

Improving on Strategy-proof School Choice Mechanisms: An Experimental Investigation[☆]

Clayton R. Featherstone^{a,*}, Muriel Niederle^b

^a*University of Pennsylvania, The Wharton School, 3620 Locust Walk – SHDH 1352,
Philadelphia, PA 19104*

^b*Stanford University, Department of Economics and NBER*

Abstract

While much of the school choice literature advocates strategyproofness, recent research has aimed to improve efficiency using mechanisms that rely on non-truthtelling equilibria. We address two issues that arise from this approach. We first show that even in simple environments with ample feedback and repetition, agents fail to reach non-truthtelling equilibria. We offer another way forward: implementing truth-telling as an ordinal Bayes-Nash equilibrium rather than as a dominant strategy equilibrium. We show that this weaker solution concept can allow for more efficient mechanisms in theory and provide experimental evidence that this is also the case in practice. In fact, truth-telling rates are basically the same whether truthtelling is implemented as an ordinal Bayes Nash equilibrium or a dominant strategy equilibrium. This provides a proof-of-concept that ordinal Bayes-Nash design might provide a middle path, achieving efficiency gains over strategy-proof mechanisms without relying on real-life agents playing a non-truth-telling equilibrium.

JEL classification: ?; ?; ?

Keywords: School Choice Mechanisms, Market Design, Experimental Economics, Matching Theory.

1. Introduction

Former Boston Public Schools superintendent Thomas Payzant described the rationale for switching away from a manipulable choice mechanism quite succinctly: “A strategy-proof algorithm ‘levels the playing field’ by diminishing the

[☆]This paper has also circulated under the title “School Choice Mechanisms under Incomplete Information: An Experimental Investigation” and “Ex Ante Efficiency in School Choice Mechanisms: An Experimental Investigation”.

*Corresponding author.

harm done to parents who do not strategize or do not strategize well.”^{1,2} This quote highlights two main assumptions implicit in much of the school choice literature. The first is that a significant fraction of parents will strategize poorly if faced with a situation where truth-telling is suboptimal. The second is that if truth-telling is not a dominant strategy, then parents might fail to submit preferences truthfully. In this paper, we present an experiment that investigates both of these assumptions in the context of two common³ school choice mechanisms: the strategy-proof Deferred Acceptance mechanism (“DA”) and the non-strategy-proof, “immediate acceptance” mechanism once used in Boston (“Boston”).

The Boston mechanism is quite popular, presumably because it maximizes the number of students that receive their reported first choice.⁴ However, it also provides participants with incentives to manipulate their preference reports, and in fact, a large number of experiments have confirmed that students do indeed deviate from truth-telling (Chen and Sönmez 2006, Pais and Pintér 2008, Calsamiglia et al. 2010)⁵. These papers mainly focus on determining what rules-of-thumb participants use when choosing their strategically manipulated preference reports. While Boston Public Schools has replaced the Boston mechanism with strategy-proof DA, several recent papers (Abdulkadiroğlu et al. 2012, 2011, Miralles 2009) have demonstrated that non-truth-telling equilibria under the old Boston mechanism can yield significant efficiency gains. For example, Abdulkadiroğlu et al. (2011), show in a stylized environment that the Boston mechanism can sustain an equilibrium that dominates the DA equilibrium by allowing students to better express their cardinal preferences.⁶ This raises a new question that current experiments fail to address, namely whether students would play close enough to the equilibrium prediction in a Boston mechanism for these gains to be realized.⁷

¹A school choice mechanism maps students’ reported preferences over schools to lotteries over possible allocations of students to schools. A strategy-proof school choice mechanism admits truthful preference revelation as a dominant strategy. A non-strategy-proof mechanism is also called manipulable.

²For more about Boston Public Schools’ decision to adopt a deferred acceptance assignment system, see Abdulkadiroğlu et al. (2005b) and Abdulkadiroğlu et al. (2006).

³DA is currently used in New York City (Abdulkadiroğlu et al. 2005a, 2009) and Boston (Abdulkadiroğlu et al. 2005b, 2006); Boston, or a close variant, is or has been used in in Charlotte-Mecklenburg (North Carolina), Lee County (Florida), Minneapolis, and Seattle (Abdulkadiroğlu and Sönmez 2003). These are, of course, non-exhaustive lists.

⁴Mechanisms like Boston are also popular in two-sided markets, see e.g. Roth (1991). While the Boston mechanism maximizes the number of students that receive the school they ranked as their first choice, it does not necessarily maximize the number of students that receive their first or second choice. For an investigation on mechanisms that seek to directly optimize the distribution of ranks received by students, see Featherstone (2013).

⁵See Calsamiglia et al. (2011) and Chen and Sönmez (2011) for comments on the data analysis in Chen and Sönmez (2006).

⁶Troyan (2012) suggests an extension that deals with more complex priority structures.

⁷Abdulkadiroğlu et al. (2006) provide suggestive evidence that some participants in Boston may not have manipulated their preference list in an optimal way.

Existing experiments cannot address this question, as they generally did not give participants sufficient information to calculate the equilibrium.⁸

The first part of our experiment tries to close this gap by considering a very simple incomplete information environment that has a unique, non-truth-telling equilibrium under the Boston mechanism.⁹ As theory predicts, students truth-tell under DA. Under Boston, we confirm former results that students do not play truth-telling strategies. However, even in our simple environment with feedback and repetition, students fail to play the theoretical equilibrium and in general are far from playing a best response. As such, we confirm the worry of Boston Public Schools superintendent Thomas Payzant that equilibria that entail untruthful reporting may be difficult to successfully implement in actual school choice problems.

Payzant’s second assumption is that strategy-proofness may be necessary to “level the playing field”. However, dominant-strategy incentive compatibility, while appealing, imposes a strong constraint on design, and it has been shown empirically that this constraint can come with costs.¹⁰ If we could successfully implement truth-telling in a weaker way, then we would open the door to the possibility of leveling the playing field with non-strategy-proof designs that are customized to induce truth-telling in particular environments.

In the second part of the experiment, we try to show that this could work by considering an environment in which truth-telling is an ordinal Bayes-Nash equilibrium under the non-strategy-proof Boston (and of course, a dominant-strategy equilibrium under DA).¹¹ In this environment, we show that participants truth-tell when it is an equilibrium to do so, whether it is implemented as an ordinal Bayes-Nash equilibrium or a dominant strategy equilibrium. Indeed, truth-telling rates turn out to be high and not statistically different across the two mechanisms. Furthermore, we empirically observe that the Boston outcome stochastically dominates the DA outcome. This does not imply that Boston is a good

⁸In Pais and Pintér (2008), students can calculate the equilibrium under the full information treatment, but not under their treatments with less information. The paper has however not addressed to what extent students play equilibrium.

⁹In our experiments, students know their own preferences, but only the distribution from which other students’ preferences are drawn. They also only know the distribution of the randomization device that is used to choose an allocation from the lottery over allocations outputted by the school choice mechanism. In other words, they don’t know the lottery tie-breakers, but they know their distribution.

¹⁰Abdulkadiroğlu et al. (2009) show that gains from the improvements suggested by Erdil and Ergin (2008) can be significant, but come at the cost of strategy-proofness. See also Kesten and Ünver (2013).

¹¹An ordinal Bayes-Nash equilibrium is a Bayes-Nash equilibrium in which each player’s equilibrium strategy yields an outcome distribution that first-order stochastically dominates the outcome distributions of all other potential strategies (when all other players are playing equilibrium strategies). See Ehlers and Massó (2007).

school choice mechanism for all environments, but rather that it is a good mechanism for our particular environment. It does suggest, however, that in more complicated environments, we might be able to design better mechanisms that still “level the playing field” in Payzant’s sense without being formally strategy-proof (i.e. dominant-strategy incentive-compatible). In the penultimate section of this paper, we investigate this idea further by providing theoretical intuition for why Boston performs so well in our experiment and by considering what sorts of mechanisms might outperform DA in other environments.

The efficiency gains of Boston over DA that are realized in our experiment are quite different from those mentioned in the other papers we’ve discussed. Our gains come out of a truth-telling equilibrium and are realized from the interim perspective, where students know their own preferences, but only the distribution over others’. Our gains are also not dependent on assumptions about cardinal utility, as the Boston outcomes dominate the DA outcomes in the first-order stochastic sense. Truth-telling is also a best-response, regardless of the underlying von Neumann-Morgenstern utilities. This stands in contrast to research that considers environments where truth-telling isn’t a dominant strategy and where gains are achieved through non-truth-telling equilibria that depend on cardinal preferences, e.g. Abdulkadiroğlu et al. (2011).

Before moving on, we highlight two theoretical points made by this paper. First, we investigate both efficiency and equilibrium from the interim information set, where agents know their own preferences, but have only incomplete information about others’ preferences. While this view is more common in the two-sided matching literature (Roth and Rothblum 1999, Kojima and Pathak 2009), it is less standard in school choice.¹² The interim approach may, however, be more realistic when considering the behavior of students. Furthermore, it may be the right way to think about the objective of a school board that wants to maximize efficiency over several cohorts of students and not simply a specific year (and hence a specific set of realized preferences). Second, we focus on ordinal Bayes-Nash incentive compatibility, which we feel to be the natural solution concept for non-strategy-proof design. Although non-strategy-proof implementation requires some assumption about underlying preferences, ordinal Bayes-Nash equilibrium only requires assumptions about underlying ordinal preferences, which can be easily estimated from reports submitted to a strategy-proof mechanism such as DA. Bayes-Nash equilibrium, although an even weaker solution concept, requires knowledge of the underlying cardinal preferences of students, which are not as easily ascertained.

The rest of the paper is organized as follows. Section 2 introduces the mecha-

¹²Notable recent exceptions include Ergin and Sönmez (2006), Erdil and Ergin (2008), Abdulkadiroğlu et al. (2011), and Troyan (2012).

nisms and environments used in the experiment and then presents the experimental design. Sections 4 and 5 discuss the results of the experiment. In Section 6, we discuss the intuition of how and when Boston can implement truth-telling, and briefly extend our findings to less symmetric environments. Section 7 concludes.

2. The experimental setup

2.1. Two school choice mechanisms

Formally, a school choice mechanism is a mapping from students' reported preferences to a lottery over allocations of students to schools. This mapping induces a preference revelation game for the students. In the mechanisms that we consider, we generate the outcome lotteries by randomizing the inputs to a deterministic algorithm; specifically, each school draws a strict priority ordering over the students from some distribution. In school choice mechanisms, this random process is usually constructed by starting with a weak ordering based on student characteristics and breaking indifferences with a lottery.¹³ Note that we could either draw the lottery independently for each school (multiple lotteries) or just once, using the same lottery for all schools (single lottery).¹⁴ While the priorities of the schools are codified by law or administrative procedure, students are strategic players.¹⁵

We use two mechanisms in the experiment. The first is based on the Deferred Acceptance algorithm (Gale and Shapley 1962) and has been introduced in Boston and New York City schools (Abdulkadiroğlu et al. 2005a,b).

The Deferred Acceptance Algorithm (DA)

Step 1: Students apply to their first choice school. Schools reject the lowest-ranking students in excess of their capacity. All other offers are held *temporarily*.

Step t: If a student is rejected in Step $t - 1$, he applies to the next school on his rank-order list. If he has no more schools on his list, he applies nowhere. Schools consider both new offers and the offers held from previous rounds. They reject the lowest ranked students in excess of their capacity. All other offers are held *temporarily*.

