

This work is distributed as a Discussion Paper by the

STANFORD INSTITUTE FOR ECONOMIC POLICY RESEARCH



SIEPR Discussion Paper No. 10-021

Comparing School Choice Mechanisms by Interim and Ex-Ante Welfare

by
Peter Troyan

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

The Stanford Institute for Economic Policy Research at Stanford University supports research bearing on economic and public policy issues. The SIEPR Discussion Paper Series reports on research and policy analysis conducted by researchers affiliated with the Institute. Working papers in this series reflect the views of the authors and not necessarily those of the Stanford Institute for Economic Policy Research or Stanford University.

Comparing School Choice Mechanisms by Interim and Ex-Ante Welfare

Peter Troyan*

March 24, 2011

Abstract

The Boston mechanism and deferred acceptance (DA) are two competing mechanisms widely used in school choice problems across the United States. Recent work has highlighted welfare gains from the use of the Boston mechanism, 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 nontrivial (weak) priorities. However, we partially restore the earlier results by showing that from an ex-ante utility perspective, the Boston mechanism once again Pareto dominates any strategyproof mechanism (including DA), even allowing for arbitrary priority structures. Thus, we suggest ex-ante Pareto dominance as a criterion by which to compare school choice mechanisms. This criterion may be of interest to school district leaders, as they can be thought of as social planners whose goal is to maximize the overall ex-ante welfare of the students. From a policy perspective, school districts may have justification for the use the Boston mechanism over a strategyproof alternative, even with nontrivial priority structures.

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

*Stanford University, Department of Economics, 579 Serra Mall, Stanford, CA, 94305. petetroyan@gmail.com. I am especially grateful to Fuhito Kojima for many helpful discussions, and to John Hatfield, Muriel Niederle, Brian Baisa and Dan Fragiadakis for comments.

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 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, Che, and Yasuda (2010) (henceforth, ACY) and Miralles (2008) show that the (symmetric) equilibrium outcomes of the Boston mechanism actually interim Pareto dominate³ that of deferred acceptance.⁴

The contribution of the current work is two-fold. First, we study the robustness of the previous interim Pareto dominance results to the assumption of no school priorities. 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 likely to be satisfied by many real-world priority structures. These results are similar in flavor to many “impossibility theorems” in the matching literature, which often construct preferences to show that certain properties of matching mechanisms will not hold in general (see, e.g., Roth and Sotomayor (1990), chapter 4).

Second, because the mechanisms will in general not be comparable on an interim welfare basis once we allow for priorities, we must search for alternative criteria by which to compare school choice mechanisms. The criterion that we propose, ex-ante Pareto dominance, examines welfare before students know their own types (cardinal utilities and priorities). From this perspective, we can once again rank mechanisms, and we show that Boston ex-ante Pareto dominates *any* strategyproof (and

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

²We will often call such priority structures “trivial” priority structures. A nontrivial priority structure is any priority structure that is not trivial.

³By 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.

⁴See 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.

anonymous) mechanism (including DA and TTC⁵), even allowing for arbitrary priority structures. Thus, there is an explicit welfare cost associated with the use of strategyproof mechanisms.

Additionally, we argue that this criterion is especially relevant for school district leaders, as they can be thought of as social planners whose goal is to maximize overall ex-ante welfare rather than the interim welfare of any one individual. We can also interpret the results from behind a Rawlsian “veil of ignorance,” in that a student who was asked to choose between mechanisms before she knew her place in society (in particular, her priorities) would pick Boston over a strategyproof mechanism. Thus, while introducing nontrivial priorities causes difficulties in ranking mechanisms from an interim perspective, our results provide a justification for the use of Boston mechanism from an ex-ante perspective, even when schools have priorities.

Since we are interested mainly in the role of priorities, we keep the common ordinal preferences assumption for most of the analysis for simplicity. This assumption is used in the literature, and may serve as a good approximation in some school choice problems where student preferences are highly correlated.⁶ Without this assumption, the model becomes much more difficult to analyze, and, in these situations, computational results may prove useful. Miralles (2008) provides a similar computational model suggesting that Boston unambiguously outperforms DA ex-ante, even with arbitrary ordinal preferences and nontrivial priorities. This, combined with the theoretical results presented here, suggests that ex-ante welfare may be a useful criterion for school districts to consider when evaluating mechanisms.

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 parents (or students) 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

⁵TTC is Gale’s top trading cycles algorithm, adapted to school choice problems by Abdulkadiroğlu and Sönmez (2003).

⁶See ACY, who impose common ordinal preferences, Featherstone and Niederle (2008), who study environments with common ordinal and completely uncorrelated preferences, and Kesten (2010) who proposes an intermediate model between common ordinal preferences and completely uncorrelated preferences.

raised by Abdulkadiroğlu, Che, and Yasuda (2008), Erdil and Ergin (2008), and Abdulkadiroğlu, Pathak, and Roth (2009).⁷

The papers most closely related to this one are Abdulkadiroğlu, Che, and Yasuda (2010) 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 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 may not apply in many contexts. Many cities classify students into several priority levels at each school,⁸ 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. As we show, when this is allowed, the interim welfare comparison between the two mechanisms is no longer clear cut, yet 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 briefly and informally describe how the Boston mechanism and DA work, and discuss a few salient features of the mechanisms. In section 3, we give two examples of school choice problems with nontrivial priority structures in which the Boston mechanism no longer interim Pareto dominates DA. Section 4 extends these examples to a general model, and identifies a sufficient condition on the priority structure under which Boston will no longer interim Pareto dominate DA. Section 5 examines welfare from an ex-ante perspective, showing that from this viewpoint, Boston Pareto dominates any strategyproof and anonymous mechanism, even with priorities. Section 6 concludes.

