, 3	1, 2, 3

Table 2
Payoffs for each student from each action when of type $\mathbf{v}^i$.

	Payoff from *ABC*	Payoff from *BAC*
Student 1	$p_L^2(\frac{1}{3} + \frac{2}{3} v_B^i) + 2 p_L p_H(\frac{1}{2}) + p_H^2(1)$	v_B^i
Student 2 or 3	$p_L^2(\frac{1}{3} + \frac{1}{6} v_B^i) + 2 p_L p_H(\frac{1}{2}) + p_H^2(1)$	$p_L^2(v_B^i) + p_L p_H(\frac{1}{2} v_B^i) + p_L p_H(0) + p_H^2(\frac{1}{2})$

('high' and 'low'), with the priority structures given in Table 1. In calculating expected utilities, we assume that ties within a priority level are broken randomly.

The priorities are public information, while the students' cardinal utility values $\mathbf{v} = (v_A, v_B, v_C)$ are privately drawn from $\mathcal{V} = \{\mathbf{v}^L, \mathbf{v}^H\}$ with probabilities p_L for $\mathbf{v}^L$ and $p_H = 1 - p_L$ for $\mathbf{v}^H$. Let v_A^H be the cardinal utility to a type $\mathbf{v}^H$ student from receiving *A*, and likewise for v_A^L, etc. Normalize $v_A^i = 1$ and $v_C^i = 0$ ($i = L, H$), and let $v_B^H > v_B^L$.[9]

Student *i*'s strategy $\sigma_i : \mathcal{V} \to \Delta(\Pi)$ is a mapping from possible types to probability distributions over all possible ordinal rankings of schools. We will write, for example, *ABC* to denote the (pure) strategy of ranking *A* first, *B* second, and *C* last. Consider the following strategies under the Boston mechanism: $\sigma_i(\mathbf{v}^L) = ABC$ and $\sigma_i(\mathbf{v}^H) = BAC$ $\forall i = 1, 2, 3$. Intuitively, these strategies will be an equilibrium when p_L and v_B^H are high and v_B^L is low. This will generate high competition at *A* in round 1, and so students of type $\mathbf{v}^H$ will find it optimal to apply to *B* in round 1. But, this makes student 1 of type $\mathbf{v}^H$ worse off under the Boston mechanism because under DA she is guaranteed a school no worse than *B*, yet has some chance at her favorite school *A*.

Table 2 gives the expected payoffs to each player from each action when of type $\mathbf{v}^i$, fixing the strategies of the other players as above.[10] For example, fixing the strategies of the other players, if student 2 plays strategy *ABC*, with probability p_L^2 the other two players are of type $\mathbf{v}^L$ and therefore rank *A* first, giving her a 1/3 chance of receiving *A* and a 1/6 chance of receiving *B*. The rest of the table is filled in similarly. Also, note that students 2 and 3 are symmetric, and that, because of her high priority, student 1 is admitted to *B* for certain if she ranks *B* first.

If we use, for example, $p_L = 0.9$ and $p_H = 0.1$,[11] we find that *ABC* is the best response for students 2 or 3 of type $\mathbf{v}^i$ when $v_B^i < 0.51$ (and *BAC* is the best response when the reverse inequality holds). The analogous cutoff value for student 1 is 0.80. Thus, for $v_B^L < 0.51$, the best response of *any* student of type $\mathbf{v}^L$ is *ABC*; similarly, for $v_B^H > 0.80$, the best response of *any* student of type $\mathbf{v}^H$ is *BAC*. Therefore, the proposed strategies constitute an equilibrium.

Under DA everyone ranks truthfully, and so the expected utility of player 1 of either type is $\frac{1}{3} v_A^i + \frac{2}{3} v_B^i$. When $1 = v_A^H > v_B^H > 0.80$, this is clearly better than the Boston mechanism outcome we constructed above for player 1 of type $\mathbf{v}^H$, who receives v_B^H with probability 1, and so the symmetric equilibrium of the Boston mechanism no longer interim Pareto dominates deferred acceptance.[12]

2.4. Example 2

Symmetry is a large driving force of the interim Pareto dominance of Boston over DA, and Example 1 may be criticized as not a true counterargument when priorities are introduced because it is not symmetric: 1 is known to have a priority that 2 and 3 do not. As a robustness check, we show that even if we keep this symmetry by allowing players to independently draw priorities as private information, Boston still may not interim Pareto dominate DA.

Everything is the same as in Example 1, except that now each student is given high priority at school *B* with probability $q_B = 1/3$, i.i.d. across students (if they do not receive high priority, they receive low priority; all students once again have low priority at schools *A* and *C*). Let *Low* denote low priority and *High* high priority at *B*. With two possible utility draws and two priority draws, there are now four possible types: $(\mathbf{v}^H, High)$, $(\mathbf{v}^H, Low)$, $(\mathbf{v}^L, High)$, $(\mathbf{v}^L, Low)$. Consider the following symmetric strategy profile where all agents play strategy σ: $\sigma(\mathbf{v}^L, Low) = \sigma(\mathbf{v}^L, High) = ABC$ and $\sigma(\mathbf{v}^H, Low) = \sigma(\mathbf{v}^H, High) = BAC$.

---

[9] The notation *H* indicates that a student of type $\mathbf{v}^H$ has a relatively high utility for *B* (compared to those of type $\mathbf{v}^L$).

[10] It is dominated to rank school *C* in either of the top 2 spots. We consider equilibria in undominated strategies, and so a strategy here is essentially just a decision to rank either *A* or *B* first. We have substituted $v_A^i = 1$ and $v_C^i = 0$ in Table 2.

[11] There is a wide set of parameter values for which our result holds. The specific numbers are chosen purely for illustration.

[12] Readers who are concerned about the uniqueness of the Boston equilibrium are referred to Proposition 1 below, which proves in a more general model that there exist type spaces for which student 1 will be better off under DA than in *any* equilibrium of the Boston mechanism.

Table 3
Expected utility for a student of type ($\mathbf{v}^H$, *High*).

Boston	DA
$p_L^2(v_B^H) + 2p_L p_H[(1-q_B)v_B^H + q_B(\frac{1}{2}v_B^H + \frac{1}{2}v_C^H)]$ $+ p_H^2[q_B^2(\frac{1}{3}v_A^H + \frac{1}{3}v_B^H + \frac{1}{3}v_C^H) + (1-q_B)^2 v_B^H$ $+ 2q_B(1-q_B)(\frac{1}{4}v_A^H + \frac{1}{2}v_B^H + \frac{1}{4}v_C^H)]$	$\frac{1}{3}v_A^H + q_B^2(\frac{1}{3}v_B^H + \frac{1}{3}v_C^H) + (1-q_B)^2(\frac{2}{3}v_B^H)$ $+ 2q_B(1-q_B)(\frac{1}{2}v_B^H + \frac{1}{6}v_C^H)$

Using the same parameter values as in Example 1, it is possible to show that the above strategies constitute a symmetric Bayesian Nash equilibrium when $v_B^L < 0.53$ and $v_B^H > 0.72$.[13]

Table 3 computes the expected utility for a student of type $\mathbf{v}^H$ with high priority at B under both the Boston equilibrium described above and under DA. Algebra shows that for any $v_B^H < 0.80$, DA is preferred to the Boston mechanism for these types of students.