¹³For example, in Boston, the weak ordering is based on whether a student is within walking distance or has a sibling at the school in question (Abdulkadiroğlu and Sönmez 2003).

¹⁴This distinction was first made in Abdulkadiroğlu et al. (2005a).

¹⁵A notable exception to this is school choice in New York City, where school principals can strategically report their preferences over students (Abdulkadiroğlu et al. 2005a).

STOP: The algorithm stops when no rejections are issued. Each school is matched to the students it is holding at the end.

Two well known results concerning the revelation game induced by DA are that truth-telling is a weakly dominant strategy (Dubins and Freedman 1981, Roth 1982a,b) and that the outcome is stable relative to reported preferences (Gale and Shapley 1962).¹⁶ The second mechanism we consider is based on a specific priority algorithm, first described as the “Boston Algorithm” by Abdulkadiroğlu and Sönmez (2003). This mechanism is commonly used in American public school choice (Abdulkadiroğlu et al. 2005b) and can be thought of as an *immediate acceptance* algorithm.

The Boston Algorithm

Step 1: Students apply to their first choice school. Schools reject the lowest-ranking students in excess of their capacity. All other offers are *immediately accepted* and become *permanent* matches. School capacities are adjusted accordingly.

Step t: If a student is rejected in Step $t - 1$, he applies to the next school on his rank-order list. If he has no more schools on his list, he applies nowhere. Schools reject the lowest ranked students in excess of their capacity. All other offers become *permanent* matches. School capacities are adjusted accordingly.

STOP: The algorithm stops when no rejections are issued.

Note that in contrast to DA, where applications are tentatively held and acceptance is deferred until the end, applications in the Boston algorithm are *immediately accepted or rejected* at each step. The chance that one might be able to lock in a school over a higher priority student by ranking that school earlier in the algorithm is intuitively why Boston, unlike DA, fails to be strategy-proof.

2.2. Two preference environments

2.2.1 The aligned preference environment

The first environment we consider is the **aligned preference environment**. In it, there are three schools – “Best”, “Second”, and “Third” – and five students who are either “Top” or “Average”. Best has two seats, while the other schools only have one, and three of the students are Top, while the other two are Average. Each school prefers a Top student to an Average one, and all students prefer Best

¹⁶Relative to strict preferences for both schools and students, an outcome is called *stable* if it is individually rational and there is no *blocking pair* (a student and a school who each prefer each other to their assigned match). See the Appendix for a formal definition.

All schools prioritize...	All students prefer...			
Top \succ Average	Best \succ Second \succ Third \succ Unmatched			
(Ties broken by lottery)	$u = 100$	67	25	0
3 students 2 students	2 seats	1 seat	1 seat	∞ seats

Figure 1: The aligned preference environment

to Second to Third. All students also have the same von Neumann-Morgenstern payoffs over schools: 100 for Best, 67 for Second, 25 for Third, and 0 for being unmatched. Each school has an independently drawn uniform lottery to order students of the same priority class. This environment is summarized in Figure 1.

The matching that results under DA in equilibrium is the unique stable matching: the two Top students with highest lottery draws at Best get Best, the other Top gets Second, and the Average with the highest lottery draw at Third gets Third, leaving the remaining Average unmatched.

Proposition 2.1. *In the aligned preference environment, a profile of strategies is a Bayes-Nash equilibrium of the revelation game induced by DA if and only if that profile has the following form:*

- Top students submit preferences that rank Best first and Second second.
- Average students submit preferences that declare Third acceptable.

Moreover, all equilibria assign the three Top students to Best and Second, one Average student to Third, and leave the other Average unmatched.

Note that truth-telling is a weakly dominant strategy under DA. Under Boston, the unique stable outcome is not achievable in equilibrium when lottery draws are not known;¹⁷ in fact, when participants are risk-neutral, the pure-strategy Bayes-Nash equilibrium outcome is unstable.

Proposition 2.2. *In the aligned preference environment, when students are risk-neutral, a profile of strategies is a pure-strategy Bayes-Nash equilibrium of the revelation game induced by Boston if and only if that profile has the following form:*

¹⁷Per Ergin and Sönmez (2006), if the lottery draws at each school are common knowledge, then the equilibrium outcome with Boston is the same as with DA. One set of supporting strategies is as follows: the two Top students with the best lottery numbers at Best rank it first, the third Top student ranks Second first, the Average student who has the better lottery number at Third ranks it first, and the last Average student submits arbitrary preferences (as he will not receive any school).

- *Top students submit preferences that rank Best first and Third second.*
- *Average students submit preferences that rank Second first.*

Moreover, all pure-strategy equilibria assign the three Top students to Best and Third, one Average student to Second, and leave the other Average unmatched.

The equilibria of Proposition 2.2 entails two types of manipulations. Top students need to misreport their second choice school, submitting Third instead of Second, that is, they must “skip the middle”. Average students need to misreport their first ranked school, submitting Second instead of Best, that is, they must “skip the top”. Note that, in equilibrium, relative to DA, the Average students are in expectation better off, while the Top students are, in expectation, worse off.

2.2.2 The uncorrelated preference environment

The second environment we consider is the **uncorrelated preference environment**. In it there are five students and four one-seat schools. Student preferences are drawn independently from the uniform distribution over all possible orderings that find all schools acceptable, and likewise, each school has a strict priority based on an independent draw from the uniform distribution over all orderings of the students. In the experiment the payoffs are such that a seat at the highest ranked school yields 110 points, the second highest 90, the third highest 67 and the fourth 25. Being unmatched yields 0 points.

In this environment, both mechanisms admit all agents truth-telling as a Nash equilibrium, but while truth-telling is a weakly dominant strategy under DA, it is only an ordinal Bayes-Nash equilibrium strategy under Boston.¹⁸ Under DA, truth-telling is the unique equilibrium in weakly undominated strategies, and under Boston, truth-telling is the unique equilibrium in *weakly undominated, anonymous strategies*. A strategy is anonymous if the reported rank of a school only depends on its true rank (and not its name).¹⁹ Now we state the formal result that all students truth-telling is an ordinal Bayes-Nash equilibrium in both environments, and that this truth-telling is, in some sense, the unique equilibrium in these environments.

Proposition 2.3. *In the uncorrelated environment, the revelation game induced by DA admits truth-telling by all students as an ordinal Bayes-Nash equilibrium. Furthermore, this is the unique Bayes-Nash equilibrium.*

Proposition 2.4. *In the uncorrelated environment, the revelation game induced by Boston admits truth-telling by all students as an ordinal Bayes-Nash equilibrium. Furthermore, this is the unique Bayes-Nash equilibrium in anonymous strategies.*

¹⁸For a definition of an an ordinal Bayes-Nash equilibrium, see Footnote 11.

¹⁹See Definition B.2 for a more formal definition.

Note that these propositions are not sensitive to cardinal utility values that we chose for the uncorrelated environment.

Figure 2a shows bootstrapped outcome distributions under the Boston and DA truth-telling equilibria.^{20,21} Note that these distributions are interim, that is, they are from the information set of a student who knows his own preference, but only the distribution of other students' preferences and the lottery. The cumulative columns make it clear that the outcome distribution under Boston first-order stochastically dominates that of DA.²² Figure 2b shows the equilibrium outcome CDFs in this environment. Note that in this representation, a distribution first-order stochastically dominates another when it is completely *above* it. In Section 5 we provide intuition and further investigate why the Boston outcome first-order stochastically dominates outcomes under DA when students submit preferences truthfully.

2.3. Experimental design

To provide an empirical test of whether agents manipulate their preference reports under the Boston algorithm, we ensure that students always have sufficient information to compute the ordinal Bayes-Nash equilibrium by informing them about the distribution from which preferences and lotteries are drawn.²³ This allows us to do three key analyses. First, we can compare submitted strategies to the equilibrium prediction. Second, we can test whether some participants manipulate their reports to maximize their payoffs given others' reports. Third, we can gain some intuition concerning which deviations from equilibrium may be more common or easy to learn. Students play the experimental game 15 times to allow for learning.²⁴

²⁰Actually, the outcome distribution under single lottery DA can be calculated in closed form, and a recursion exists for calculation the outcome distribution under Boston. Deriving these formulae would take us too far afield from the task at hand, so it is easier to think of these numbers as coming from a bootstrap, as they did for multiple lottery DA, which does not have a neat solution.

²¹There is only one column for Boston because that mechanism does not yield a different distribution over the two lottery structures in this environment.

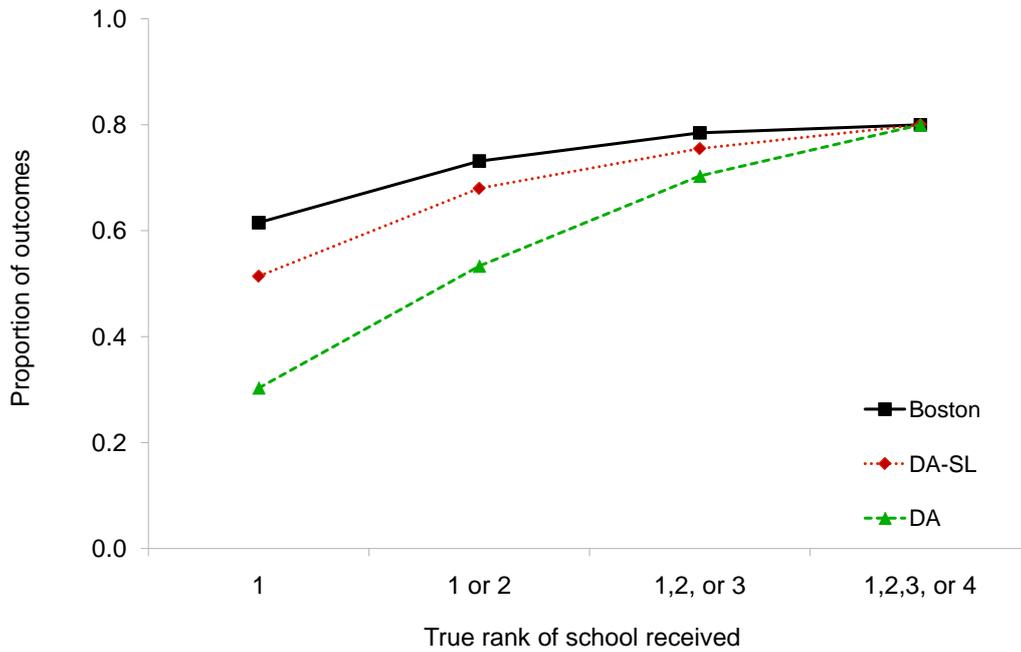
²²We focus on ordinal preferences, as students, in general, submit ordinal rankings. When cardinal preferences are considered, and efficiency is measured as the ex post sum of student welfares, then Boston may dominate DA in some environments, generally by way of a non-truth-telling equilibrium (Abdulkadiroğlu et al. 2012, Miralles 2009).

²³Testing whether and how agents manipulate reports under the Boston algorithm is difficult to address in the field since the preferences of students are unobserved. Abdulkadiroğlu et al. (2006) analyzes submitted preferences and find some evidence of potentially sub-optimal behavior. Barros et al. (2010) use a survey to assess preferences of students who then participate in a Boston-like mechanism and compare preference reports. This method is also used by Budish and Cantillon (2012).