2 The Boston Mechanism and Deferred Acceptance

Since the workings of both the Boston mechanism and DA are well-known in the literature, we give only an informal description of how they assign students to schools.⁹ Students submit ordinal rankings over schools and schools have a priority ordering over students, which is usually given by law (hence schools are not strategic in our models). Both mechanisms require the schools to have

⁷While most of the above works focus on specific aspects of the mechanisms, Kojima and Ünver (2010) takes a general axiomatic approach to understanding the Boston mechanism, while Kojima and Manea (2007) does the same for deferred acceptance.

⁸For example, Boston, Minneapolis, New York City, Seattle, and San Francisco all divide students into priority classes based on various characteristics such as distance from a school, sibling attendance, etc.

⁹For a more complete description of these mechanisms in the context of school choice, see, for example, Abdulkadiroğlu and Sönmez (2003).

a strict priority ordering over students. For cases in which a school’s priority ordering is weak (i.e., the law gives many students the same priority) we use a random tie-breaking procedure to construct a strict ordering with which the mechanism is then run.

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 $t \geq 1$, any student who was rejected in round $t - 1$ applies to their next most preferred school who has not yet rejected them. The schools then 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.

Boston Mechanism: In round 1 of the Boston mechanism, all students apply to the most preferred school on their list, and again, each school accepts the highest ranked set (according to its priority ordering) of students who apply to it 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 applies to the t^{th} ranked school on their list, and each school accepts the best set of applicants up to the new capacity. The algorithm finishes when every student is either accepted or has exhausted all schools on the list.

A few important features of these mechanisms should be pointed out. First, DA is strategyproof, while Boston is not. Effectively, 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. While strategyproofness is an important feature, our results, as well as those of previous papers on this subject, show that in the context of school choice, strategyproofness has a cost in terms of welfare. The Boston mechanism allows students to express a high cardinal utility for a school (by promoting it over its true ordinal rank if they foresee lots of competition for it), and it is this logic that leads to the welfare gains associated with the Boston mechanism.

3 Two Examples with Weak Priorities

We retain the assumptions of common ordinal preferences and incomplete information about others’ cardinal utility values as in Abdulkadiroğlu, Che, and Yasuda (2010) and part of Miralles (2008), but relax the assumption of no school priorities and show that 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 (i.e., 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 are not able to Pareto rank the mechanisms in either the public or private priority case. 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.

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 (‘high’ and ‘low’), with the priority structures given in the table below. In calculating expected utilities, we assume that ties within a priority level are broken randomly.

	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

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 $v_B^H > v_B^L$.¹⁰

Student i ’s strategy $\sigma_i : \mathcal{V} \rightarrow \Delta(\Pi)$ 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:¹¹ $\sigma_i(\mathbf{v}^L) = ABC$ and $\sigma_i(\mathbf{v}^H) = BAC \forall i = 1, 2, 3$.

Before formally checking that this is an equilibrium of the Boston mechanism for some parameter values, we briefly discuss the intuition. If type \mathbf{v}^L ’s are prevalent enough in the population, the above

¹⁰The 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).

¹¹Note that it is dominated to rank school C in the top 2 spots. We consider equilibria in undominated strategies, and so a strategy here is essentially a choice to rank either A or B first.

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

Table 1: Calculating expected utilities

strategies will generate a lot of competition for A . Ranking A first will nevertheless be optimal for these students if \mathbf{v}^L places a relatively high utility on A (i.e., if v_B^L is close to 0). On the other hand, if \mathbf{v}^H places relatively high weight on B (v_B^H close to 1), students of this type will find it optimal to apply to B in round 1 to avoid the competition at A . 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 .

Now, to find parameter values for which the above strategies are an equilibrium, we fix the strategies of the other players and compute the expected utility for each student from each of her two possible actions (recall a strategy is essentially a decision to rank A or B first). The utilities for each student from each action when of type \mathbf{v}^i are given in Table 1.¹²

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. In this case, student 2 has a $1/3$ chance of receiving A but only a $1/6$ chance of receiving B because it is possible for her to receive B only if the lottery is such that player 1 receives A (since 1 has a higher priority at B and ranks it the same); this gives overall utility $\frac{1}{3} + \frac{1}{6}v_B^i$. 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 $p_H = 0.1$,¹³ 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 at 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 . But, this means the proposed strategies are an equilibrium of the Boston mechanism.

Under DA everyone ranks truthfully, and so the expected utility of player 1 of either type is $\frac{1}{3}v_A^i + \frac{2}{3}v_B^i$. For $1 = v_A^H > v_B^H > 0.80$, 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. Thus, the symmetric equilibrium of the Boston mechanism no longer interim Pareto dominates deferred acceptance.¹⁴

¹²We have substituted $v_A^i = 1$ and $v_C^i = 0$ in the table.