The intuition behind both examples is straightforward. If a student with high priority at the middle-ranked school B has a high enough cardinal utility for it, she is better off ranking B first under Boston and receiving it for sure. If she instead ranked A first and B is taken in round 1 by another student, she will be assigned C if she has a poor lottery draw; the marginal gain from A to B is not worth the risk of ending up with C. Under DA, she applies to A without giving up her priority at B. Since she has a chance at A but is guaranteed no worse than B, she is better off.[14]

DA also does not interim Pareto dominate Boston in either example due to the mechanism identified by ACY and Miralles (2008): students of type $\mathbf{v}^H$ without priority prefer Boston because they can strategize to gain a better chance at B in round 1. Thus, these two mechanisms will not be comparable on an interim welfare basis.

3. Interim welfare in a general model with weak priorities

The intuition of these simple examples can be extended to a large class of priority structures which includes many real-world examples. To formalize this, let $S = \{s_1, \ldots, s_m\}$ be a set of schools and $N = \{1, \ldots, n\}$ a set of students, with $m, n \geqslant 3$.[15] $q = (q_1, \ldots q_m)$ is a vector of school capacities. We assume that $\sum_{s=1}^m q_s = n$, so that all students can be accommodated in some school,[16] and that $q_s + q_t < n$ for all distinct s, t so that no two schools can accommodate all students. Additionally, all students are acceptable to all schools.

As in ACY, each student i privately draws a vector of vNM utility values $\mathbf{v} = (v_1, \ldots, v_m)$ from a finite set $\mathcal{V} \subset [0,1]^m$; we will often refer to this vector as an agent's "type". Preferences are strict, and we normalize the best and worst schools to have utilities 1 and 0 for each type; formally, for any $\mathbf{v} \in \mathcal{V}$ (i) $s \neq t \Longrightarrow v_s \neq v_t$, (ii) $\max_s v_s = 1$ and (iii) $\min_s v_s = 0$. Let $p_\mathbf{v}$ be the probability a student draws type $\mathbf{v}$, with $\sum_{\mathbf{v} \in \mathcal{V}} p_\mathbf{v} = 1$. Denote this prior, which is the same for all students and is common knowledge, as $\mathbf{p} = \{p_\mathbf{v}\}_{\mathbf{v} \in \mathcal{V}}$.

Let $\rho_s^i \in \mathbb{N}$ be the priority level of student i at school s. $\rho_s^i = 1$ means that i is in the best possible priority level at s, though each level can have many (possibly all) students. Let $\gamma_{sl} = \{j \in N: \rho_s^j = l\}$ be the set of students in the l-th priority level at school s, with cardinality $|\gamma_{sl}|$. As in Example 1, the priority structure is public information, and ties within a priority level are broken randomly when calculating expected utility.

Let l_s^* be the critical priority level at school s in the sense that any student with $\rho_s^i \leqslant l_s^*$ is guaranteed to be admitted to school s if they rank s first under the Boston mechanism, independent of the strategies of other students; these students also cannot be admitted to a worse school than s under DA. Formally, $l_s^* = \max\{l \in \mathbb{N}: \sum_{j=1}^l |\gamma_{sj}| \leqslant q_s\}$. It is possible that no such l_s^* exists, in which case we let $l_s^* = 0$.

A strategy σ for each player is a mapping $\sigma : \mathcal{V} \to \Delta(\Pi)$, where $\Delta(\Pi)$ denotes the set of probability distributions over Π, the set of ordinal rankings of schools. The solution concepts are dominant strategies for DA and Bayesian Nash equilibrium (in undominated strategies) for Boston.[17] We analyze welfare from an interim perspective, after students learn their own types, but before they learn the types of other students or the outcome of the lottery used to break ties.

The result is the following.

---

[13] The calculations are much more cumbersome, and so we do not print them here. They are available from the author. Once again, the exact choice of numbers is purely for illustration, and the idea holds more generally.

[14] This is actually complicated a bit in Example 2 with independent priority draws for all students, because a student is no longer **guaranteed** a seat at B, as other students may also have high priority there. However, since it is not "likely" that there will be other high priority students, the same general intuition works. This is also the reason there is an upper bound on the v_B^H's for which students of type $\mathbf{v}^H$ with high priority are better off under DA.

[15] When $m = 2$ it is a weakly dominant strategy for every student to rank truthfully under the Boston mechanism, and the outcome is trivially the same as under DA.

[16] Since all students are guaranteed public education, it is reasonable to assume $\sum_s q_s \geqslant n$. Restricting to equality is without loss of generality, and our results will hold in that case as well.

[17] Such an equilibrium exists by standard arguments.

Proposition 1. *Assume that there exist some student i and distinct schools s, t such that $\rho_s^i \leqslant l_s^*$ and $\rho_t^i > l_t^*$. Then, there exists a type space $\mathcal{V}$, together with prior **p**, such that some types of students are interim strictly better off under deferred acceptance than under any equilibrium outcome of the Boston mechanism.*

As mentioned in the introduction, this is similar in spirit to "impossibility results" common in the matching literature in that we are able to construct preferences for which the interim Pareto dominance of Boston over DA will not hold. The proof of the proposition is in Appendix A, but the intuition is similar to the examples above: we construct a type space such that i finds it optimal to rank the school she is guaranteed a seat at (school s) first under Boston, rather than risk trying for a better school (t) and missing, thereby ending up at a school worse than either. Under DA, she is guaranteed a school no worse than s, but, has some strictly positive chance of receiving t. Because under DA she does not have to give up her high priority at s when applying for t, she is better off.

The sufficient condition on the priority structure is likely to be satisfied in many real-world settings. It only requires that there is *some* student who is guaranteed admission to *some* school, but not to another. In situations where priorities are based on factors such as walk zones, this will likely be satisfied, as students will be guaranteed admission to their neighborhood school, but not to other schools outside of their neighborhood.[18]

While the proof uses a type space consisting of only three types, similar intuition will hold in larger markets where we expect many different types of students. Additionally, we only show explicitly that one student is better off under DA, but in general there can be many students for which this is true. Intuitively, any student whose highest priority is at a school in the middle of her ordinal rankings will be worse off under Boston if they perceive high competition for schools they like more and decide not to take the risk of applying to them. At the same time, other students will prefer Boston over DA because of the mechanism identified by Abdulkadiroğlu et al. (2011), and so we will not be able to rank one mechanism over the other based on interim Pareto dominance.