²⁴Repetition allows us to see whether behavior remains stable and whether participants may eventually learn to manipulate in a Boston mechanism and report truthfully in a DA mechanism.

Mechanism	Boston		DA (single lottery)		DA (multiple lotteries)	
	1 st choice	0.610		0.500		0.328
2 nd choice	0.117	(0.727)	0.167	(0.667)	0.222	(0.550)
3 rd choice	0.055	(0.782)	0.083	(0.750)	0.150	(0.700)
4 th choice	0.018	(0.800)	0.050	(0.800)	0.100	(0.800)
No school	0.200	(1.000)	0.200	(1.000)	0.200	(1.000)

(a) Outcome distributions (CDF in parentheses)



(b) CDF plots (DA-SL is single lottery; DA is multiple)

Figure 2: Theoretical equilibrium outcomes in the uncorrelated environment (100% truth-telling)

Our treatments form a 2×2 factorial design: (DA, Boston) \times (Aligned preference environment, Uncorrelated preference environment). The environments and payoffs were discussed in Section 2.2. Payoffs were chosen to provide sufficient incentive for students to play equilibrium.²⁵ In all environments, each school independently draws its own lottery. Note that under DA (but not under Boston), using a single lottery instead of multiple lotteries may have an effect, which we will address in the analysis of Section 4.2.

We ran four sessions under DA and seven sessions under Boston. In each session, five Stanford undergraduate students played for 15 periods in the aligned preference environment, during which players kept their role, as either a Top or an Average student. Then, after a pause to learn about the new environment, they played another 15 periods in the uncorrelated preference environment, where preferences for all students and lotteries for all schools are redrawn in each period.

In our design, participants always see the aligned preference environment before they see the uncorrelated preference environment. This may lead to less truth-telling in the Boston-uncorrelated treatment, which would actually work *against* what we are trying to show. The reason we chose the within-subjects design is that it captures the flavor of the likely strategic transition that students face when a district changes its school choice mechanism: they have to learn to truth-tell after being conditioned to strategically manipulate.

Within a session, the mechanism was held constant, and each subject participated in only one session. The experiment was conducted on computers, using z-Tree (Fischbacher 2007). At the start of a session, after reading instructions concerning the environment and the mechanism, we checked each player’s understanding by having them solve the outcome of a complete information test environment, where participants were given submitted preferences and had to determine the outcome of the relevant mechanism. We repeatedly checked understanding by correcting and explaining outcomes through each subsequent step of the algorithm. Participants earned 1.5 cents for every point and were paid based on their cumulative earnings over all 30 periods of the experiment. In addition participants received a seven dollar show-up payment, bringing average payments to about twenty-five dollars for an average of 60 minutes in the lab.

While parents in general participate in school choice mechanisms only once, they often draw from the experiences of other parents, and many districts have school choice mechanisms at several points in a child’s education. In using multiple rounds, we also follow the tradition of two-sided matching experiments (Kagel and Roth 2000, Ünver 2001, McKinney et al. 2005).

²⁵As a reminder, the cardinal payoffs only matter for equilibrium predictions in one treatment group: when we look at behavior under the Boston mechanism in the aligned preference environment.

3. Results from the aligned preference environment

In the aligned preference treatment, theory predicts truth-telling under DA and equilibrium deviation from truth-telling under Boston. While theory is born out under DA, under Boston, we see deviation from truth-telling, but not according to equilibrium and not in a way easily ascribed to best-responses. To assess the behavior of participants relative to equilibrium play we focus the analysis of this section on the last five periods, and hence on periods where we expect that ample learning opportunities may have helped participants converge to equilibrium (or at least some sort of stable play). The results are qualitatively the same, though slightly worse for the Boston equilibrium prediction, when we focus on all 15 periods instead.

3.1. Strategic behavior under Boston

In the aligned environment, under the Boston mechanism, we did not observe a single period in which every player used risk-neutral Bayes-Nash equilibrium strategies.²⁶ What’s more, the strategic deviation from equilibrium leads to an outcome that significantly deviates from equilibrium, as can be seen in Table 1. Even in this simple environment, participants do not seem to play equilibrium, which casts doubt on the notion that non-truth-telling equilibria might predict behavior in actual school choice settings.

Given that participants do not play equilibrium, we investigate two alternative theories of behavior. The first is that agents are simply truth-telling. This is not the case: the truth-telling rate is only 1.43% for Average students, confirming existing experimental results (Chen and Sönmez 2006, Pais and Pintér 2008, Cal-samiglia et al. 2010).²⁷ Investigating the distribution of submitted first choices, we

School	Best	Second	Third	No school
Top	0.67	0.11	0.05	0.17
	$(\frac{2}{3})$	(0)	$(\frac{1}{3})$	(0)
Average	0.00	0.33	0.43	0.24
	(0)	$(\frac{1}{2})$	(0)	$(\frac{1}{2})$

Table 1: Empirical distribution of outcomes under Boston (equilibrium predictions in parentheses)

²⁶This includes all 15 periods in all 7 sessions.

²⁷We find that 13 out of 14 Average students deviate from truth-telling 80% of the time or more. Eleven (of 14) Average students skip school Best 100% of the time.

School	Best	Second	Third
Top	0.92	0.07	0.01
Average	0.06	0.67	0.27

Table 2: Empirical distribution of first choices under Boston

see in Table 2 that the vast majority of strategies submitted by Average students rank some school other than Best as their first choice. Not ranking unachievable schools seems easy to learn, and Average students do so almost immediately.²⁸ For Top students, the equilibrium response to this behavior is to report Third as their second choice. Combining the facts that 1) Tops report Best as their favorite school 92% of the time, and 2) Tops truth-tell 65.7% of the time, it seems that the Tops are having a much harder time learning to “skip the middle” than Average students did learning to “skip the top”.²⁹

An alternative theory of behavior is that some Top students are failing to best-respond, but that other, more sophisticated Tops are best-responding to the non-equilibrium behavior of the others. To test this, one option would be to look, round by round, for ex-post best-responses; however, this seems too stringent a test, given that participants did not have sufficient information to compute the ex-post best response. Instead, we apply a weaker test by identifying instances in the data where behavior can be unequivocally classified as suboptimal.

One of these is when Top players rank Third as their first choice, but only one player ever does this, and only in one period. This clearly sub-optimal behavior does not seem difficult for Tops to avoid. There are also other, more frequent instances in the data where we can clearly identify sub-optimal behavior. For instance, if all Tops rank Best first, and both Averages rank Second as their first choice and Third as their second choice, then Tops are constrained to best-respond by ranking Third as their second choice. To investigate whether Tops might be best-responding to each other, we will look for instances in sessions where the Tops could best-respond with the same strategy for each of the last 5 periods, that is, we ignore sessions in which the best-response for Tops is still fluctuating in the last 5 rounds.³⁰ The details of this analysis are in Appendix Appendix A; we summarize the results here.

In the one session in our analysis where the Tops’ best-response is truth-telling

²⁸Even in the first five rounds, Average student truth-telling was at 6%.

²⁹In fact, a Wilcoxon signed-rank test comparing truth-telling rates for Top students across the aligned and uncorrelated environments indicates that the difference is not significant ($p = 0.61$).

³⁰These limits are not overly restrictive: we are able to look at 5 of the 7 sessions in the Boston-Aligned treatment.

in the 5 final rounds, all three truth-tell 80% of the time or more. In contrast, in the four sessions of our analysis where the Tops’ best-response is “skipping the middle” only 2 of 12 do so 80% of the time or more. Of those that remain, 7 truth-tell 80% of the time, while the rest fail to play a consistent strategy. So even if we can’t explicitly rule out some “sophisticated” Tops who best-respond to the suboptimal behavior of the others, we see that they are neither a large part of our experimental population nor are they able to steer outcomes closer to equilibrium. It is worth noting that, when analyzing reports submitted to a Boston mechanism in the field, it is a lack of manipulation of this sort (“skipping the middle”) that may be the most identifiable, even if true preferences of students are not known. Our experimental results thus dovetail nicely with Abdulkadiroğlu et al. (2006), which finds suggestive evidence that students sometimes rank schools below their first choice that are expected to be filled in the first step of the algorithm, that is, that many students fail to “skip the middle”.³¹

In short, under Boston, participants are not playing equilibrium, even when we focus only on the last five periods of play, after participants had ample feedback and repetition. This suggests that a non-truth-telling equilibrium might not be predictive of behavior in the field, and hence, gains from such an equilibrium might fail to materialize.

3.2. Truth-telling under DA

The outcome in the last five periods under the DA mechanism is given by Table 3, which shows for each student type, the fraction of matches at each school. The outcome corresponds exactly to the equilibrium: the three Top students receive seats at Best and Second, while Average students either receive a seat at Third or remain unmatched. Note that under DA, compared to Boston, Top students are better off and Average students are worse off.

We begin our investigation of strategies with the submitted first choices (Table 4). Sufficient conditions for the outcome under DA to yield the stable match constrain Top students to submit preferences $(Best, Second, x)$, while only constraining Average students to rank Third as acceptable. In spite of the fact that truth-telling is not a *strict* best response, we observe that 92% of Top and 63% of Average student strategies are truth-telling.^{32,33}

³¹If a school is filled in the first step of a Boston mechanism, then no student who has not ranked the school first has a chance of matching to it. If a student believes that a school will fill in the first step of the algorithm, then it is a clear mistake to rank it lower than first.

³²Note that under DA, 11 out of 12 Top and 4 out of 8 Average students truth-tell 80% of the time or more. In contrast, under Boston the numbers were 13 out of 21 for Top students and 0 out of 14 for Average students.

³³When truth-telling is a weakly dominant strategy, overall truth-telling rates are not significantly different as we change the environment. They are 80% in the aligned and 66% in the uncorrelated environment, ($p = 0.14$) using a Wilcoxon signed rank test.

School	Best	Second	Third	No school
Top	0.67 $(\frac{2}{3})$	0.33 $(\frac{1}{3})$	0.00 (0)	0.00 (0)
Average	0.00 (0)	0.00 (0)	0.50 $(\frac{1}{2})$	0.50 $(\frac{1}{2})$

Table 3: Empirical outcome distribution under DA (equilibrium predictions in parentheses)

School	Best	Second	Third
Top	1.00	0.00	0.00
Average	0.70	0.05	0.25

Table 4: Empirical distribution of first choices under DA

To compare truth-telling rates between Boston and DA, we use Mann-Whitney tests on session averages from our 7 Boston and 4 DA sessions. Running the test for the two student types separately, in the aligned environment, we find that both Top ($p = 0.07$) and Average students ($p < 0.01$) are significantly more likely to use truth-telling strategies under DA than under Boston. This shows that the manipulations of both Average and Top players under the Boston mechanism are not due to the environment alone, but rather are driven by the combination of environment and mechanism.

4. Results from the uncorrelated preference environment

In the uncorrelated environment, under the Boston mechanism, everyone truth-telling is an ordinal Bayes-Nash equilibrium, while under DA, everyone truth-telling is a dominant strategy equilibrium. If students truthfully reveal their preferences, we expect the Boston outcome to first-order stochastically dominate that of DA. In the previous section, we looked at the last 5 periods to allow for ample feedback and repetition, effectively stacking the cards against our finding that the equilibrium is not predictive under Boston. In this section, we hope to show that truth-telling is an easy equilibrium to find, even without much feedback and repetition. To this end, we will stack the cards against ourselves by considering all rounds of play. Furthermore, the theoretical gains of Boston over DA arise in expectation *across* realizations of student preferences, so analyzing all rounds provides us with more data for these gains to take effect.