¹³The existence of cardinal utilities that support the equilibrium is not sensitive to the choice of p_k 's. These numbers are purely for illustration.

¹⁴Readers who are concerned about the uniqueness of the Boston equilibrium are referred to Proposition 1 below,

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 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 $q_B = 1/3$ (if they do not receive high priority, they receive low priority; all students once again have low priority at schools A and C). As with the utility draws, each student only observes the outcome of her own priority draw, and the draws are i.i.d. across students. Let *Low* denote low priority and *High* high priority at B . With two possible utility draws and 2 priority draws, there are now four possible types for each student: $(\mathbf{v}^H, High)$, (\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 $\sigma(\mathbf{v}^H, Low) = \sigma(\mathbf{v}^H, High) = BAC$.

Using the same parameter values as in Example 1 for illustration,¹⁵ it is possible to show that the above strategies constitute a symmetric Bayesian Nash equilibrium for any $v_B^L < 0.53$ and $v_B^H > 0.72$.

Now, all that remains is to show that some types of students are better off under DA. The following table computes the expected utility for a student of type \mathbf{v}^H with high priority at B under both mechanisms:

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

Given the specified parameters, algebra shows that for any $v_B^H < 0.80$, DA dominates the Boston mechanism for these types of students. So, to summarize: we have given an example of a priority structure and type distribution such that, for any $0 < v_B^L < 0.53$ and $0.72 < v_B^H < 0.80$, we have an equilibrium of the Boston mechanism under which some students are strictly worse off than under DA.

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 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.

¹⁵The calculations are much more cumbersome, and so we do not print them all here. They are available from the author. Once again, the exact choice of numbers is purely for concreteness, and the idea holds more generally.

under Boston and getting it for sure (because of her priority). This is true because 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 utility 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.¹⁶

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.

4 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. Intuitively, the condition on the priority structure we impose guarantees two things: first, that there exists some student i whose highest guaranteed school (i.e., a school at which she has high enough priority that she is certain to be admitted to it if she ranks it first under Boston) is in the middle of her ordinal rankings, and second, that i has a nonzero chance of being admitted to *some* more preferred school.¹⁷ If these conditions are satisfied, we can construct a type space for which the equilibrium outcomes of deferred acceptance and the Boston mechanism are incomparable by interim Pareto dominance.

To formalize this, let $S = \{s_1, \dots, s_m\}$ be a set of schools and $N = \{1, \dots, n\}$ a set of students, with $m, n \geq 3$.¹⁸ $q = (q_1, \dots, q_m)$ 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,¹⁹ 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.

Here only, we will not impose common ordinal preferences; this is not fundamentally important,

¹⁶This is actually complicated a bit in Example 2 with independent priority draws for all students, because a student is no longer **guaranteed** a spot at B , as other students may also have priority there. However, since it is not “likely” that there will be other 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 priority are better off under DA.

¹⁷This is probably satisfied in Boston, for example, where 50% of seats are set aside for walk-zone priority students. Boston also has sibling priorities, but in general, some seats at every school will likely be open to all students.

¹⁸When $m = 2$ 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.

¹⁹Since all students are guaranteed public education, it is reasonable to assume $\sum_s q_s \geq n$. Restricting to equality is without loss of generality, and the result presented here will hold in that case as well.

and we discuss what happens with this assumption after the formal proposition. Each student i privately draws a vector of vNM utility values $\mathbf{v} = (v_1, \dots, v_m)$ 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) $s \neq t \implies v_s \neq v_t$, (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 $\sum_{\mathbf{v} \in \mathcal{V}} p_{\mathbf{v}} = 1$. This prior is the same for all students, and is common knowledge.

Finally, we need some notation related to the priority structure. 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 $|\gamma_{sl}|$. The priority structure is public information, and ties within a priority level are broken randomly when calculating expected utility.

For any priority structure, let l_s^* be the critical priority level at s in the sense that any student with $\rho_s^i \leq l_s^*$ 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 under DA. Formally, $l_s^* = \max\{l \in \mathbb{N} : \sum_{j=1}^l |\gamma_{sj}| \leq q_s\}$. It is possible that no such l_s^* exists, in which case we let $l_s^* = 0$.²⁰

To achieve our result, we must make the following assumption on the priority structure: there exists a student i and distinct schools s, t such that $\rho_s^i \leq l_s^*$ and $\rho_t^i > l_t^*$. In words, all this means is that there is *some* student who is guaranteed admission to *some* school, but not to another. We will focus on such a student who prefers t to s .

A strategy σ for each player is $\sigma : \mathcal{V} \rightarrow \Delta(\Pi)$, 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 equilibria (in undominated strategies) for Boston.²¹ Here, 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. In the proof of the proposition, we construct a type space such that Boston does not Pareto dominate DA in any equilibrium.

Proposition 1. *There exist type spaces \mathcal{V} 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 the appendix, but

²⁰This is the case with no school priorities, as no student is guaranteed admission to any school.