4. Ex-ante welfare

The results of the previous section necessitate a search for other criteria by which to compare mechanisms, and the criterion we propose is ex-ante Pareto dominance. Consider Example 2 above, in which each student was distributed a 'high' priority at school B independently with probability $q_B = 1/3$. In Example 2, we showed that conditional on having high priority at school B, a student may prefer DA. However, if we calculate the expected utility of any agent *before* she knows whether she has high priority or not, it is possible to show that all students will strictly prefer the Boston mechanism to DA. Since all students are symmetric from this point of view, we immediately see that the Boston mechanism once again (strictly) Pareto dominates DA, even with a nontrivial priority structure.

This insight will in fact hold more generally from the ex-ante perspective. To show this, we must expand our model from the previous section to incorporate a student's priority as part of his (privately observed) type.[19] Let each school s have L_s priority levels $\{1, \ldots, L_s\}$. Nature will randomly assign each student to exactly one priority level at each school, where the probability any student is assigned to level k at school s is α_{sk} (so $\sum_{k=1}^{L_s} \alpha_{sk} = 1$ for all s). The α_{sk}'s are common knowledge, and the priority draws are independent across both students and schools. For a student i, let $\rho^i = (\rho_1^i, \ldots, \rho_m^i)$ be a vector denoting his priority level at each of the m schools, and let $\mathcal{P}$ be the set of all such feasible vectors; formally, $\mathcal{P} = \times_{s=1}^m \{1, \ldots, L_s\}$. For a priority vector $\rho = (\rho_1, \ldots, \rho_m) \in \mathcal{P}$, define g_ρ as the probability any student is assigned this priority vector.[20]

Thus, a student's type $(\mathbf{v}^i, \rho^i) \in \mathcal{V} \times \mathcal{P}$ consists of a "utility type" $\mathbf{v}^i$ and a "priority type" ρ^i. We assume that the utilities and priorities are drawn independently, so that the probability that any student i is of overall type $(\mathbf{v}^i, \rho^i)$ is $p_{\mathbf{v}^i} g_{\rho^i}$. Note that it is not known a priori how many students will be in each priority level at each school; that is, the fact that agent i is in the first priority level at school s does not change the probability that some agent j also is in the first priority level at s. One could imagine an alternative formulation in which the size of each priority class at each school is fixed and known, and nature draws the priority structure uniformly across all those that satisfy these constraints. This formulation leads to equivalent results to those presented here; what is important is that the students are ex-ante symmetric. We choose the formulation here to avoid extra notational complications that provide no additional insight.

Because of the lotteries required to break ties in priorities, it will be simpler to model the mechanisms we consider as returning a *random assignment*. A *random assignment* is an $(n \times m)$-dimensional matrix A such that (i) $0 \leqslant A_{is} \leqslant 1$ for all $i \in N$ and $s \in S$, (ii) $\sum_{s=1}^m A_{is} = 1$ for all $i \in N$, and (iii) $\sum_{i=1}^n A_{is} = q_s$ for all $s \in S$. Let $\mathcal{A}$ be the set of random assignment matrices. The rows of A correspond to the (marginal) distribution with which a student is assigned to each school (and thus must sum to 1) and the columns similarly correspond to the schools.

Since each school can have multiple seats, the columns sum to the capacity of each school. In the case in which each school has one seat, the columns also all sum to 1, and the matrix is bistochastic. The Birkhoff–von Neumann theorem then states that every bistochastic matrix can be written as a (not necessarily unique) convex combination of permutation

---

[18] With no school priorities as in ACY, $l_s^* = 0$ for all s, and, since $\rho_s^i \geqslant 1$ for all i, s, the sufficient condition will not hold.

[19] All notation used in the previous section carries over here.

[20] That is, $g_\rho = \alpha_{1\rho_1} \alpha_{2\rho_2} \cdots \alpha_{m\rho_m}$. Clearly, $\sum_{\rho \in \mathcal{P}} g_\rho = 1$.

matrices. Since permutation matrices correspond to deterministic assignments, this means that any random assignment can be feasibly implemented as a lottery over deterministic assignments. When schools have more than 1 seat, the columns of A may sum to some integer more than 1. Budish et al. (2009) provide a generalization the Birkhoff–von Neumann theorem which shows that we are still able to implement any random assignment in our setting as a lottery over deterministic assignments, and thus we are justified in considering mechanisms that provide random assignment matrices as outputs.

Agent i of utility type $\mathbf{v}^i$ evaluates random assignments by simply calculating her expected utility. Letting A_i denote the i-th row of a random assignment matrix, define the preference relation $\succcurlyeq_{\mathbf{v}^i}$ such that $A \succcurlyeq_{\mathbf{v}^i} A'$ if and only if $\sum_{s=1}^{m} A_{is} v_s^i \geq \sum_{s=1}^{m} A_{is}' v_s^i$.[21]

A *mechanism* then is a function $\psi : (\Pi \times \mathcal{P})^n \to \mathcal{A}$, where, as before, Π is the set of all ordinal rankings over schools in S. We let $(\pi, \rho) = (\pi^1, \rho^1, \ldots, \pi^n, \rho^n)$ denote a vector of preference and priority inputs, one for each agent, and, as is standard, will often write $(\pi, \rho) = (\pi^i, \rho^i, \pi^{-i}, \rho^{-i})$ when we wish to separate the vector into the inputs for agent i and the remaining agents $-i$. We let $\psi_i(\pi, \rho)$ be i's random assignment when the agents submit (π, ρ); that is, $\psi_i(\pi, \rho)$ is the i-th row of the random assignment matrix $\psi(\pi, \rho)$. We will often use the dot product $\psi_i(\pi, \rho) \cdot \mathbf{v}^i = \sum_{s \in S} \psi_{is}(\pi, \rho) v_s^i$ to calculate the expected utility of agent i of utility type $\mathbf{v}^i$ under assignment $\psi(\pi, \rho)$. Examples of mechanisms include deferred acceptance, the Boston mechanism, or TTC with some arbitrary priority tie-breaking rule.