		True Preferences							
		BOSTON				DA			
		1 st	2 nd	3 rd	4 th	1 st	2 nd	3 rd	4 th
Submitted preferences	1 st	0.76	0.22	0.01	0.01	0.74	0.19	0.07	0.00
	2 nd	0.16	0.61	0.13	0.10	0.14	0.69	0.15	0.02
	3 rd	0.06	0.07	0.80	0.07	0.07	0.11	0.77	0.05
	4 th	0.02	0.10	0.06	0.82	0.05	0.01	0.01	0.93

Table 5: Empirical distribution of submitted ranks in the uncorrelated environment

4.1. Strategies

To compare strategies between Boston and DA, note that basically all submitted strategies rank all schools.³⁴ The proportion of truth-telling strategies is 58% under Boston, compared with 66% under DA.³⁵ This difference is not significant: a Mann-Whitney test across mechanisms, comparing mean truth-telling rates in each session, yields a p -value of 0.70 ($n = 11$).³⁶

To address what manipulations are submitted, we check truth-telling rates at each ranked position, that is, how often a participant’s submitted k^{th} -ranked school corresponds to his true k^{th} -ranked school (see Table 5). Note that there is little discernible difference across mechanisms, as predicted by equilibrium, even though the equilibria are implemented according to different solution concepts. A plausible explanation is that truth-telling holds special sway as a focal strategy.³⁷

4.2. Outcomes

Before we begin our analysis, note that although the DA mechanism we ran in the lab used multiple lotteries, recent work by Abdulkadiroğlu et al. (2009) provides simulations that indicate that DA with a single lottery may produce a better outcome in some environments. While this is the case in our environment, Boston

³⁴Only one student, for three rounds at the beginning of the uncorrelated preference environment, failed to rank all schools.

³⁵That the students truth-tell at a much higher rate in the uncorrelated environment for Boston tells us that they are strategically responding to the environments. A Wilcoxon signed-rank test comparing truth-telling rates for Average students across the aligned and uncorrelated environments indicates that the difference is highly significant ($p = 0.02$).

³⁶Truth-telling rates declined slightly from the first five periods to the last. The rates in the first five periods were 74% under DA and 61% under Boston.

³⁷This idea is not a new one; Pathak and Sönmez (2008) look at a model in which some player are optimally strategic and others sub-optimally truth-tell.

still theoretically dominates DA with either multiple or single lottery tiebreakers. To make sure that Boston is outperforming the single lottery DA (that we did not use in the lab) as well as the multiple lottery DA (that we did use in the lab), we compute the counter-factual outcome had we used a single lottery in the laboratory as a further test.³⁸ We do this by taking the participants’ submitted preferences and using the lottery draw for one of the schools as the universal lottery. Since we have four such lotteries (one per school) each period, we use each of them as a universal lottery and then average the resulting outcome distributions. In the figures, we denote this by “DA-SL”. DA with multiple lotteries, as was used in the actual experiment, is denoted simply by “DA”.

We begin our analysis by recalling Figure 2b, which shows just how much gain over DA is possible if all subjects truth-tell, averaged across all possible preference and lottery draws. Of course, in the lab, it is possible that we get “bad draws” which undermine our potential to recreate the theoretical gains. To show that this is not the case, we construct a truth-telling counterfactual, calculating the outcome distribution given the preference and lottery draws realized in the lab, but constraining the truth-telling rate to be 100%. The CDF for this counterfactual is in Figure 3a. Fortunately, the gains are still quite large and are essentially identical to those that theory predicts: participants would have a 61.5% chance to receive their first choice under Boston, which is significantly larger than the 30.3% under DA with multiple lotteries ($p < 0.01$) and the 51.4% under DA with a single lottery ($p < 0.01$).³⁹ When we compare the proportion of students who receive either their true first or second choice, once more Boston has significantly higher proportions than either DA or DA-SL ($p < 0.01$ in both cases).⁴⁰

Now, we move to the actual outcome distribution observed in the lab. Even with a less than perfect truth-telling rate, we observe a significantly larger fraction of students receive their true first choice with Boston than with DA ($p < 0.01$), or DA-SL ($p = 0.06$).⁴¹ When we compare the proportion of students who receive either their first or second choice, Boston once more significantly outperforms both DA ($p < 0.01$) and DA-SL ($p = 0.09$).⁴² These findings can be seen graphically in

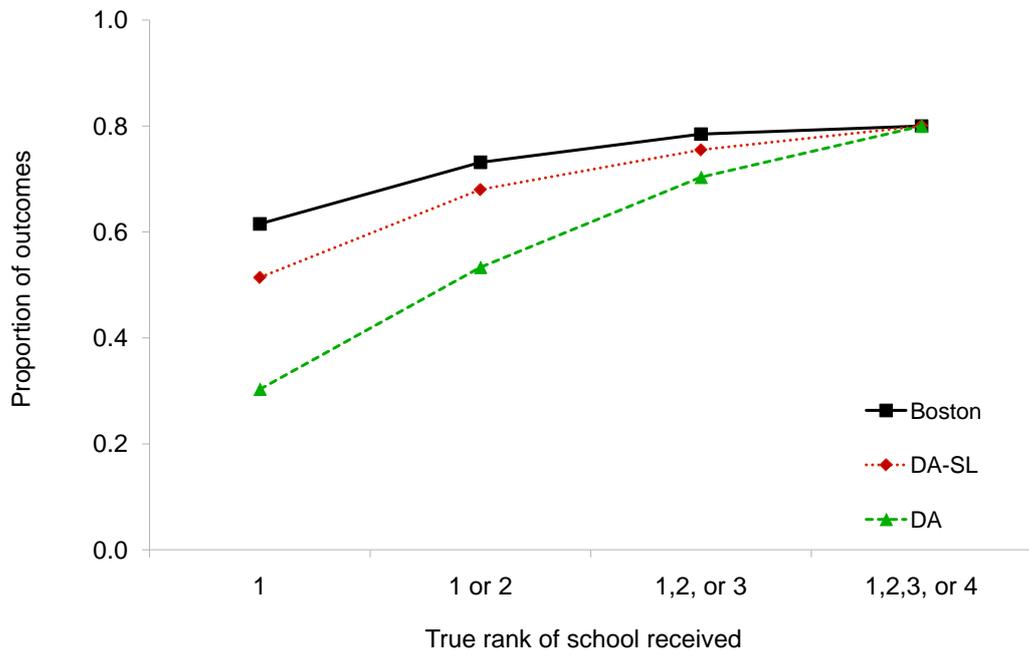
³⁸It is not necessary to consider multiple versus single lottery setups under Boston: both yield the same distribution over outcomes.

³⁹We use Mann-Whitney tests, where the session mean is a data point, that is, we have 7 data points for Boston, and 4 for DA and DA-SL. When we only consider the last five periods, the p -values are 0.01 and less than 0.01, respectively.

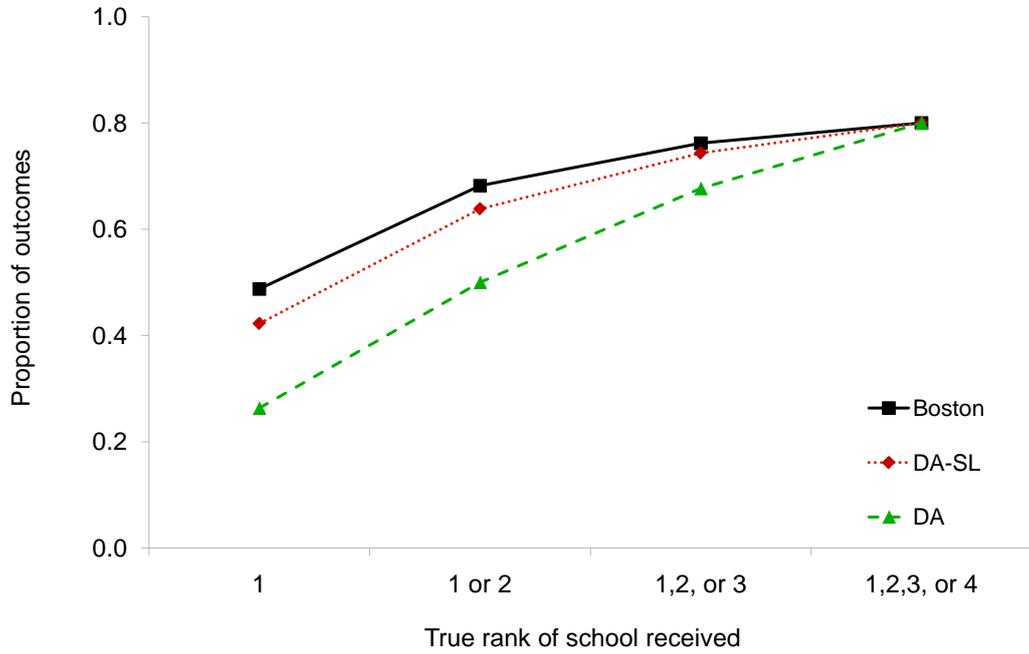
⁴⁰When we look at the last five rounds only, the difference between Boston and DA is still significant at $p < 0.01$, while the one between Boston and DA-SL is not ($p = 0.12$).

⁴¹When we only consider the last five periods, Boston still gives a higher fraction of participants their first choice. The difference is significant when compared to DA ($p < 0.01$), though not when compared to DA-SL, ($p = 0.18$). Note that a one-sided test would yield significance in all the two-sided tests that we have discussed.

⁴²When we consider only the last five periods, the p values are < 0.01 and 0.03 respectively.



(a) Lab outcomes under 100% truth-telling counterfactual



(b) Actual lab outcomes

Figure 3: Outcome CDFs with preferences realized in lab (DA-SL is the single lottery counter-factual)

Figure 3b.

Hence, it seems that the theoretical advantages of Boston over DA can indeed be realized in practice, at least in our specialized environment. This provides a proof-of-concept that non-strategy-proof mechanisms might perform well in the field if they are ordinal Bayes-Nash incentive-compatible.

5. Why did Boston succeed?

5.1. The art and science schools example

To understand why Boston outperforms DA in the uncorrelated environment, it helps to simplify. Suppose there are three students and two one-seat schools – an art school and a science school. Each student is equally likely to be an artist (**a**) who prefers the art school or a scientist (**s**) who prefers the science school. All students find both schools acceptable, and a common lottery determines priorities at both schools. In this environment, corollaries to the same Appendix Appendix B propositions mentioned in Section 2.2 tell us that truth-telling is the unique equilibrium under both Boston and DA.

We will consider the outcomes of the two mechanisms both from the interim and ex ante points of view. States of the world from the ex ante perspective can be represented as a string of three **a**'s or **s**'s; for instance, **aas** represents the state where the students with the highest two lottery draws are both artists, and the remaining student is a scientist.⁴³ For interim states,⁴⁴ we will use a capital letter to denote the student whose perspective we are considering; for instance, **aaS** represents the state of the world we just mentioned, from the perspective of the scientist with the lowest lottery draw.