²¹Such an equilibrium exists by standard arguments.

the intuition is similar to the examples above. When the condition on the priority structure holds, we can construct a type space such that i finds it optimal to rank the school she is guaranteed 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 , and thus is strictly better off. We use non-common ordinal preferences to ensure that student i has a *strictly* positive probability of receiving t under DA. Restricting to common ordinal preferences, we could only say that DA is *weakly* better than Boston for this student, because she would receive s under both mechanisms. Alternatively, we could strengthen the condition on the priority structure slightly to ensure that i has some positive chance at receiving t under DA.²² In summary, we are trying to highlight the role played by priorities, and the general idea holds regardless of the assumption on preferences: students with high priority at schools in the middle of their ordinal rankings for which they have relatively high cardinal utilities will prefer deferred acceptance because they do not have to give up this priority when applying for higher ranked schools.

Combining this with ideas of earlier works, we see that it will generally not be possible to argue for one mechanism over the other based on interim Pareto dominance. Also, note that the condition on the priority structure is somewhat weak, and is likely to be satisfied in a wide range of real life scenarios. In situations where priorities are based on factors such as walk zones, the condition will likely be satisfied, as students will be guaranteed admission to their neighborhood school, but not to other schools outside of their neighborhood.

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 a lot of competition for schools they like more and decide not to take the risk of applying to them.

²²Formally, this would mean that $\exists i$ such that $\rho_i^i = l_i^* + 1$ and $\sum_{j=1}^{l_i^*} |\gamma_{tj}| < q_t$. Many real world priority structures do satisfy this criterion, as schools often only have a few very broad priority classes. See, for example, <http://www.bostonpublicschools.org/files/IntroBPS09%20English.pdf> for a description of the priority structure used by the city of Boston.

5 Boston Ex-Ante Pareto Dominates Any Strategyproof Assignment Rule

As the previous discussion has shown, the interim Pareto dominance results will in general fail with the introduction of nontrivial priorities. This necessitates a search for other criteria by which to compare the two 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 take the equilibrium found there and calculate the expected utility of any agent *before* she knows whether she has high priority or not, it is possible to show that her expected utility will be (strictly) higher under the Boston mechanism than under DA. Since all students are symmetric from this point of view, we immediately see that the Boston mechanism once again (strictly) Pareto dominates DA from this perspective, even with a nontrivial priority structure.

Rather than showing this calculation explicitly for this example, we will move immediately to a general model. We can actually show more than the above and show that in situations with common ordinal preferences (but arbitrary priority structures), Boston will ex-ante (weakly) Pareto dominate *any* strategyproof mechanism that is anonymous in the sense that the assignment depends only on the submitted preferences and priorities, and not on the labels of the agents. This will include, for example, the DA algorithm with random priority tie-breaking and the top trading cycles (TTC) algorithm as defined for school choice problems by Abdulkadiroğlu and Sönmez (2003) (also modified to include random tie-breaking). This ex-ante notion is useful because we are able to allow for arbitrary priority structures and yet can still Pareto rank various mechanisms. The notion is also relevant to a school district superintendent who must decide between mechanisms and is interested in maximizing the overall ex-ante welfare of students in the school district. 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.

To formalize these statements, we must expand our earlier model. As in section 4, we let S and N be sets of schools and students and q be the school capacity vectors with $\sum_{s=1}^m q_s = n$. Π once again denotes the set of all possible (strict) ordinal rankings of the schools in S , and, as above, $\mathcal{V} \subset [0, 1]^m$ is a finite set of vNM utility vectors over the schools. We impose the restriction of common ordinal preferences, so that, without loss of generality, $1 = v_1 > \dots > v_m = 0$ for all $\mathbf{v} = (v_1, \dots, v_m) \in \mathcal{V}$. 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.²³ Without this assumption, the model is much more difficult to analyze, and we briefly discuss these situations in the next section.

Since the priorities of each student are no longer common knowledge, we will incorporate them as part of a student’s type. Formally, it is known a priori that each school s has L_s priority levels (let $\{1, \dots, L_s\}$ denote the set of priority levels at 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 $\rho^i = (\rho_1^i, \dots, \rho_m^i)$ 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, $\mathcal{P} = \times_{s=1}^m \{1, \dots, L_s\}$. For a priority vector $\rho^i = (\rho_1^i, \dots, \rho_m^i) \in \mathcal{P}$, define g_{ρ^i} as the probability any student is assigned this priority vector.²⁴

Thus, a type of a student is composed of a “utility type” \mathbf{v}^i and a “priority type” ρ^i , with the overall type being $(\mathbf{v}^i, \rho^i) \in \mathcal{V} \times \mathcal{P}$. We assume that the utility and priority types are drawn independently, so that the probability that any student i is of overall type (\mathbf{v}^i, ρ^i) is $p_{\mathbf{v}^i} g_{\rho^i}$.²⁵

We will in general denote $\pi^i \in \Pi$ as a (submitted) preference profile for agent i and $\pi = (\pi^1, \dots, \pi^n)$ as a collection of profiles for all agents. Similarly, Let $\rho = (\rho^1, \dots, \rho^n)$ be the collection of all students’ priority vectors. For notational simplicity, we will sometimes denote the type of student i as $\theta^i = (\mathbf{v}^i, \rho^i) \in \mathcal{V} \times \mathcal{P}$, and the type profile of all students other than i as θ^{-i} . Let $\theta = (\theta^1, \dots, \theta^n)$.