We will call a mechanism *anonymous* if the assignments to the agents depend only on their submitted preferences and priority vectors, and not on their labels. Formally, let $\beta : N \to N$ be any permutation of the agents. Then, ψ is anonymous if $\psi_{\beta(i)}(\pi^{\beta^{-1}(1)}, \rho^{\beta^{-1}(1)}, \ldots, \pi^{\beta^{-1}(n)}, \rho^{\beta^{-1}(n)}) = \psi_i(\pi^1, \rho^1, \ldots, \pi^n, \rho^n)$ for all submitted preferences and priority vectors; that is, permuting the indices of the agents simply permutes the rows of the assignment matrix. Note that this also implies that any two agents who submit the same preferences and have the same priority standing at all schools will receive the same random assignment. From a fairness perspective, this is a desirable feature of a mechanism. Most real-world school choice mechanisms use only the preferences submitted by the agent and their priorities, and so will satisfy anonymity. For example, when deferred acceptance (or Boston) is used in the presence of weak priorities for the schools, ties in priority are usually broken by assigning each student a unique random number. This procedure clearly satisfies anonymity as defined here.

Every mechanism ψ induces a game of incomplete information in which the students submit an ordinal preference list over schools and their priority vector, and, given these submissions, assignments are given according to ψ. For notational purposes, we define the action space of this game to be $\Pi \times \mathcal{P}$, and so a strategy for each agent is a mapping $\sigma : \mathcal{V} \times \mathcal{P} \to \Delta(\Pi \times \mathcal{P})$; however, while students can report any ordinal ranking $\pi \in \Pi$ they wish, their priority vector ρ is "hard information" in the sense that it cannot be misreported.

Let $o : \mathcal{V} \to \Pi$ be a function that assigns to any cardinal utility vector $\mathbf{v} \in \mathcal{V}$ its associated ordinal ranking of schools. We will say that a mechanism is *strategyproof* if truthfully reporting $o(\mathbf{v}^i)$ is a dominant strategy for every agent and every type $(\mathbf{v}^i, \rho^i)$. That is, $\psi_i(o(\mathbf{v}^i), \rho^i, \pi^{-i}, \rho^{-i}) \succcurlyeq_{\mathbf{v}^i} \psi_i(\hat{\pi}^i, \rho^i, \pi^{-i}, \rho^{-i})$ for all $(\mathbf{v}^i, \rho^i) \in \mathcal{V} \times \mathcal{P}$, and $\hat{\pi}^i \in \Pi$, $(\pi^{-i}, \rho^{-i}) \in (\Pi \times \mathcal{P})^{n-1}$. For a strategyproof mechanism, we will consider the dominant strategy equilibrium in which every agent always truthfully reports her ordinal preferences, i.e. every agent plays the strategy σ^{SP} defined by $\sigma^{SP}(\mathbf{v}, \rho) = (o(\mathbf{v}), \rho)$ for all $(\mathbf{v}, \rho) \in \mathcal{V} \times \mathcal{P}$.

As the Boston mechanism is *not* strategyproof, we must allow for mixing to guarantee the existence of an equilibrium in the induced preference revelation game. We restrict attention to symmetric Bayesian Nash equilibria. Since our game is finite and the payoffs are symmetric (because the Boston mechanism with symmetric tie-breaking is anonymous), a symmetric equilibrium exists by standard arguments.

Here, we do impose the common ordinal preferences assumption of ACY, so that $\mathcal{V} = \{(v_1, \ldots, v_m) \in [0, 1]^m : 1 = v_1 > v_2 > \cdots > v_m = 0\}$; that is, $\mathcal{V}$ is some finite set of vectors where, for each vector, the first coordinate is higher than the second, which is higher than the third, and so forth. This assumption has been used in the literature as an approximation of situations of high conflict amongst students, common in many school choice problems, and allows us to highlight the role of priorities and ex-ante welfare while still keeping the analysis tractable.[22]

We analyze welfare by calculating students' expected utilities before they learn their own types, which we call "ex-ante welfare". The main result is the following (for a more formal statement, see Appendix A).

Proposition 2. *If agents have common ordinal (but different cardinal) preferences, then every student is ex-ante weakly better off under any symmetric equilibrium of the Boston mechanism with random tie-breaking than under any strategyproof and anonymous mechanism, even with arbitrary priority structures.*

The method of the proof is to assume all other players are playing their equilibrium strategy under Boston, and then for any player, identify a strategy that gives her at least as high an ex-ante expected utility as she would receive under the strategyproof rule ψ^{SP}. Since the equilibrium strategy must be at least as good as this strategy, all players must be weakly better off under the Boston mechanism. The strategy that we use has each student, when of priority ρ^i, play the "average" equilibrium strategy of agents in the population with the same priority type, where the average is taken over all utility

---

[21] Clearly, only the utility type matters for evaluating random assignments, which is why $\succcurlyeq_{\mathbf{v}^i}$ does not depend on ρ^i.

[22] See, for example, Abdulkadiroğlu et al. (2011), Miralles (2008), and Featherstone and Niederle (2008). Also, see footnote 23 for a discussion of non-common ordinal preferences.

types in $\mathcal{V}$. Since everyone is symmetric before the types are known, this strategy achieves the same level of expected utility as under ψ^{SP}.

Since both DA and top trading cycles with random tie-breaking are strategyproof and anonymous mechanisms, the following corollary is immediate.

Corollary 1. *Every student is ex-ante weakly better off under any symmetric equilibrium of the Boston mechanism than under either deferred acceptance or top trading cycles with random tie-breaking, even under arbitrary priority structures.*

While the propositions claim that Boston is weakly better than strategyproof mechanisms for all students, the generalization of Example 2 discussed above shows that the converse does not hold (i.e., in some cases everyone will strictly prefer Boston, so that DA does not weakly Pareto dominate Boston). Thus, the welfare criterion has content, and can be used to meaningfully rank mechanisms.

The intuition behind these results is that when agents are ex-ante symmetric, strategyproof mechanisms ignore agents' cardinal preferences. While in general strategyproof rules can favor agents with priorities at certain schools, this is only advantageous to students who *know* they have these priorities. Taking the viewpoint before priorities are realized, this advantage of strategyproof mechanisms to such agents is lost, and the Boston mechanism performs better because it allows agents to express a relatively high cardinal utility for a school by promoting it over its true ordinal rank.[23] It should also be noted that the result in Abdulkadiroğlu et al. (2011) which finds that the Boston mechanism interim Pareto dominates deferred acceptance can also be easily extended via a similar argument to the one presented here to conclude that Boston will interim Pareto dominate not just DA but also any strategyproof and anonymous mechanism when there are no school priorities.

This ex-ante notion is useful because we are able to allow for arbitrary priority structures and yet can still Pareto rank various mechanisms. It is particularly relevant to a school district that must decide between mechanisms and is interested in maximizing the overall ex-ante welfare of its students. Alternatively, we can make a "veil of ignorance" type argument for the use of the Boston mechanism, in that a student who did not know her place in society (in particular, her priorities) would choose the Boston mechanism over any strategyproof and anonymous alternative.