The ex ante perspective is that of a social planner (the “superintendent”) who thinks about the quality of overall matches in terms of the distribution of how students rank their assigned schools.⁴⁵ In this environment, such social preferences are equivalent to maximizing the probability that an agent gets his first choice.⁴⁶ When comparing Boston and DA, the superintendent only cares about the two states in which the two mechanisms yield different matches: **aas** and **ssa** (see Figure 4a). Intuitively, these are states where DA gives a school to a student who

⁴³Since we will only consider anonymous mechanisms, there is no need to keep track of whether the first artist is Student 1, 2, or 3. This simplifies the state space considerably without losing any of the intuition.

⁴⁴That is, states from the perspective of a student who knows her own preferences, but not those of others and not the lottery draw.

⁴⁵This is a common and natural way to aggregate welfare with limited information. See Featherstone (2013) for more.

⁴⁶Everyone is unmatched one-third of the time by symmetry, so the ex ante probability of getting a first choice is sufficient for the entire distribution of ranks received.

aaa aas ass sss asa sas saa ssa		S aa a S a aa S S as a S s as S S sa s S a sa S S ss s S s ss S	
<p>(a) The 8 (equally likely) states of the world from the ex ante perspective. Boxed cells indicate cells where Boston and DA give different overall matches.</p>		<p>(b) The 12 (equally likely) states of the world from the (interim) perspective of scientist S. Boxed cells indicate that Boston and DA give different schools to S.</p>	

Figure 4: States of the world: ex ante and interim.

ranks it second, even though there is an unmatched student who ranked it first. Note that Boston never does this *in any environment*.⁴⁷ So, assuming truth-telling, from the ex ante perspective, Boston is either the same as or better than DA in every state of the world. This is the precise way in which Boston outperforms DA.

Instead of Boston, however, consider running DA with the addition of a simple ex post reallocation rule: when the state of the world is **aas**, take the science school from the second artist and give it to the scientist. Assuming truth-telling, this clearly makes the superintendent better off, as it is the same as or better than DA in all states of the world. While ex ante students are better off, from the interim perspective the science student **S** is better off while an art student **A** is worse off. This may give an artist who is almost indifferent between the art school and the science school an incentive to masquerade as a scientist, since such an artist is more worried about being unmatched when the state is **aas** than about sometimes getting the science school instead of the art school. Analogously, we can also consider the “mirror image” of our **aas**-contingent reallocation: an **ssa**-contingent rearrangement in which we take the art school from the second scientist and give it to the artist. By symmetry, this would make the superintendent and the artist better off, but would hurt the scientist and potentially give her incentive to misreport her preferences. It is not surprising that trying to improve upon DA ex post could cause incentive problems (Erdil and Ergin 2008, Abdulkadiroğlu et al. 2009, Featherstone 2013). What is surprising in our example is that if we combine *both* of our problematic reallocation rules, we return to implementing truth-telling, although as a Bayes-Nash equilibrium instead of a dominant strategy equilibrium. In other words, Boston is equivalent to running DA and reallocating in the “bad” states of the world (**aas** and **ssa**).

To get a better look at why this works, consider the perspective of a scientist

⁴⁷Kojima and Ünver (2013) show that Boston “respects preference rankings”, which implies the statement above. Intuitively, this is because, under Boston, if a student is rejected by his k^{th} ranked school, it must have filled up by being assigned students who ranked it k^{th} or higher.

(capital-S) considering her preference report when both the **aas** and **ssa** reallocation rules are in effect. She cares most about the interim states **sSa** and **aaS**, as in those states, one of the reallocation rules changes the school she gets (see Figure 4b). Although both **sSa** and **aaS** are equally likely from her perspective, the help she gets in **aaS** (moving from unmatched to a first choice) is bigger than the hurt she gets in **sSa** (moving from second choice to unmatched). So, she is interim better off under the new rules. Moreover, she has no incentive to misrepresent herself as an artist: if she did so, the hurt (having a first choice taken away because it was *declared* as a second choice) would be bigger than the help (moving from unmatched to a second choice because it was *declared* as a first choice). By symmetry, an artist also prefers to truthfully report his preference.

Interpreting Boston as DA plus two ex post reallocation rules shows both why Boston dominates and why it admits a truth-telling equilibrium. However, much of this analysis depends on the fact that the two reallocation rules are equally likely to be triggered from the interim perspective. In the next section, we consider generalizing to a less symmetric environment.

5.2. Generalizing to less symmetric environments

Let the ex ante probability that a student is a scientist be $k > \frac{1}{2}$. Assuming truth-telling, in this new environment, Boston would still dominate DA state by state from the ex ante perspective; the fact that $k > \frac{1}{2}$ only changes the relative likelihood of those states. Because of this, however, Boston will no longer admit a truth-telling equilibrium. This is because a scientist, **S**, might not be pleased that both rules are being added to DA, since from his perspective, the state in which he is hurt, **sSa**, is more likely ($p_{\text{sSa}} = \frac{1}{3} \cdot k \cdot (1 - k)$) than the state in which he is helped, **aaS** ($p_{\text{aaS}} = \frac{1}{3} \cdot (1 - k)^2$). A simple calculation shows that a scientist who just barely prefers the science school would prefer to masquerade as an artist in order to increase the probability that the **ssa** reallocation rule would save him from being unmatched.⁴⁸

If we could equalize the likelihood that the two rules change a scientist's assignment, however, he would again prefer to truthfully report his preferences, by the same logic as the previous section. This can be accomplished by only implementing the **ssa** reallocation rule with probability $\frac{1-k}{k}$, since $p_{\text{sSa}} \cdot \frac{1-k}{k} = p_{\text{aaS}}$. Intuitively, the superintendent is tempted to reallocate 100% of the time, but he

⁴⁸If the scientist tells the truth, then his gain from the reallocation rules is $\frac{1}{3} \cdot k \cdot (1 - k) \cdot [u(\emptyset) - u(a)] + \frac{1}{3} \cdot (1 - k)^2 \cdot [u(s) - u(\emptyset)]$ which approaches $\frac{1}{3} \cdot (1 - k) \cdot (2k - 1) \cdot [u(\emptyset) - u(s)] < 0$ as $u(a) \rightarrow u(s)$. If the scientist misrepresents himself as an artist, his gain is $\frac{1}{3} \cdot k \cdot (1 - k) \cdot [u(\emptyset) - u(a)] + \frac{1}{3} \cdot k^2 \cdot [u(s) - u(\emptyset)]$ which approaches $\frac{1}{3} \cdot (1 - k) \cdot (2k - 1) \cdot [u(s) - u(\emptyset)] > 0$ as $u(a) \rightarrow u(s)$.

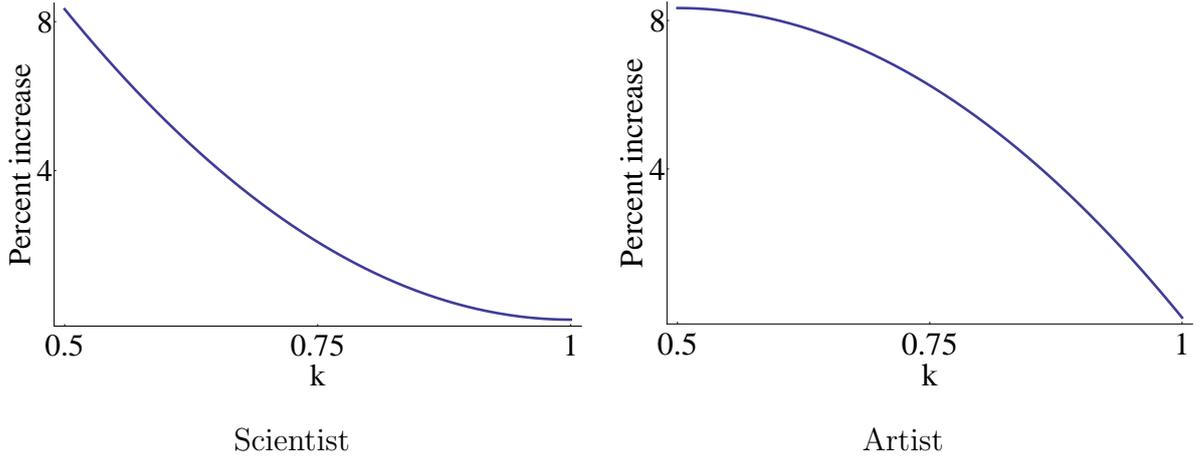


Figure 5: Increase in chance to receive the first choice school from DA to the Boston-like new mechanism

must commit himself to only do so probabilistically in order to induce truthful preference revelation.

Proposition 5.1. *Consider the extension of the arts and science school example in this section, where each student is independently a scientist with probability $k > \frac{1}{2}$. Suppose we run DA and then apply the following two ex post reallocation rules:*

- *If the state is **aas**, then give the science school to the scientist, leaving the artist with the lower lottery draw unmatched.*
- *If the state is **ssa**, then with probability $1 - \frac{1-k}{k}$ do nothing, and with probability $\frac{1-k}{k}$, give the art school to the artist, leaving the scientist with the lower lottery draw unmatched.*

Then, all students truth-telling is an ordinal Bayes-Nash equilibrium. Further, all students interim⁴⁹ prefer the equilibrium outcome of this mechanism to the outcome of the truth-telling equilibrium under DA.

When $k = \frac{1}{2}$, the second rule in the theorem reallocates with probability 1, that is, the mechanism in the theorem becomes Boston. Figure 5 shows just how much mechanism of Proposition 5.1 dominates DA.⁵⁰

While the example is only suggestive, the intuition of probabilistically randomizing over whether a school is allowed to ex post reallocate seems generalizable and is a possible area of future research.

⁴⁹That is, evaluated from the information set where students know their own preference realization, but only the distribution over others' preferences and the lottery draw.

⁵⁰It is sufficient to show only how much more likely a student is to receive his first choice school, because truth-telling would not be an ordinal Bayes-Nash equilibrium if artists and scientists were unmatched differentially, contradicting Proposition 5.1.

6. Conclusion

We had two goals for our experiment. The first concerns the new and active literature on mechanisms that rely on equilibrium manipulation of preference reports to implement outcomes superior to those of DA (Abdulkadiroğlu et al. 2012, Miralles 2009, Troyan 2012). To assess the viability of these approaches for practical market design, it is important to determine whether participants are likely to at least come close to equilibrium play. Empirical work suggests that this may not be the case for at least a few participants (Abdulkadiroğlu et al. 2005b). Existing experiments have focused on whether participants submit preferences truthfully in non-strategy-proof mechanisms, and while these papers confirm that participants would manipulate their preferences, they in general do not provide participants with sufficient information to calculate the equilibrium (Chen and Sönmez 2006, Pais and Pintér 2008, Calsamiglia et al. 2010).

The first result of the present paper is that non-truth-telling equilibria might be hard to implement in the real world. In the aligned preference environment, while DA elicited truth-telling, Boston elicited sub-optimal manipulation of preference reports. This was the case even in the last five periods of play, where participants had ample experience and received precise feedback in a very simple environment. This leaves two potential ways to move the “improving on strategy-proof mechanisms” literature forward. The first is to give serious thought as to how students can be pushed in the direction of equilibrium. Inquiry into this question seems like a fruitful and understudied line of experimental research, but is not what we pursue in this paper. The second is to weaken the implementation of the truth-telling equilibrium.⁵¹

The second part of our paper shows that truth-telling equilibria that are not implemented in dominant strategies do have the potential to succeed, both in inducing truthful reporting and in yielding efficiency gains. We think of our results as showing that there is *potential* for the incentive-compatible-but-not-strategy-proof research agenda to succeed. Further theoretical development is clearly required, but, as we hope our experiment has demonstrated, such a research agenda could potentially have a large impact on student welfare.