In all of the mechanisms we consider, students submit an ordinal preference ranking over schools, and based on these submitted rankings and priorities, the mechanism assigns the students to schools. The information structure of the game is such that each student will learn her own preferences and priorities, but will not know the preferences or priorities of any other student at the time her

²³See, for example, Abdulkadiroğlu, Che, and Yasuda (2010), Miralles (2008), and Featherstone and Niederle (2008), all of whom impose this assumption for some of their results. Kesten (2010) considers a hybrid model in which schools are partitioned into “quality classes” such that all students rank schools from different partitions the same (and this is common knowledge), but within a partition, any student’s ordinal preferences are drawn uniformly at random.

²⁴That is, $g_{\rho^i} = \alpha_{1\rho_1^i} \alpha_{2\rho_2^i} \dots \alpha_{m\rho_m^i}$. Clearly, $\sum_{\rho^i \in \mathcal{P}} g_{\rho^i} = 1$.

²⁵Note that in this formulation, 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 . Rather than letting each student receive priority k at school s with some fixed probability, we could imagine a formulation in which the size of each priority class at each school is fixed and known a priori, 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 in the text to avoid extra notational complications.

preferences are submitted. This is similar to Example 2 in section 3 above, and captures a realistic situation in which students do not know the exact priorities of every other student, but instead may only have an estimate of how many students will be in each priority level at each school (i.e., they know the underlying distribution by which priorities are distributed).²⁶

Because of the lotteries required to break ties, it will be simpler to model the mechanisms we consider as returning a *random assignment* for any preference-priority vector input. A *random assignment* is an $n \times m$ dimensional matrix A such that (i) $0 \leq A_{is} \leq 1$ 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. A random assignment is a matrix where the rows 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 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, Che, Kojima, and Milgrom (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 expected utility. Letting A_i denote the i^{th} row of a random assignment matrix, define the preference relation $\succeq_{\mathbf{v}^i}$ such that $A \succeq_{\mathbf{v}^i} A'$ if and only if $\sum_{s=1}^m A_{is} v_s^i \geq \sum_{s=1}^m A'_{is} v_s^i$.²⁷

An *assignment rule* here is a function $\psi : (\Pi \times \mathcal{P})^n \rightarrow \mathcal{A}$ which, for every submitted preference profile and priority structure, gives a random assignment. We let $(\pi, \rho) = (\pi^1, \rho^1, \dots, \pi^n, \rho^n)$ 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

²⁶As in the previous footnote, allowing the students to publicly observe the priorities after they are drawn by Nature would lead to similar results at the expense of more cumbersome notation. What is important is that the students are ex-ante symmetric and so from this perspective do not know what priorities they will be assigned.

²⁷Clearly, only the utility type matters for evaluating random assignments, which is why $\succeq_{\mathbf{v}^i}$ does not depend on ρ^i .

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 assignment rules include deferred acceptance, the Boston mechanism, or TTC with some arbitrary priority tie-breaking rule.

Let $o : \mathcal{V} \rightarrow \Pi$ 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 an assignment rule is *strategyproof* if truthfully reporting $o(\mathbf{v}^i)$ is a dominant strategy for every agent and every type (\mathbf{v}^i, ρ^i) . That is, $\psi_i(o(\mathbf{v}^i), \rho^i, \pi^{-i}, \rho^{-i}) \succeq_{\mathbf{v}^i} \psi_i(\hat{\pi}^i, \rho^i, \pi^{-i}, \rho^{-i})$ for all $(\mathbf{v}^i, \rho^i) \in \mathcal{V} \times \mathcal{P}, \hat{\pi}^i \in \Pi, (\pi^{-i}, \rho^{-i}) \in (\Pi \times \mathcal{P})^{n-1}$.

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 $\beta : N \rightarrow N$ be any permutation of the agents. Then, ψ is anonymous if $\psi_{\beta(i)}(\pi^{\beta^{-1}(1)}, \rho^{\beta^{-1}(1)}, \dots, \pi^{\beta^{-1}(n)}, \rho^{\beta^{-1}(n)}) = \psi_i(\pi^1, \rho^1, \dots, \pi^n, \rho^n)$ 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 anonymity will be satisfied by most standard school choice mechanisms. 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 assignment rule ψ induces a game of incomplete information (as described above) 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 to be $\Pi \times \mathcal{P}$; however, while students can report any ordinal rankings they wish, any student's priority vector ρ^i is "hard information" in the sense that it cannot be misrepresented. In this incomplete information game, the type space for all agents is $\mathcal{V} \times \mathcal{P}$ with common prior $p_{\mathbf{v}g\rho}$ as described above. The strategy space of the "expanded" incomplete information game is Σ , the space of mappings $\sigma : \mathcal{V} \times \mathcal{P} \rightarrow \Pi \times \mathcal{P}$, with the restriction that the second component of $\sigma(\mathbf{v}^i, \rho^i)$ is equal to ρ^i for all $\sigma \in \Sigma$ (i.e., students cannot misreport their priorities). Since both $\mathcal{V} \times \mathcal{P}$ and $\Pi \times \mathcal{P}$ are finite, Σ is also finite.