5. Conclusion

This paper studies the welfare consequences of various school choice mechanisms. Recent work has renewed interest in the Boston mechanism from a welfare perspective by showing that once we incorporate the lotteries necessary to break priority ties and cardinal utility values into the model, Boston may unambiguously outperform deferred acceptance on interim welfare grounds in situations with common ordinal preferences and no school priorities. However, these interim Pareto dominance results are not robust to the introduction of nontrivial (weak) priorities, and so we must search for other criteria to compare the mechanisms when schools may have more complicated priority structures. We find that from an ex-ante perspective, Boston Pareto dominates not just DA, but in fact any strategyproof and anonymous mechanism, even when schools have arbitrary priority structures. Thus, we suggest ex-ante Pareto dominance as a possible welfare criterion for evaluating school choice mechanisms. This criterion is relevant to school districts and to hypothetical students behind a "veil of ignorance", and gives some justification for the use of the Boston mechanism in school choice problems.

This is not to say that the Boston mechanism should replace DA or other strategyproof alternatives, as the incentive problem of the Boston mechanism remains a significant shortcoming. However, previous work and the current paper all highlight the ever present trade-off between welfare and incentive properties, especially when weak priorities lead to difficulties over the best way to break ties. We do not definitively endorse any mechanism here, but simply inform the debate between them, in particular by introducing the notion of ex-ante Pareto dominance into the school choice literature. More theoretical, experimental, and empirical work is needed to either decide between these mechanisms or to find new mechanisms that do a better job of balancing these trade-offs. These are important areas of future research for academics and parents alike.

Acknowledgments

---

[23] Note that even without common ordinal preferences, agents would still be ex-ante symmetric, and so one might expect that similar results would hold with non-common ordinal preferences from the ex-ante perspective. While theoretical results are difficult to obtain, Miralles (2008) builds a computational model and indeed finds that Boston outperforms DA ex-ante with non-common ordinal preferences and nontrivial priority structures.

Appendix A

A.1. Proof of Proposition 1

Fix i, s, t as in the statement of the proposition. We will construct a type space consisting of three vectors $\mathcal{V} = \{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$ for which student i chooses to apply to school s in equilibrium in round 1 of the Boston mechanism even when it is not his favorite school. Let $u_t = v_t = w_s = 1$ (so that t is the best school for types $\mathbf{u}$ and $\mathbf{v}$, and s is the best for types $\mathbf{w}$). Let $p_{\mathbf{u}} > 0$ be the probability that a student is of type $\mathbf{u}$, and similarly for $p_{\mathbf{v}} > 0$ and $p_{\mathbf{w}} > 0$, with $p_{\mathbf{u}} + p_{\mathbf{v}} + p_{\mathbf{w}} = 1$.

We will make the argument in several steps. First, define the set X as follows. If $\sum_{k=1}^{l_t^*} |\gamma_{tk}| = q_t$, then $X = \bigcup_{k=1}^{l_t^*} \gamma_{tk}$. If $\sum_{k=1}^{l_t^*} |\gamma_{tk}| < q_t$, then $X = \{\bigcup_{k=1}^{l_t^*} \gamma_{tk}\} \cup Z$, where Z is any set such that $Z \subseteq \gamma_{t(l_t^*+1)}$, $|Z| = q_t - \sum_{k=1}^{l_t^*} |\gamma_{tk}|$, and $i \notin Z$.

In words, the set X is a set of exactly q_t students other than i who have some positive probability of receiving school t when they rank it first, even if everyone else is also ranking t first.[24] The set is well defined because of the definition of l_t^*.

Step 1. *There exists a type $\mathbf{u}$ such that for all $j \in X$ any strategy that does not rank t first when of type $\mathbf{u}$ is dominated.*

Proof. Let $\mathbf{u}$ be any vector that satisfies $1 = u_t > 1/n - \varepsilon_u > (\max_{k \neq t} u_k) > \cdots > (\min_{k \neq t} u_k) = 0$ for some $\varepsilon_u \in (0, 1/n)$. Consider some $j \in X$ and some profile of strategies of the other players σ_{-j}. Ranking t first gives j a probability $\Pr_j(\sigma_{-j}) \geqslant 1/n$ of being admitted to t.[25] There are two cases.

Case (i): $\Pr_j(\sigma_{-j}) = 1$. If $\Pr_j(\sigma_{-j}) = 1$, then j is guaranteed a seat at t if he ranks t first. Since t is j's favorite school when he is of type $\mathbf{u}$, it is obvious that any strategy which does not rank t first is dominated.

Case (ii): $1/n \leqslant \Pr_j(\sigma_{-j}) < 1$. When j is of type $\mathbf{u}$, ranking t first gives an expected utility that is bounded below by $\Pr_j(\sigma_{-j})$. $\Pr_j(\sigma_{-j}) < 1$ implies that school t is oversubscribed in round 1, and hence will be unavailable at round 2. Thus, ranking some other school first means that j cannot receive t, which implies that his expected utility is bounded above by $1/n - \varepsilon_u$. Since $\Pr_j(\sigma_{-j}) \geqslant 1/n > 1/n - \varepsilon_u$, any strategy that does not rank t first is dominated by any strategy that does. $\quad\square$

Now, consider a modified priority structure at s where the set of students with priority l is defined as $Y_{sl} = (\gamma_{sl} \setminus X) \setminus \{i\}$. Define $\hat{l}_s$ analogously to l_s^*, that is $\hat{l}_s = \max\{l: \sum_{k=1}^l |Y_{sk}| \leqslant q_s\}$. Define the set X' exactly as we defined X above, only replacing all t's with s's, $|\gamma_{tk}|$'s with $|Y_{sk}|$'s, and l_t^*'s with $\hat{l}_s$'s.

In words, the set X' is a set of q_s students who have a nonzero probability of receiving s when they rank it first, given that all students in X and student i are not competing for s in round 1.[26]

Step 2. *Assume that i ranks t first when of type $\mathbf{v}$. Then, given step 1, there exists a $\mathbf{w}$ such that the set of best responses of all students $j \in X'$ consists only of strategies that rank s first when of type $\mathbf{w}$.*

Proof. Let $1 = w_s > \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{v}}}{n} - \varepsilon_w > (\max_{k \neq s,t} u_k) > \cdots > w_t = 0$ for some $\varepsilon_w \in (0, \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{v}}}{n})$. Consider some $j \in X'$ and some strategy profile σ_{-j} of the other players satisfying the stated assumptions. Ranking s first gives j a probability $\Pr_j(\sigma_{-j}) \geqslant \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{v}}}{n}$ of being admitted to s. Apply the same argument used in step 1. $\quad\square$

Step 3. *There exists a type $\mathbf{v}$ such that $1 = v_t > v_s > \cdots$, yet student i of type $\mathbf{v}$ ranks s first in any equilibrium.*

Proof. Pick some $\varepsilon_v \in (0, 1)$ and let $\alpha = \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}}{2} \varepsilon_v + [1 - \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}}{2}]$. Note that $\varepsilon_v < \alpha < 1$, and let $\mathbf{v}$ be any vector such that $1 = v_t > v_s > \alpha > \varepsilon_v > \max_{k \neq s,t} v_k > \cdots > \min_{k \neq s,t} v_k = 0$.