In addition to simply pointing out that Boston can dominate DA in some environments, we have also tacitly made two suggestions for non-strategy-proof design. The first is to implement truth-telling as an ordinal Bayes-Nash equilibrium, as it requires information about the distribution of ordinal preferences (which can be observed under a strategy-proof mechanism), but not about the cardinal prefer-

⁵¹Barros et al. (2010) runs an experiment with a decision aid that pursues this line of inquiry.

ences, which are more elusive.⁵² The second is to analyze efficiency and incentives from the interim information set, in which a student knows his own preferences, but only the distribution of others'. Not only does this seem like the natural assumption for what real-world students know, but from an efficiency standpoint, it seems like the sort of “best outcome averaged over many years” sort of calculus that seems natural for a school board.

References

- Abdulkadiroğlu, A., Che, Y.-K., Yasuda, Y., 2011. Resolving Conflicting Preferences in School Choice: The “Boston” Mechanism Reconsidered. *American Economic Review* 101(1), 399–410.
- Abdulkadiroğlu, A., Che, Y.-K., Yasuda, Y., 2012. Expanding “Choice” in School Choice. working paper.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., 2005a. The New York City High School Match. *American Economic Review* 95 (2), 364–367.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., 2009. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *The American Economic Review* 99 (5), 1954–1978.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., Sönmez, T., 2005b. The Boston Public School Match. *The American Economic Review* 95 (2), pp. 368–371.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., Sönmez, T., 2006. Changing the Boston School Choice Mechanism. working paper.
- Abdulkadiroğlu, A., Sönmez, T., 2003. School Choice: A Mechanism Design Approach. *The American Economic Review* 93 (3), pp. 729–747.
- Barros, R., Featherstone, C. R., Salcedo, A., 2010. Matching university students to majors: analysis of an old mechanism and redesign of the new. Working Paper.
- Budish, E., Cantillon, E., 2012. The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. *American Economic Review* 102 (5), 2237–2271.
- Calsamiglia, C., Haeringer, G., Klijn, F., 2010. Constrained School Choice: An Experimental Study. *American Economic Review* 100(4), 1860–74.
- Calsamiglia, C., Haeringer, G., Klijn, F., January 2011. A comment on “School choice: an experimental study” [J. Econ. Theory 127 (1)(2006) 202-231]. *Journal of Economic Theory* 146 (1), 392–396.
- Chen, Y., Sönmez, T., 2006. School choice: an experimental study. *Journal of Economic Theory* 127 (1), 202–231.
- Chen, Y., Sönmez, T., 2011. Corrigendum to “School choice: an experimental study” [J. Econ. Theory 127 (1)(2006) 202-231]. *Journal of Economic Theory* 146 (1), 397–399.
- Dubins, L. E., Freedman, D. A., 1981. Machiavelli and the Gale-Shapley Algorithm. *The American Mathematical Monthly* 88 (7), pp. 485–494.

⁵²For instance, in any school district currently running a strategy-proof mechanism, we have good reason to believe that submitted ordinal preferences are truthful. From these, we can construct a model of the distribution from which the ordinal preferences came. Note that these preference reports contain no information about the underlying cardinal preferences (except that they should rationalize the submitted ordinal preferences).

- Ehlers, L., Massó, J., 2007. Incomplete information and singleton cores in matching markets. *Journal of Economic Theory* 136 (1), 587–600.
- Erdil, A., Ergin, H., 2008. What’s the Matter with Tie-Breaking? Improving Efficiency in School Choice. *The American Economic Review* 98 (3), 669–689.
- Ergin, H., Sönmez, T., 2006. Games of School Choice under the Boston Mechanism. *Journal of Public Economics* 90 (1-2), 215–237.
- Featherstone, C. R., 2013. Rank efficiency: Investigating a widespread ordinal welfare criterion. working paper.
- Fischbacher, U., 2007. z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics* 10 (2), 171–178.
- Gale, D., Shapley, L. S., 1962. College Admissions and the Stability of Marriage. *American Mathematical Monthly* 69 (1), 9–15.
- Kagel, J. H., Roth, A. E., 2000. The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment. *Quarterly Journal of Economics* 115 (1), 201–235.
- Kesten, O., Ünver, M. U., 2013. A Theory of School Choice Lotteries. working paper.
- Kojima, F., Pathak, P. A., 2009. Incentives and Stability in Large Two-Sided Matching Markets. *The American Economic Review* 99 (3), 608–627.
- Kojima, F., Ünver, M. U., 2013. The “Boston” school-choice mechanism: an axiomatic approach. *Economic Theory*, 1–30.
- McKinney, C. N., Niederle, M., Roth, A. E., 2005. The Collapse of a Medical Labor Clearinghouse (and why such failures are rare). *American Economic Review*, 878–889.
- Miralles, A., 2009. School Choice: The Case for the Boston Mechanism. *Auctions, Market Mechanisms and Their Applications*, 58–60.
- Pais, J., Pintér, Á., 2008. School Choice and Information: An Experimental Study on Matching mechanisms. *Games and Economic Behavior* 64 (1), 303–328.
- Pathak, P. A., Sönmez, T., 2008. Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *The American Economic Review* 98 (4), 1636–1652.
- Roth, A. E., 1982a. Incentive Compatibility in a Market with Indivisible Goods. *Economics Letters* 9 (2), 127–132.
- Roth, A. E., 1982b. The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research* 7 (4), 617–628.
- Roth, A. E., 1991. A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom. *The American Economic Review* 81 (3), 415–440.
- Roth, A. E., Rothblum, U. G., 1999. Truncation Strategies in Matching Markets – in search of advice for participants. *Econometrica* 67 (1), 21–43.
- Troyan, P., 2012. Comparing school choice mechanisms by interim and ex-ante welfare. *Games and Economic Behavior* 75(2), 936–947.
- Ünver, M. U., 2001. Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets. *Journal of Economic Dynamics and Control* 25 (6-7), 1039–1080.

Appendix A. Best-response analysis for aligned environment under Boston

There are two cases present in the data which give Top students strict incentives to rank either Second or Third as their second choice, given that they ranked Best as their first choice

Case 1 (Skip the middle): All Top students rank Best as their first choice, at least one Average ranks Second as his first choice, and one of the Averages who might not be assigned in the first round ranks Third as his second choice. Tops are then constrained to rank Third as their second choice.

Case 2 (Do not skip the middle): All Top students rank Best as their first choice, at least one Average ranks Third as his first choice, and one of the Averages who might not be assigned in the first round ranks Second as his second choice. Tops are then constrained to rank Second second.

We restrict ourselves to studying sessions in which a large fraction (4 out of 5) or at least the majority (3 out of 5) of periods are either in Case 1 or Case 2 and in which the remaining periods do *not* fall into any of these two cases. That is, we focus on sessions in which Top students weakly best respond by submitting $(Best, Second, x)$ or $(Best, Third, x)$ in all five periods, and, in 4 (or 3) of those periods, this behavior is a *strict* best response. Table A.6 shows the number of sessions that fall into these categories and whether Top students submit truthfully (i.e. submit $(Best, Second, x)$) or best respond as a primary strategy. Note that table below analyzes 4 distinct “Skip the middle” sessions.

Best response for Top students (BR)	# sessions where BR is weakly optimal in each of the last 5 rounds (# rounds where the BR is strict)	# of Tops who truth-tell at least 4 of the last 5 rounds	# of Tops who BR in at least 4 of the last 5 rounds
Skip the middle	2 (4/5)	5/6	0/6
Skip the middle	2 (3/5)	2/6	2/6
Don't skip the middle	1 (4/5)	3/3	3/3

Table A.6: Unambiguously suboptimal behavior of Top students

Appendix B. Proofs

B.1. The model

Let there be a finite set of students, \mathcal{I} , and a finite set of schools, \mathcal{S} . Also, let $\emptyset \notin \mathcal{S}$ represent the outcome of being unmatched. Each student i has a

strict preference P_i over the elements of $\mathcal{S} \cup \{\emptyset\}$, and each school has a strict priority ordering \succ_s over \mathcal{I} .⁵³ Denote the number of acceptable schools in a preference P_i by $|P_i| \equiv |\{s \in \mathcal{S} : s P_i \emptyset\}|$. As we are dealing with individually rational mechanisms, we will only consider the ordering of schools in P_i that are ranked above \emptyset , that is, if two preferences order their acceptable schools in the same way, then they are considered the same preference. A utility function u_i rationalizes a preference P_i if $u_i(P_i(a)) > u_i(P_i(b)) \Leftrightarrow a < b$, where $P_i(a)$ denotes the a^{th} ranked school in P_i . $q \equiv (q_s)_{s \in \mathcal{S}}$ is a vector of capacities for each school. The unmatched outcome, \emptyset , can be thought of as a special “school” with $q_\emptyset = \infty$. $P \equiv (P_i)_{i \in \mathcal{I}}$ and $\succ \equiv (\succ_s)_{s \in \mathcal{S}}$ are drawn from a joint probability distribution $\lambda(P, \succ)$.

A *school choice problem* is defined by $\langle \mathcal{I}, \mathcal{S}, q, \lambda \rangle$. A *matching* $\mu : \mathcal{I} \mapsto \mathcal{S} \cup \{\emptyset\}$ is a function such that $|\{i : \mu(i) = s\}| \leq q_s$ for all $s \in \mathcal{S}$. A *deterministic mechanism* is a function φ that maps reported student preferences \tilde{P} and school priorities \succ to a matching. Let $\varphi_i(\tilde{P}; \succ)$ be shorthand for $\varphi(\tilde{P}; \succ)(i)$. For a student, a strategy in the Bayesian game induced by a school choice problem and a deterministic mechanism is a mapping from true preferences to reported preferences, σ_i ; let $\sigma \equiv (\sigma_i)_{i \in \mathcal{I}}$ denote strategy profile. A strategy profile is a *Bayes-Nash equilibrium* if no student has a profitable unilateral deviation, that is, $\forall i, P_i, \tilde{P}_i$

$$\mathbb{E} \{u_i[\varphi_i(\sigma(P); \succ) | P_i] \geq \mathbb{E} \left\{ u_i \left[\varphi_i \left(\tilde{P}_i, \sigma_{-i}(P_{-i}); \succ \right) \right] \middle| P_i \right\}$$

A set of strategies is an *ordinal Bayes-Nash equilibrium* if a student does weakly worse (in the first-order stochastic sense) when he deviates from equilibrium, that is, $\forall i, P_i, \tilde{P}_i, |P_i| \geq b \geq 1$

$$\sum_{a=1}^b \Pr \{ \varphi_i(\sigma(P); \succ) = P_i(a) | P_i \} \geq \sum_{a=1}^b \Pr \left\{ \varphi_i \left(\tilde{P}_i, \sigma_{-i}(P_{-i}); \succ \right) = P_i(a) \middle| P_i \right\}$$

B.2. The aligned environment

Proof of Proposition 2.1. Recall that truth-telling is a weakly dominant strategy for the proposing side in two-sided matching. For our one-sided setting, this means that truth-telling does at least as well as any other strategy *ex post*, that is, after other students’ preferences and the lottery draws are known. Hence, if we can show one lottery draw that occurs with positive probability, where truth-telling is strictly better for some student i than some alternative strategy \tilde{P}_i , regardless of

⁵³Note that \succ_s is not an ordering over $\mathcal{I} \cup \{\emptyset\}$, that is, schools are required to find all students acceptable.

the other students' submitted preferences, then we have shown that \tilde{P}_i is strictly dominated by truth-telling for i . Using this approach, we apply iterated deletion of strictly dominated strategies (IDOSDS).