For a strategyproof assignment rule, 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 $\sigma^{SP}(\tilde{\mathbf{v}}, \tilde{\rho}) = (o(\tilde{\mathbf{v}}), \tilde{\rho})$ for all $(\tilde{\mathbf{v}}, \tilde{\rho}) \in \mathcal{V} \times \mathcal{P}$.

As the Boston mechanism is *not* a strategyproof assignment rule, we must allow mixing over deterministic mappings in Σ to guarantee existence of an equilibrium in the preference revelation

game induced by the Boston mechanism. We will restrict attention to symmetric Bayesian Nash equilibria $x \in \Delta(\Sigma)$ where every agent chooses the same strategy x . Since our game is finite and the payoffs are symmetric (which follows because the Boston mechanism with symmetric tie-breaking is anonymous), a symmetric equilibrium exists by standard arguments.²⁸

We analyze welfare from the perspective before students learn their own types, which we call “ex-ante welfare”. The main result is the following (for a formal definition of the ex-ante expected utilities, see the appendix).

Proposition 2. *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 assignment rule, even under arbitrary priority structures.*

Since all agents are ex-ante symmetric, under a strategyproof assignment rule, the ex-ante probability that they are assigned to any school is simply $\frac{q_s}{n}$. Because they all submit the same ordinal preferences, we can write the expected utility as $\sum_{s \in S} \sum_{v \in \mathcal{V}} p_v v_s \frac{q_s}{n}$.²⁹ 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 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 TTC (as defined in Abdulkadiroğlu and Sönmez (2003)) with random tie-breaking are strategyproof and anonymous mechanisms, the following corollary is immediate.

Corollary 3. *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

²⁸Also note that since our type and action spaces are finite, this “ex-ante” formulation is equivalent to a situation in which each agent chooses her action at an interim stage after her type is realized, but before she knows the types of any other agents. However, we evaluate welfare from the ex-ante perspective, and so this formulation is more convenient.

²⁹This formula is formally derived in the appendix.

(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 in this environment where agents are ex-ante symmetric, strategyproof assignment rules essentially allocate agents purely at random, ignoring the 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 assignment rules 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. It should also be noted that the result in Abdulkadiroğlu, Che, and Yasuda (2010) 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 assignment rule when there are no school priorities.

5.1 Non-common ordinal preferences

We finish by briefly discussing the assumption of common ordinal preferences. This restriction was needed to get the clear theoretical results in the previous section. As noted in ACY, this assumption may be regarded as a good approximation in some real life school choice problems where preferences are highly correlated. Without this restriction, the model becomes much more difficult to analyze.³⁰ One possibility to gain tractability while allowing for non-common ordinal preferences is to use computational models. Symmetry is a large driving force for the Pareto dominance of the Boston mechanism both here and in earlier works, and it was exactly the lack of symmetry created by the introduction of priorities that led to the interim Pareto incomparability of section 3. Thus, it may be reasonable to expect that Boston will still ex-ante Pareto dominate strategyproof mechanisms in the model of section 5 even with arbitrary ordinal preferences, as all students would still be ex-ante symmetric. Miralles (2008) analyzes a computational model similar to the model here and finds evidence that Boston does indeed outperform deferred acceptance ex-ante in school choice problems

³⁰There are other stylized restrictions that can make such models more tractable. For example, Roth and Rothblum (1999) and Ehlers (2008) consider models of “symmetric” information where students have very little information about the ordinal preferences of others, which is at the opposite extreme of our assumption of common ordinal preferences. Kesten (2010) proposes a hybrid model where schools are partitioned into “quality classes”, but within each class, preferences are “symmetric” in the Roth and Rothblum sense. Featherstone and Niederle (2008) show that in symmetric environments where the preferences of other students are drawn uniformly from Π , truth-telling is an ordinal Bayesian Nash equilibrium, and Boston outperforms DA at the interim stage both in theory and in experiments.

with arbitrary ordinal preferences and priority structures. The theoretical results presented here and elsewhere, combined with such computational models, give some justification for the use of the Boston mechanism from an ex-ante welfare perspective, even in the most general situations.

6 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 any strategyproof and anonymous mechanism, even when schools may have arbitrary priority structures. Thus, we suggest ex-ante Pareto dominance as a welfare criterion to rank school choice mechanisms. This criterion is relevant to real-world school superintendents and to hypothetical students behind a “veil of ignorance,” and gives justification for the use of the Boston mechanism in school choice problems.

Previous work and the current paper all highlight the ever present tradeoff 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 tradeoffs. These are important areas of future research for academics and parents alike.

References

- ABDULKADIROĞLU, A., Y.-K. CHE, AND Y. YASUDA (2008): “Expanding ‘Choice’ in School Choice,” Unpublished manuscript.
- (2010): “Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered,” forthcoming, *American Economic Review*.