Assume there existed an equilibrium in which student i of type $\mathbf{v}$ ranked t first, rather than s.[27] By step 2, in this equilibrium, all students $j \in X'$ must rank s first when of type $\mathbf{w}$. Thus, in this equilibrium, there is a probability $p \geqslant \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}}{2}$ that i is not admitted to either s or t, and in this case, the most utility he can get is ε_v.[28] Thus, an upper bound on i's expected utility from ranking t first is $\frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}}{2} \varepsilon_v + [1 - \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}}{2}] v_t = \alpha$. If i ranks s first, she is guaranteed admission, giving expected utility $v_s > \alpha$. This means that ranking s first is a profitable deviation, which is a contradiction. $\quad\square$

---

[24] Note that $i \notin X$ because $\rho_t^i > l_t^*$ and we have specified $i \notin Z$.

[25] This is true because we assumed that every $j \in X$ has a positive probability of being admitted to t even when everyone else ranks t first. Since there are only n students, this probability cannot be less than $1/n$.

[26] Such a set exists because we have assumed that $n > q_s + q_t$.

[27] Ranking any school besides t or s first is dominated by ranking s first.

[28] For example, when all $j \in X$ are of type $\mathbf{u}$ and all $k \in X'$ are of type $\mathbf{w}$, schools s and t are filled in round 1, by steps 1 and 2. This happens with probability $p_{\mathbf{u}}^{q_t} p_{\mathbf{w}}^{q_s}$. The extra $1/2$ factor is necessary because if $\rho_t^i = l_{t+1}^*$, then i has some positive probability (less than or equal to $1/2$) of being admitted to t when he is competing against all other people in X for it.

To complete the argument, consider any $\mathcal{V} = \{\mathbf{u}, \mathbf{v}, \mathbf{w}\}$ satisfying steps 1–3. By steps 1–3, in any equilibrium of the Boston mechanism, student i ranks of type $\mathbf{v}$ ranks s first and receives utility v_s, where $v_s < v_t = 1$. Under deferred acceptance, all students rank truthfully, and student i is guaranteed a school no worse for her than s. However, she has some positive probability of receiving school t,[29] meaning her expected utility under DA is $p v_t + (1 - p) v_s$ for some $p > 0$. Thus, student i is strictly better off under DA.

A.2. Proof of Proposition 2

Before proving the proposition, we introduce some additional notation and formally state the definition of expected utility. As is standard in finite games of incomplete information, we can take the (pure) strategy space to be $\Pi \times \mathcal{P}$ and allow *each type* of agent to mix by choosing some probability distribution in $\Delta(\Pi \times \mathcal{P})$ (as in Section 4). Alternatively, we can take the (pure) strategy space to be Σ, the space of *deterministic mappings* $\sigma : \mathcal{V} \times \mathcal{P} \to \Pi \times \mathcal{P}$, and allow each agent to mix over mappings in Σ by choosing some probability distribution from $\Delta(\Sigma)$. Both formulations are equivalent, and in the proof below, we alternate between the two, depending on which is more convenient.

To shorten notation, let $\theta^i = (\mathbf{v}^i, \rho^i)$. Given a profile of pure strategies of the expanded game $s \in \Sigma^n$, let $s(\theta) \in (\Pi \times \mathcal{P})^n$ denote the profile of actions taken when the realized type is $\theta = (\theta^1, \ldots, \theta^n)$. Letting $s^j(\theta^j) \in \Pi \times \mathcal{P}$ denote the action taken by agent j when of type θ^j (and similarly, letting $s^{-j}(\theta^{-j})$ be the profile of actions of all agents other than j when the realized types are θ^{-j}), we can write the ex-ante expected utility to any agent i under a mechanism ψ as $\sum_{\theta \in (\mathcal{V} \times \mathcal{P})^n} \Pr(\theta) \psi_i(s^i(\theta^i), s^{-i}(\theta^{-i})) \cdot \mathbf{v}^i$.

Let ψ^{SP} denote any strategyproof and anonymous assignment rule, and let $\sigma^{SP} \in \Sigma$ be the truthful mapping defined in Section 4. Then, the ex-ante expected utility to any agent i under ψ^{SP} can be written as

$$EU^{SP} = \sum_{\theta \in (\mathcal{V} \times \mathcal{P})^n} \Pr(\theta) \psi_i^{SP}\big(\sigma^{SP}(\theta^i), \sigma^{SP}(\theta^{-i})\big) \cdot \mathbf{v}^i \tag{A.1}$$

Let ψ^{BM} be the Boston mechanism (with random tie-breaking of priorities). Let $z = (z^1, \ldots, z^n) \in \Delta(\Sigma)^n$ be a vector of mixed strategies, and, for $\sigma^i \in \Sigma$, write $\Pr_{z^i}(\sigma^i)$ for the probability that strategy z^i assigns to the mapping σ^i, and let $\Pr_z(\sigma) = \prod_{j=1}^n \Pr_{z^j}(\sigma^j)$ for the probability that the vector of (pure) strategies is $\sigma = (\sigma^1, \ldots, \sigma^n)$ under mixed strategy profile z. Then, we can write the ex-ante expected utility to agent i under a strategy profile z as

$$EU_i^{BM}(z) = \sum_{\sigma \in \Sigma^n} \sum_{\theta \in (\mathcal{V} \times \mathcal{P})^n} \Pr_z(\sigma) \Pr(\theta) \psi_i^{BM}\big(\sigma^i(\theta^i), \sigma^{-i}(\theta^{-i})\big) \cdot \mathbf{v}^i$$

With slight abuse of notation, let $z^* = (z^*, \ldots, z^*)$ be a symmetric equilibrium of the Boston mechanism where each agent plays strategy $z^* \in \Delta(\Sigma)$. By symmetry, we can drop the dependence of the expected utility on i, and simply write $EU^{BM}(z^*)$. With these definitions, Proposition 2 can now be formally written as:

Proposition 2. *$EU^{BM}(z^*) \geqslant EU^{SP}$ for any symmetric equilibrium of the Boston mechanism z^*.*

This will be proved via two claims.