Round 1: Delete any strategies played by Tops that fail to rank Best first. This is valid because a Top is ranked first by all schools with positive probability. When this happens, ranking Best first is strictly better than ranking anything else first.

Round 2: Delete any strategies played by Tops that fail to rank Best first and Second second. This is valid because a Top is ranked third by all schools with positive probability. With the strategies that remain, the Tops who are ranked first and second at Best will get Best, and our third-ranked Top will get whatever he ranks second. Ranking Second second is strictly better than anything else that remains.

Round 3: Delete any strategies played by Averages that fail to report Third acceptable. This is valid because an Average is ranked fourth by all schools with positive probability. With the strategies that remain after Round 2, the Tops will take Best and Second, leaving one seat at Third. If our fourth-ranked Average reports Third acceptable, he gets it, which is strictly better than going unmatched, which is what would happen if he failed to report Third acceptable.

So far, IDOSDS has left us with the family of strategy profiles in the statement of the proposition. As IDOSDS cannot remove strategies used in equilibrium, any equilibrium profile must be in our family. Now, we show that all profiles in the family are equilibria. First, note that relative to any of these profiles, any stable match must give Best to the two Tops with the best lottery draw at Best, Second to the other Top, and Third to the Average with the best lottery draw at Third. The remaining Third is unmatched, which gives us the outcome equivalence result in the statement of the proposition. As the lottery is fair and independent of reported preferences, any profile in our family will lead to a Top getting Best $\frac{2}{3}$ of the time and Second $\frac{1}{3}$ of the time, and an Average getting Third half the time and remaining unmatched the other half of the time. Hence, all students are equally well off in any of the profiles of our family, which shows that all profiles in the family of profiles that survive IDOSDS are equilibria. \square

Proof of Proposition 2.2. We begin by showing that all three Tops must rank Best first in equilibrium. We do so by ruling out the other possibilities. First, by way of contradiction, assume an equilibrium where two or more Tops fail to rank Best

first. Of those Tops, one must not be getting Best 100% of the time. If this Top were to rank Best first, he would get Best 100% of the time, which would be a profitable deviation, a contradiction.

Now, by way of contradiction, assume an equilibrium in which exactly two Tops rank Best first. The Top who doesn't rank Best first will get whatever else he ranks first; hence, he must rank Second first. Given these rankings from the Tops, the Averages can only get Third or remain unassigned. In equilibrium, they must both rank Third first. To see this, consider the cases. If exactly one Average were ranking Third first, then the other Average could profitably deviate by ranking Third first as well. If no Averages rank Third first, then one of them who isn't getting Average 100% of the time could do so by ranking Average first, a profitable deviation. Hence, we have two Tops ranking Best first, one Top ranking Second first, and both Averages ranking Third first. But in this scenario, the Top who ranks Second first can profitably deviate by ranking Best first and Second second, since Second will still remain for him in round two of the Boston algorithm. This is a contradiction of our initial hypothesis.

So, we have shown that there must be three Tops ranking Best first in any pure-strategy equilibrium. This means that both Averages must rank Second first in equilibrium. To see this, we assume otherwise and check the cases. If neither Average ranks Second first, then one of them must not be getting Second 100% of the time. But that Average could guarantee himself Second by ranking Second first, a contradiction. If exactly one Average ranks Second first, then the other Average is getting a payoff of 25 (from Third) at best. But by ranking Second first, that Average could guarantee himself a payoff of at least $\frac{1}{2} \cdot 67 = 33.5 > 25$, a contradiction.

So, we have shown that in any pure-strategy equilibrium, all Tops rank Best first, and both Averages rank Second first. This means Best and Second will be filled in the first round of the Boston algorithm, leaving one Top and one Average behind to vie for the seat at Third in the second round. In equilibrium, the Average cannot get the seat, as if he did, the Top could profitably deviate by ranking Third second. Moreover, in equilibrium, the Top must get Third in the second round of the algorithm, as if he didn't, the Average could guarantee himself the seat by ranking Third second, contradicting the "lemma" of the previous sentence. And since all Tops have a $\frac{1}{3}$ chance of being the Top in the second round of the algorithm, they must all rank Third second.

Hence, in any pure-strategy equilibrium, all Tops rank Best first and Third second, and all Averages rank Third first, demonstrating the "only if" part of the proposition. For the "if" part, it is easy to check that in any profile of strategies that meets our conditions, there is no profitable unilateral deviation for any player.

Finally, note that the conditions pin down the outcome of the match: two Tops get Best, one Top gets Third, one Average gets Second, and one Average goes unmatched. \square

B.3. The uncorrelated environment

We begin our analysis of the uncorrelated environment by looking at DA. Under DA, there is always a truth-telling ordinal Bayes-Nash equilibrium.

Lemma B.1. *Under DA, for any school choice problem, there exists an ordinal Bayes-Nash equilibrium where all students are truth-telling.*

Proof. Since truth-telling is weakly dominant for the proposing side in two-sided DA, we know that in our one-sided DA setup, by reporting truthfully, a student guarantees himself a weakly better assignment for every realization of reported preferences and priorities, $(\widetilde{P}_{-i}; \succ)$. This state-by-state weak dominance of truth-telling is sufficient for truth-telling to yield a distribution over outcomes that weakly first-order stochastically dominates any other response to any profile of other students' strategies. Hence, all students truth-telling is an ordinal Bayes-Nash equilibrium. \square

In the uncorrelated environment, this equilibrium is unique.

Lemma B.2. *The revelation game induced by DA in the uncorrelated environment is dominance solvable.*

Proof. To show that truth-telling is the unique Bayes-Nash equilibrium, we will implement the same strategy from the proof of Proposition 2.1, using the following rounds of iterated deletion of strictly dominated strategies (IDOSDS):

Round 1: Delete all strategies that don't report rank 1 truthfully. If a student who failed to do so were ranked first at all schools, which occurs with positive probability, then he would do strictly better to tell the truth.

Round 2: Delete all strategies that don't report ranks 1 and 2 truthfully. If all students share the same preference, and one of them were ranked second at all schools, then that student would do strictly better to report rank 2 truthfully (as he will not get his truthfully reported rank 1 school). This event occurs with positive probability.

\vdots

The logical end of this progression is that all strategies except for truth-telling are deleted. In other words, the revelation game induced by DA in the uncorrelated environment is dominance solvable. \square

These two lemmas establish Proposition 2.3 from the main text.

Proof of Proposition 2.3. Lemma B.1 establishes that all students truth-telling is an ordinal Bayes-Nash equilibrium, and Lemma B.2 established uniqueness, since dominance solvable games have a unique equilibrium. \square

Now, we move onto the behavior of Boston in the uncorrelated environment. We will focus on school choice problems that have a certain symmetry property. Let $P_{-i}^{s \leftrightarrow s'}$ denote the preferences constructed by switching schools s and s' everywhere in all students' preferences, save for i . Similarly, let $\succ^{s \leftrightarrow s'}$ be the vector of schools' orderings where the orderings of schools s' and s are switched.

Definition B.1. A probability distribution λ is **school-symmetric with respect to i** if, for any P_i that occurs with positive probability and any $s, s' \in \mathcal{S}$, we have $\lambda[(P; \succ) | P_i] = \lambda[(P_i, P_{-i}^{s \leftrightarrow s'}; \succ^{s \leftrightarrow s'}) | P_i]$. A school choice problem is school-symmetric with respect to i if its distribution is school-symmetric and $q_s = q_{s'}, \forall s, s' \in \mathcal{S}$. We simply call a distribution or school-choice problem **school-symmetric** if it is school-symmetric with respect to i for all $i \in \mathcal{I}$.

This definition does not allow for interchanges that involve \emptyset , as $\emptyset \notin \mathcal{S}$. Also, note that school-symmetry is not the same as drawing preferences from the uniform distribution; for instance, it allows for a student to know that it is unlikely that the other students find as many schools acceptable as he does. With school-symmetric reported preferences, we can show that truth-telling does at least as well as any other strategy a student might play. We do this via two lemmas. Let $B_i[P; \succ]$ denote student i 's assignment under Boston when the reported preferences are P and the school priorities are \succ .

Lemma B.3. *Let $s' P_i s$. Then,*

- a) *If i fails to get s or s' when reporting either P_i or $P_i^{s' \leftrightarrow s}$, then he gets the same school when reporting either preference, that is,*

$$B_i \left[P_i^{s' \leftrightarrow s}, P_{-i}; \succ \right] \notin \{s', s\} \quad \text{and} \quad B_i [P; \succ] \notin \{s', s\}$$

$$\Rightarrow \quad B_i \left[P_i^{s' \leftrightarrow s}, P_{-i}; \succ \right] = B_i [P; \succ]$$

- b) *If i fails to get s' under P_i , then he will also fail to get it under $P_i^{s' \leftrightarrow s}$, that is,*

$$B_i [P; \succ] \neq s' \quad \Rightarrow \quad B_i [P_i^{s' \leftrightarrow s}, P_{-i}; \succ] \neq s'$$

- c) *If i gets s under P_i , then he also gets it under $P_i^{s' \leftrightarrow s}$, that is,*

$$B_i [P; \succ] = s \quad \Rightarrow \quad B_i [P_i^{s' \leftrightarrow s}, P_{-i}; \succ] = s$$

- d) *It never hurts i to report a school acceptable if and only if it is truly acceptable.*

Proof. Recall that what the Boston algorithm does in a given round depends solely on who isn't matched yet and which seats are left. Define j , k , and m by $P_i(j) = s'$, $P_i(k) = s$, and $P_i(m) = B_i[P; \succ]$, respectively.