- ABDULKADIROĞLU, A., P. PATHAK, AND A. E. ROTH (2005): “The New York City High School Match,” *American Economic Review, Papers and Proceedings*, 95, 364–367.
- (2009): “Strategyproofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match,” *American Economic Review*, 99, 1954–1978.
- ABDULKADIROĞLU, A., P. PATHAK, A. E. ROTH, AND T. SÖNMEZ (2005): “The Boston Public School Match,” *American Economic Review, Papers and Proceedings*, 95, 368–371.
- ABDULKADIROĞLU, A., AND T. SÖNMEZ (2003): “School Choice: A Mechanism Design Approach,” *American Economic Review*, 93, 729–747.
- BUDISH, E., Y.-K. CHE, F. KOJIMA, AND P. MILGROM (2009): “Designing Random Allocation Mechanisms: Theory and Applications,” Stanford University, unpublished mimeo.
- DUBINS, L., AND D. FREEDMAN (1981): “Machiavelli and the Gale-Shapley Algorithm,” *American Mathematical Monthly*, 88, 485–494.
- EHLERS, L. (2008): “Truncation Strategies in Matching Markets,” *Mathematics of Operations Research*, 33, 327–335.
- ERDIL, A., AND H. ERGIN (2008): “What’s the Matter with Tie-Breaking? Improving Efficiency with School Choice,” *American Economic Review*, 98, 669–689.
- ERGIN, H. (2002): “Efficient Resource Allocation on the Basis of Priorities,” *Econometrica*, 70, 2489–2498.
- ERGIN, H., AND T. SÖNMEZ (2006): “Games of School Choice under the Boston Mechanism,” *Journal of Public Economics*, 90, 215–237.
- FEATHERSTONE, C., AND M. NIEDERLE (2008): “Ex-Ante Efficiency in School Choice Mechanisms: An Experimental Investigation,” *NBER Working Paper 14618*.
- GALE, D., AND L. S. SHAPLEY (1962): “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 69, 9–15.
- KESTEN, O. (2010): “School Choice with Consent,” *Quarterly Journal of Economics*, 125, 1297–1348.
- KOJIMA, F., AND M. MANEA (2007): “Axioms for Deferred Acceptance,” *Econometrica*, 78, 633–653.
- KOJIMA, F., AND U. ÜNVER (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., AND A. PINTÉR (2008): “School Choice and Information: An Experimental Study on Matching Mechanisms,” *Games and Economic Behavior*, 64, 303–328.
- ROTH, A., AND M. SOTOMAYOR (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.
- ROTH, A. E. (1982): “The Economics of Matching: Stability and Incentives,” *Mathematics of Operations Research*, 7, 617–628.
- ROTH, A. E., AND U. G. ROTHBLUM (1999): “Truncation Strategies in Matching Markets-In Search of Advice for Participants,” *Econometrica*, 67, 21–43.

A Proof of Proposition 1

Fix i, s, t as in the statement of the proposition. We will construct a type space consisting of two 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 $X = \cup_{k=1}^{l_t^*} \gamma_{tk}$. If $\sum_{k=1}^{l_t^*} |\gamma_{tk}| < q_t$, then $X = \{\cup_{k=1}^{l_t^*} \gamma_{tk}\} \cup Z$, 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.³¹ The set is well defined because of the definition of l_t^* .

Step 1. *There exists a type \mathbf{u} such that it is a dominant strategy for any $j \in X$ to rank t first when of type \mathbf{u} .*

³¹Note that $i \notin X$ because $\rho_t^i > l_t^*$ and we have specified $i \notin Z$.

Proof. Let \mathbf{u} be any vector that satisfies $1 = u_t > 1/n - \varepsilon_u > (\max_{k \neq t} u_k) > \dots > (\min_{k \neq t} u_k) = 0$ 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 $p_j(\sigma_{-j}) \geq 1/n$ of being admitted to t .³² There are two cases.

Case (i): $p_j(\sigma_{-j}) = 1$

If $p_j(\sigma_{-j}) = 1$, 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 obviously dominant to do so.

Case (ii): $1/n \leq p_j(\sigma_{-j}) < 1$

When j is of type \mathbf{u} , ranking t first gives an expected utility that is bounded below by $p_j(\sigma_{-j})$. $p_j(\sigma_{-j}) < 1$ 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 $p_j(\sigma_{-j}) \geq 1/n > 1/n - \varepsilon_u$, any strategy that does not rank t first is dominated by any strategy that does. □

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 $\hat{l}_s = \max\{l : \sum_{k=1}^l |Y_{sk}| \leq q_s\}$. 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 1.³³

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 best responses of all students $j \in X'$ consist only of strategies that rank s first when of type \mathbf{w} .*

Proof. Let $1 = w_s > \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{v}}}{n} - \varepsilon_w > (\max_{k \neq s, t} u_k) > \dots > w_t = 0$ 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 $p_j(\sigma_{-j}) \geq \frac{p_{\mathbf{u}}^{q_t} p_{\mathbf{v}}}{n}$ of being admitted to s . Apply the same argument used in step 1. □

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

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

³²This 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 $1/n$.

³³Such a set exists because we have assumed that $n > q_s + q_t$.

Assume there existed an equilibrium in which student i of type \mathbf{v} ranked t first, rather than s .³⁴ 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 $p \geq \frac{p_{\mathbf{u}}^{qt} p_{\mathbf{w}}^{qs}}{2}$ that i is not admitted to either s or t , and in this case, the most utility he can get is ε_v .³⁵ Thus, an upper bound on i 's expected utility from ranking t first is $\frac{p_{\mathbf{u}}^{qt} p_{\mathbf{w}}^{qs}}{2} \varepsilon_v + [1 - \frac{p_{\mathbf{u}}^{qt} p_{\mathbf{w}}^{qs}}{2}] v_t = \alpha$. 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. \square

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 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 ,³⁶ 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.