Claim 1. *The ex-ante expected utility to any agent under ψ^{SP} is $EU^{SP} = \sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \frac{q_s}{n}$.*

Proof. Let $\bar{\pi} = (o(\mathbf{v}), \ldots, o(\mathbf{v})) \in \Pi^n$ denote a vector of submitted preferences where each agent submits the same common ordinal preference vector $o(\mathbf{v})$ (since we assume common ordinal preferences, $\mathbf{v}$ can be any element of $\mathcal{V}$). Because ψ^{SP} is strategyproof, we know the agents will submit this preference profile in equilibrium. This strategy and the distribution over priority types induce a distribution over actions $(\pi, \rho) \in (\Pi \times \mathcal{P})^n$ taken by the agents in equilibrium (recall that agents must report their priority type truthfully). Let $\Pr(\pi, \rho)$ denote the probability that the action profile is (π, ρ). Since we know the submitted preference vector will be $\bar{\pi}$ with probability 1, we can drop the dependence on π and write just $\Pr(\rho)$. So, the ex-ante expected utility for any agent i can be written as

$$\sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} \sum_{\rho \in \mathcal{P}^n} \Pr(\rho) \psi_i^{SP}(\bar{\pi}, \rho) \cdot \mathbf{v}^i = \sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \sum_{\rho \in \mathcal{P}^n} \Pr(\rho) \psi_{is}^{SP}(\bar{\pi}, \rho) \tag{A.2}$$

Focus on the last term $\sum_{\rho \in \mathcal{P}^n} \Pr(\rho) \psi_{is}^{SP}(\bar{\pi}, \rho)$. Letting $\beta : N \to N$ be any permutation of the agents, partition the set of priority vectors $\mathcal{P}^n$ into R subsets $(\mathcal{H}_1, \ldots, \mathcal{H}_R)$ such that $\hat{\rho} = (\hat{\rho}^1, \ldots, \hat{\rho}^n)$ and $\tilde{\rho} = (\tilde{\rho}^1, \ldots, \tilde{\rho}^n)$ belong to the same partition member if and only if there exists a permutation β such that $(\hat{\rho}^1, \ldots, \hat{\rho}^n) = (\tilde{\rho}^{\beta^{-1}(1)}, \ldots, \tilde{\rho}^{\beta^{-1}(n)})$. Note that by our assumption that priority vectors are independent across agents, if $\hat{\rho}$ and $\tilde{\rho}$ both belong to the same partition member, then $\Pr(\hat{\rho}) = \Pr(\tilde{\rho})$. This means that we can write

---

[29] For example, when all students are of type w and therefore rank t last.

$$\sum_{\rho \in \mathcal{P}^n} \Pr(\rho) \psi_{is}^{SP}(\bar{\pi}, \rho) = \sum_{r=1}^{R} \Pr(\mathcal{H}_r) \sum_{\rho \in \mathcal{H}_r} \Pr(\rho | \mathcal{H}_r) \psi_{is}^{SP}(\bar{\pi}, \rho) \tag{A.3}$$

Consider an arbitrary $\mathcal{H}_r$, and fix any $\hat{\rho} = (\hat{\rho}^1, \ldots, \hat{\rho}^n) \in \mathcal{H}_r$. For any other $\tilde{\rho} \in \mathcal{H}_r$, by anonymity of ψ^{SP}, we can write $\psi_{is}^{SP}(\bar{\pi}, \tilde{\rho}) = \psi_{\beta(i)s}^{SP}(\bar{\pi}, \hat{\rho})$ for the permutation β such that $(\hat{\rho}^1, \ldots, \hat{\rho}^n) = (\tilde{\rho}^{\beta^{-1}(1)}, \ldots, \tilde{\rho}^{\beta^{-1}(n)})$. Since all permutations are represented and are equally likely, we express every term in the second summation as a function of the fixed $\hat{\rho}$ by simply changing the index corresponding to the permutation, and so we can write

$$\sum_{\rho \in \mathcal{H}_r} \Pr(\rho | \mathcal{H}_r) \psi_{is}^{SP}(\bar{\pi}, \rho) = \frac{1}{n} \sum_{i=1}^{n} \psi_{is}^{SP}(\bar{\pi}, \hat{\rho}) = \frac{q_s}{n} \quad \forall \mathcal{H}_r \tag{A.4}$$

for the fixed $\hat{\rho}$. The last equality follows because we are now simply summing the s-th column of the assignment matrix $\psi^{SP}(\bar{\pi}, \hat{\rho})$, which by definition is q_s.[30] Now, since $\sum_{r=1}^{R} \Pr(\mathcal{H}_r) = 1$, Eq. (A.3) is just $\frac{q_s}{n}$. Plugging this into (A.2), we get that the ex-ante expected utility under any strategyproof and anonymous mechanism is simply $EU^{SP} = \sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \frac{q_s}{n}$. $\quad \square$

Claim 2. *Consider any symmetric equilibrium of the Boston mechanism. Then, assuming all other agents follow their equilibrium strategies, any agent has a strategy that gives ex-ante expected utility equal to EU^{SP}.*

Proof. Let $z^* \in \Delta(\Sigma)$ be any symmetric equilibrium of the game induced by ψ^{BM}. It will be convenient to represent this strategy in an equivalent manner as a function $\sigma^* : \mathcal{V} \times \mathcal{P} \to \Delta(\Pi \times \mathcal{P})$; that is, rather than randomizing over deterministic mappings in Σ, a strategy is a function from types into randomizations over actions in $\Pi \times \mathcal{P}$ (with the restriction that students cannot misreport their priorities). From the perspective of agent i, the equilibrium strategies again induce a probability distribution over action profiles of the other agents $(\pi^{-i}, \rho^{-i}) \in (\Pi \times \mathcal{P})^{n-1}$. Denote the probability that the agents $-i$ play action profile (π^{-i}, ρ^{-i}) as $\Pr(\pi^{-i}, \rho^{-i})$, and note that since types are independent, this distribution is independent of the type of agent i. Given this distribution, we can define a quantity $P_{s,\rho^i}(y)$ to be the probability that a student is admitted to school s when she randomizes according to $y \in \Delta(\Pi \times \mathcal{P})$, conditional on having priority vector ρ^i, and assuming all other agents follow their equilibrium strategies.

$$n \sum_{\rho^i \in \mathcal{P}} \sum_{\mathbf{v}^i \in \mathcal{V}} g_{\rho^i} p_{\mathbf{v}^i} P_{s,\rho^i}(\sigma^*(\mathbf{v}^i, \rho^i)) = q_s \quad \Longrightarrow \quad \sum_{\rho^i \in \mathcal{P}} \sum_{\mathbf{v}^i \in \mathcal{V}} g_{\rho^i} p_{\mathbf{v}^i} P_{s,\rho^i}(\sigma^*(\mathbf{v}^i, \rho^i)) = \frac{q_s}{n} \tag{A.5}$$

The double summation is just the ex-ante probability that any given student is assigned to school s in equilibrium, which, by symmetry, is the same across students. Multiplying this by n gives the total number of seats assigned at school s, which, in equilibrium, must be equal to q_s.