- a) By assumption, $m \neq j$. If $m < j$, then i must also be assigned in round m when he reports $P_i^{s' \leftrightarrow s}$, as switching s and s' won't change what seats and students are still available at the start of the round m . If $m > j$, then by our assumptions that $B[P_i^{s' \leftrightarrow s}, P_{-i}; \succ](i) \notin \{s', s\}$ and $B[P; \succ](i) \notin \{s', s\}$, the seats at both s and s' must be exhausted at the start of round j . Hence, the set of students and seats that remain available is the same for both reports going forward from round j . So i is assigned in round m under both preference reports.
- b) Since i is rejected in round j when he reports P_i , there must not be any seats available at s' at the end of round j . Since he didn't take a seat at s' when he reported P_i , there must also not be any seats available at the end of round j when he reports $P_i^{s' \leftrightarrow s}$. Hence he cannot get s' by ranking it $k > j$.
- c) P_i and $P_i^{s' \leftrightarrow s}$ are identical down to rank $j - 1$, so i must be available at the start of round j , since he isn't matched until round $k > j$ when he reports P_i . This also tells us that there must be at least one unclaimed seat at s at the end of round j when i reports P_i . Reporting $P_i^{s' \leftrightarrow s}$ allows i to claim this seat.
- d) If i fails to report a truly acceptable school s'' acceptable, ranking it between the lowest ranked acceptable school and \emptyset in the original report can only either not affect what i gets, or allow i to get s'' instead of \emptyset . Now, consider the lowest unacceptable school s''' that i reports acceptable. Reporting s''' unacceptable and moving all schools originally declared inferior to s''' up one position in the reported rank order can only either not affect what i gets or allow him to get something he prefers to s''' instead of s''' . \square

Lemma B.4. *Under the Boston mechanism, if the distribution of others' submitted preferences and school priorities is school-symmetric with respect to i and $q_s = q_{s'}$, $\forall s, s' \in \mathcal{S}$, then for i , truth-telling weakly first-order stochastically dominates any other preference report.*

Proof. Note that the proof of this theorem borrows heavily from the methods of Roth and Rothblum (1999). To start, Part 4 of Lemma B.3 tells us that there is no loss in limiting ourselves to reports that declare a school acceptable if and only if it is truly acceptable. Now, consider a student, i , who truthfully prefers school

		Lie: $B_i[P_i^{s' \leftrightarrow s}, P_{-i}; \succ]$		
		$= u \notin \{s, s'\}$	$= s$	$= s'$
	$= u \notin \{s, s'\}$	Case A	Case B	Impossible
Truth: $B_i[P_i, P_{-i}; \succ]$	$= s$	Impossible	Case C	Impossible
	$= s'$	Case D	Case E	Case F

(a) Table of cases

	Truth		Lie (swap s and s')	
	$B_i[P_i, P_{-i}; \succ]$	$B_i[P_i, P_{-i}^{s' \leftrightarrow s}; \succ^{s' \leftrightarrow s}]$	$B_i[P_i^{s' \leftrightarrow s}, P_{-i}; \succ]$	$B_i[P_i^{s' \leftrightarrow s}, P_{-i}^{s' \leftrightarrow s}; \succ^{s' \leftrightarrow s}]$
Case A	u	u	u	u
Case B	u	s'	s	u
Case C	s	s'	s	s'
Case D	s'	u	u	s
Case E	s'	s'	s	s
Case F	s'	s	s'	s

(b) Outcomes for the cases

Table B.7: Cases and outcomes for Lemma B.4

s' to s , that is $s' P_i s$. Table B.7a lists the possibilities for what i gets with and without switching s and s' . Parts 2 and 3 of Lemma B.3 rule out the cells labeled “Impossible”, while Part 1 tells us that u must be the same for both Truth and Lie in Case A.

Now, for each of the cases in Table B.7a, we can write out what i gets when he tells the truth and what he gets when he swaps s and s' . These are the first and third columns of Table B.7b. Looking at just these columns, however, it is not clear that i doesn't want to tell the lie: in Cases B and D, it could potentially make him better off. To show that he does better with the truth, note that Boston is label-independent, that is, if we merely change the names of two schools with the same capacity, then the output of Boston will be the same, modulo the name

change. Mathematically, if $q_s = q_{s'}$, then

$$\begin{aligned} B_i [P; \gamma] = u \notin \{s', s\} &\Leftrightarrow B_i [P^{s' \leftrightarrow s}; \gamma^{s' \leftrightarrow s}] = u \\ B_i [P; \gamma] = s' &\Leftrightarrow B_i [P^{s' \leftrightarrow s}; \gamma^{s' \leftrightarrow s}] = s \\ B_i [P; \gamma] = s &\Leftrightarrow B_i [P^{s' \leftrightarrow s}; \gamma^{s' \leftrightarrow s}] = s' \end{aligned}$$

This allows us to use the first column of Table B.7b to fill in the fourth, and the third to fill in the second. Now, by school-symmetry, the two columns under Truth in Table B.7b are equally likely if i truthfully orders s and s' , and the two columns under Lie are equally likely if i reverses their ordering. Row by row, truthfully ordering s and s' is at least as good as, and sometimes better than, switching their ordering. In other words, we can partition the states-of-the-world in such a way that, for each element of the partition, ordering s and s' truthfully weakly first-order stochastically dominates reversing their order. This implies that truthfully ordering s and s' weakly first-order stochastically dominates any other report. Hence, truth-telling weakly first-order stochastically dominates any other strategy. \square

We have thus established conditions under which Boston admits a truth-telling equilibrium.

Proposition B.1. *Under the Boston mechanism, if a school choice problem is school-symmetric, then there exists an ordinal Bayes-Nash equilibrium where all students are truth-telling.*

Proof. If all students but i are truth-telling, then their reports are school-symmetric with respect to i . By Lemma B.4, truth-telling weakly first-order stochastically dominates any other strategy for i . Hence, all students truth-telling is an ordinal Bayes-Nash equilibrium. \square

To get uniqueness, we will need to restrict ourselves to a class of strategies that is natural for symmetric environments.

Definition B.2. A strategy σ_i is **anonymous** if for any $s, s' \in \mathcal{S}$ and for any P_i , $\sigma_i (P_i^{s' \leftrightarrow s}) = (\sigma_i (P_i))^{s' \leftrightarrow s}$.

Anonymous strategies are useful to us because they translate school-symmetric true preferences into school-symmetric reported preferences, allowing us to leverage Lemma B.4.

Lemma B.5. *If the distribution of true preferences and school priorities is school-symmetric with respect to i when all students besides $i \in \mathcal{I}$ play anonymous strategies, then the distribution of students' reported preferences and school orderings is school-symmetric with respect to i .*

Proof. Strategies for the students induce a distribution over reported preferences and school priorities:

$$\tilde{\lambda}_\sigma (P_{-i}; \succ | P_i) \equiv \sum_{\{\Pi_{-i} : \sigma_{-i}(\Pi_{-i}) = P_{-i}\}} \lambda(\Pi_{-i}; \succ | P_i).$$

Using this definition, we get

$$\begin{aligned} \tilde{\lambda}_\sigma (P_{-i}^{s' \leftrightarrow s}; \succ^{s' \leftrightarrow s} | P_i) &= \sum_{\{\Pi_{-i} : \sigma_{-i}(\Pi_{-i}) = P_{-i}^{s' \leftrightarrow s}\}} \lambda(\Pi_{-i}; \succ^{s' \leftrightarrow s} | P_i) \\ &= \sum_{\{\Pi_{-i} : \sigma_{-i}(\Pi_{-i}^{s' \leftrightarrow s}) = P_{-i}^{s' \leftrightarrow s}\}} \lambda(\Pi_{-i}^{s' \leftrightarrow s}; \succ^{s' \leftrightarrow s} | P_i) \\ &= \sum_{\{\Pi_{-i} : (\sigma_{-i}(\Pi_{-i}))^{s' \leftrightarrow s} = P_{-i}^{s' \leftrightarrow s}\}} \lambda(\Pi_{-i}; \succ | P_i) \quad , \\ &= \sum_{\{\Pi_{-i} : (\sigma_{-i}(\Pi_{-i})) = P_{-i}\}} \lambda(\Pi_{-i}; \succ | P_i) \\ &= \tilde{\lambda}_\sigma (P_{-i}; \succ | P_i) \end{aligned}$$

where the second equality comes from reindexing the sum, the third comes from the definition of an anonymous strategy, the third comes from applying the $s' \leftrightarrow s$ operator to both sides of the condition in the index set, and the final inequality comes from the definition of $\tilde{\lambda}$. So, the distribution of students' reported preferences and school orderings is school-symmetric with respect to i . \square

Now, we are in a position to prove the uniqueness of the truth-telling equilibrium under Boston in the uncorrelated environment.

Proof of Proposition 2.4. Assume an equilibrium in anonymous strategies. By Lemmas B.4 and B.5, a student i in such an equilibrium can't be hurt by truth-telling. We will now proceed by ruling out strategies that can't be used in such an equilibrium.

Round 1a: Delete all strategies that fail to report some school acceptable. If student i is the top-priority student at his first choice school, then he does strictly better to rank it first.

Round 1b: Delete all strategies that don't truthfully report the first choice school. Say that i falsely reports his first choice school. Since students are playing

anonymous strategies and ranking at least one school, it is possible that all students rank the same school first, and that i is the top priority student at that school. Falsely reporting his first choice school is strictly worse than truthfully reporting.

Round 2a: Delete all strategies that fail to report two schools acceptable. Say that all students share i 's preference and he is has second priority at his second choice school. Failing to report that school acceptable is strictly worse than truth-telling.

Round 2b: Delete all strategies that don't truthfully report the second choice school. Say that i falsely reports his second choice school. Since students are playing anonymous strategies and ranking at least two schools, it is possible that all students rank the same schools first and second, and that i is the second priority student at his true second choice. Falsely reporting his second choice is strictly worse than truthfully reporting.

⋮

The logical end of this process is to eliminate all strategies except for truth-telling. Hence, the truth-telling equilibrium is the unique equilibrium in anonymous strategies for the revelation game induced by Boston in the uncorrelated environment. \square

B.4. The arts and science schools example

Proof of Proposition 5.1. Conditional on the rules not changing a student's assigned school, he does better truth-telling, state-by-state. This is true by the strategy-proofness of DA. Now, we focus on states of the world in which the rules do change the student's outcome. Assume the other students truth-tell. If our student reports himself a scientist, then with probability $\frac{1-k}{k} \cdot p_{\text{SSa}} = \frac{1-k}{k} \cdot \frac{1}{3} \cdot k \cdot (1-k) = \frac{1}{3} \cdot (1-k)^2$ the **ssa** rule is triggered and changes his utility by $u(\emptyset) - u(a)$, and with probability $p_{\text{aas}} = \frac{1}{3} \cdot (1-k)^2$ the **aas** rule is triggered and changes his utility by $u(s) - u(\emptyset)$. Since $\frac{1-k}{k} \cdot p_{\text{SSa}} = p_{\text{aas}}$, the rule allows him to net $\frac{1}{3} \cdot (1-k)^2 \cdot [u(s) - u(a)] > 0$ in utility above what he would have gotten under DA. Similarly, if he reports himself an artist, then with probability $\frac{1-k}{k} \cdot p_{\text{SSa}} = \frac{1-k}{k} \cdot \frac{1}{3} \cdot k^2 = \frac{1}{3} \cdot k \cdot (1-k)$ the **ssa** rule is triggered and changes his utility by $u(a) - u(\emptyset)$, and with probability $p_{\text{aAs}} = \frac{1}{3} \cdot k \cdot (1-k)$ the **aas** rule is triggered and his utility changes by $u(\emptyset) - u(s)$. Since $\frac{1-k}{k} \cdot p_{\text{SSa}} = p_{\text{aAs}}$, the rule allows him to net $\frac{1}{3} \cdot k \cdot (1-k) \cdot [u(a) - u(s)] < 0$ in utility above what he would have gotten under DA. So, conditional on the rules changing a student's

outcome, truthfully reporting his type first-order stochastically dominates a false report, as well as the outcome under plain DA. Hence, regardless of whether the rules change his outcome, when faced with other students truth-telling, he does first-order stochastically better to truth-tell. Thus, we have shown that truth-telling is an ordinal Bayes-Nash equilibrium and that all students interim prefer their outcomes under this equilibrium to their outcomes under the truth-telling equilibrium of plain DA. \square