B Proof of Proposition 2

Before proving the proposition, we first formally state the definition of expected utility.

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 θ . Letting $s^j(\theta^j) \in \Pi \times \mathcal{P}$ denote the action taken by agent j when of type θ^j (and similarly, $s^{-j}(\theta^{-j})$ is 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 an assignment rule ψ as $\sum_{\theta \in (\mathcal{V} \times \mathcal{P})^n} \Pr(\theta) \psi_i(s^i(\theta^i), s^{-i}(\theta^{-i})) \cdot \mathbf{v}^i$ (recall the notation $\theta^i = (\mathbf{v}^i, \rho^i)$).

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

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

³⁴Ranking any school besides t or s first is dominated by ranking s first.

³⁵For 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 $p_{\mathbf{u}}^{qt} p_{\mathbf{w}}^{qs}$. 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.

³⁶For example, when all students are of type w and therefore rank t last. This is the only place in the argument where different ordinal preferences are necessary.

Let ψ^{BM} be the Boston mechanism assignment rule (with random tie-breaking of priorities). Let $z = (z_1, \dots, z_n) \in \Delta(\Sigma)^n$ 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 $\sigma = (\sigma^1, \dots, \sigma^n)$ under mixed strategy profile z . Then, we can write the ex-ante expected utility to agent i under a strategy profile z as

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

With slight abuse of notation, let $z^* = (z^*, \dots, z^*)$ 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. $EU^{BM}(z^*) \geq EU^{SP}$ 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 $\bar{\pi} = (o(\mathbf{v}), \dots, o(\mathbf{v})) \in \Pi^n$ 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

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

Focus on the last term $\sum_{\rho \in \mathcal{P}^n} \Pr(\rho) \psi_{is}^{SP}(\bar{\pi}, \rho)$. Letting $\beta : N \rightarrow N$ be any permutation of the agents, partition the set of priority vectors \mathcal{P}^n into R subsets $(\mathcal{H}_1, \dots, \mathcal{H}_R)$ such that $\hat{\rho} = (\hat{\rho}^1, \dots, \hat{\rho}^n)$ and $\tilde{\rho} = (\tilde{\rho}^1, \dots, \tilde{\rho}^n)$ belong to the same partition member if and only if there exists a permutation β such that $(\hat{\rho}^1, \dots, \hat{\rho}^n) = (\tilde{\rho}^{\beta^{-1}(1)}, \dots, \tilde{\rho}^{\beta^{-1}(n)})$. 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

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

Consider an arbitrary \mathcal{H}_r , and fix any $\hat{\rho} = (\hat{\rho}^1, \dots, \hat{\rho}^n) \in \mathcal{H}_r$. 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 $(\hat{\rho}^1, \dots, \hat{\rho}^n) = (\tilde{\rho}^{\beta^{-1}(1)}, \dots, \tilde{\rho}^{\beta^{-1}(n)})$. 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

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

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 .³⁷ Now, since $\sum_{r=1}^R \Pr(\mathcal{H}_r) = 1$, equation (3) is just $\frac{q_s}{n}$. Plugging this into (2), we get that the ex-ante expected utility under any strategyproof and anonymous assignment rule is simply $EU^{SP} = \sum_{s \in S} \sum_{\mathbf{v}^i \in \mathcal{V}} p_{\mathbf{v}^i} v_s^i \frac{q_s}{n}$.

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} .*

Let $z^* \in \Delta(\Sigma)$ be any symmetric equilibrium of the game induced by ψ^{BM} (i.e., all students play the same strategy z^*). It will be convenient to represent this strategy in an equivalent manner as a function $\sigma^* : \mathcal{V} \times \mathcal{P} \rightarrow \Delta(\Pi)$; that is, rather than randomizing over deterministic mappings in Σ , a strategy is a function from types into randomizations over actions in Π . 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 playing strategy $y \in \Delta(\Pi)$, conditional on having priority vector ρ^i , and assuming all other agents follow their equilibrium strategies.³⁸ In equilibrium, the following must hold:

³⁷To 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 $\psi^{SP}(\bar{\pi}, \hat{\rho})$.

³⁸More formally, for all $\pi^i \in \Pi$, let $\Pr_y(\pi^i)$ denote the probability that strategy y assigns to action π^i . Then, $P_{s, \rho^i}(y) = \sum_{\pi^i \in \Pi} \sum_{(\pi^{-i}, \rho^{-i}) \in (\Pi \times \mathcal{P})^{n-1}} \Pr_y(\pi^i) \Pr(\pi^{-i}, \rho^{-i}) \psi_{is}^{BM}(\pi^i, \pi^{-i}, \rho^i, \rho^{-i})$.

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

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 $\hat{\sigma} : \mathcal{V} \times \mathcal{P} \rightarrow \Delta(\Pi)$ 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

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

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

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

Substituting from (6), this becomes

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

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

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

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

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.