Consider some agent i who, rather than playing the prescribed equilibrium strategy σ^*, deviates to the strategy $\hat{\sigma}$ defined as follows: $\hat{\sigma}(\mathbf{v}, \rho') := \sum_{\tilde{\mathbf{v}} \in \mathcal{V}} \sigma^*(\tilde{\mathbf{v}}, \rho') p_{\tilde{\mathbf{v}}}$. That is, she plays the "average" strategy of agents in the population with priority vector $\rho' \in \mathcal{P}$, averaging over $\tilde{\mathbf{v}}$. The agent does this for all $\mathbf{v}$, so that effectively, her strategy depends only on her priority ρ'. Then, we can write

$$P_{s,\rho^i}(\hat{\sigma}) = \sum_{\tilde{\mathbf{v}} \in \mathcal{V}} P_{s,\rho^i}(\sigma^*(\tilde{\mathbf{v}}, \rho^i)) p_{\tilde{\mathbf{v}}} \tag{A.6}$$

Thus, the (ex-ante) expected utility to such a deviation is

$$\sum_{\mathbf{v}^i \in \mathcal{V}} \sum_{\rho^i \in \mathcal{P}} p_{\mathbf{v}^i} g_{\rho^i} \sum_{s \in S} P_{s,\rho^i}(\hat{\sigma}) v_s^i$$

Substituting from (A.6), this becomes

$$\sum_{\mathbf{v}^i \in \mathcal{V}} \sum_{\rho^i \in \mathcal{P}} p_{\mathbf{v}^i} g_{\rho^i} \sum_{s \in S} \sum_{\tilde{\mathbf{v}} \in \mathcal{V}} P_{s,\rho^i}(\sigma^*(\tilde{\mathbf{v}}, \rho^i)) p_{\tilde{\mathbf{v}}} v_s^i = \sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \sum_{\rho^i \in \mathcal{P}} \sum_{\tilde{\mathbf{v}} \in \mathcal{V}} p_{\tilde{\mathbf{v}}} g_{\rho^i} P_{s,\rho^i}(\sigma^*(\tilde{\mathbf{v}}, \rho^i))$$

where the equality is just a rearrangement of the summations. But now, the double sum can be substituted from (A.5) to yield

$$\sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \frac{q_s}{n}$$

which means that strategy $\hat{\sigma}$ gives the same ex-ante expected utility as EU^{SP}. $\quad \square$

---

[30] To put it more simply, conditional on $\mathcal{H}_r$, the matrix $\psi^{SP}(\bar{\pi}, \tilde{\rho})$ is the same as $\psi^{SP}(\bar{\pi}, \hat{\rho})$ only with the rows permuted in the same manner as the priority vectors. All priority vectors are equally likely, so ex-ante, all agents are equally likely to be "assigned" any row of the (fixed) matrix $\psi^{SP}(\bar{\pi}, \hat{\rho})$.

Since the $\hat{\sigma}$ constructed in Claim 2 is not necessarily an equilibrium strategy, the equilibrium strategy σ^* must give a weakly higher expected utility than that under ψ^{SP}, which completes the proof.

References

Abdulkadiroğlu, A., Che, Y.-k., Yasuda, Y., 2008. Expanding 'choice' in school choice. Unpublished manuscript.

Abdulkadiroğlu, A., Che, Y.-K., Yasuda, Y., 2011. Resolving conflicting preferences in school choice: The "Boston mechanism" reconsidered. Amer. Econ. Rev. 101 (1), 399–410.

Abdulkadiroğlu, A., Pathak, P., Roth, A.E., 2005a. The New York City high school match. In: Papers and Proceedings of the One Hundred Seventeenth Annual Meeting of the American Economic Association. Amer. Econ. Rev. 95, 364–367.

Abdulkadiroğlu, A., Pathak, P., Roth, A.E., 2009. Strategyproofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. Amer. Econ. Rev. 99, 1954–1978.

Abdulkadiroğlu, A., Pathak, P., Roth, A.E., Sönmez, T., 2005b. The Boston public school match. In: Papers and Proceedings of the One Hundred Seventeenth Annual Meeting of the American Economic Association. Amer. Econ. Rev. 95, 368–371.

Abdulkadiroğlu, A., Sönmez, T., 2003. School choice: A mechanism design approach. Amer. Econ. Rev. 93, 729–747.

Budish, E., Che, Y.-k., Kojima, F., Milgrom, P., 2009. Designing random allocation mechanisms: Theory and applications. Stanford University, unpublished mimeo.

Dubins, L., Freedman, D., 1981. Machiavelli and the Gale–Shapley algorithm. Amer. Math. Mon. 88, 485–494.

Erdil, A., Ergin, H., 2008. What's the matter with tie-breaking? Improving efficiency with school choice. Amer. Econ. Rev. 98, 669–689.

Ergin, H., 2002. Efficient resource allocation on the basis of priorities. Econometrica 70, 2489–2498.

Ergin, H., Sönmez, T., 2006. Games of school choice under the Boston mechanism. J. Public Econ. 90, 215–237.

Featherstone, C., Niederle, M., 2008. Ex-ante efficiency in school choice mechanisms: An experimental investigation. NBER Working Paper 14618.

Gale, D., Shapley, L.S., 1962. College admissions and the stability of marriage. Amer. Math. Mon. 69, 9–15.

Kojima, F., Manea, M., 2007. Axioms for deferred acceptance. Econometrica 78, 633–653.

Kojima, F., Ünver, U., 2010. The 'Boston' school choice mechanism. Mimeo.

Miralles, A., 2008. School choice: The case for the Boston mechanism. Boston University, unpublished mimeo.

Özek, U., 2008. Games of school choice under uncertainty. University of Florida, unpublished mimeo.

Pais, J., Pintér, A., 2008. School choice and information: An experimental study on matching mechanisms. Games Econ. Behav. 64, 303–328.

Roth, A.E., 1982. The economics of matching: Stability and incentives. Math. Oper. Res. 7, 617–628.

Roth, A., Sotomayor, M., 1990. Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis. Cambridge University Press.