



School choice: an experimental study

Yan Chen^{a,*}, Tayfun Sönmez^b

^a*School of Information, University of Michigan, 1075 Beal Avenue, Ann Arbor, MI 48109-2112, USA*

^b*Department of Economics, Koç University, Sarıyer, 34450, İstanbul, Turkey*

Received 26 October 2004; final version received 27 October 2004

Available online 28 January 2005

Abstract

We present an experimental study of three school choice mechanisms. The Boston mechanism is influential in practice, while the Gale–Shapley and Top Trading Cycles mechanisms have superior theoretical properties. Consistent with theory, this study indicates a high preference manipulation rate under Boston. As a result, efficiency under Boston is significantly lower than that of the two competing mechanisms in the designed environment. However, contrary to theory, Gale–Shapley outperforms Top Trading Cycles and generates the highest efficiency. Our results suggest that replacing the Boston mechanism with either Gale–Shapley or Top Trading Cycles mechanism might significantly improve efficiency.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: C78; C92; D82

Keywords: School choice; Experiment; Matching; Gale–Shapley; Top trading cycles; Boston mechanism

1. Introduction

School choice has been one of the most important and widely debated topics in education research and policy in the past 15 years. In the current debate on school reform, choice has moved to the top of the U.S. national agenda [40]. The most prevalent form of school choice in urban areas is open enrollment [9]. Well-known pioneer open enrollment programs, such as the intra-district choice plan in Boston and Cambridge, Massachusetts and the

* Corresponding author. Fax: +1 734 764 1555.

E-mail address: yanchen@umich.edu (Y. Chen).

inter-district choice plan in Minnesota, have been widely praised in the education literature [40]. However, there has been little research rigorously assessing the performance of these plans. Empirical work on school choice tends to focus on the effects of choice on housing prices, student performance and tax revenue (e.g., [20]), rather than on the mechanisms themselves.¹ As more states have passed legislation mandating intra- or inter-district choice, the urgency of such research is apparent.

In a recent paper, Abdulkadiroğlu and Sönmez formulate school choice as a mechanism design problem [2]. In their model, there are a number of students each of whom should be assigned a seat at one of a number of schools. Each student has a strict preference ranking of all schools. Each school, on the other hand, has a maximum capacity and a strict priority ordering of all students. The school choice model is closely related to the well-known *college admissions* model [13]. The key difference between the two models is that, in school choice, school seats are indivisible objects, which shall be assigned to students based on student preferences and school priorities, whereas in college admissions, both sides of the market are agents who have preferences over the other side and whose welfare are taken into consideration. In college admissions, both sides of the market are strategic actors who can misrepresent their preferences, whereas in school choice, only students are strategic actors and school priorities are exogenously determined by the school district based on state/local laws (and/or education policies). Nevertheless, based on the isomorphism between the two models, one can treat school priorities as school preferences and therefore, many concepts/findings in college admissions have their counterparts in school choice. Stability is one such concept. A matching (i.e. the outcome of both models) is *unstable* in the context of college admissions, if there exists a student–college pair (i, s) , where student i prefers college s to her assignment and college s prefers student i to at least one of its assignments. The counterpart of this situation in the context of school choice involves a student–school pair (i, s) , where student i prefers school s to her assignment, although she has higher priority than another student who is assigned a seat at school s . Therefore, stability in the context of college admissions corresponds to elimination of *justified-envy* [3] in the context of school choice.

It is well-known that not only does there exist a stable matching in the context of college admissions, but indeed there exists a stable matching that is weakly preferred to any stable matching by each student [13]. Moreover, truthful preference revelation is a dominant strategy for each student under the mechanism which selects this *student-optimal stable matching* [10,31]. Since students are the only strategic agents in the context of school choice, the corresponding mechanism, namely the *Gale–Shapley student-optimal stable mechanism* (*GS* hereafter), is strategy-proof in this setup. Moreover, it eliminates justified envy. However, it has a caveat in the context of school choice. Although it is Pareto efficient when the welfare of colleges are taken into consideration, its outcome may be Pareto inefficient when only the welfare of students are considered.² As a remedy, Abdulkadiroğlu and

¹ Another line of discussion has been on private school vouchers. See [19,11,38,27].

² This is because, while the student-optimal stable matching is weakly preferred to any stable matching by each student in the context of college admissions, there can be an unstable matching that is weakly preferred to the student-optimal stable matching by each student [31].

Sönmez [2] propose an extension of the *Top Trading Cycles mechanism* (hereafter *TTC*) which plays a central role in *housing markets* [41]. *TTC* is also strategy-proof [2,30,32] and Pareto-efficient, but it does not eliminate justified-envy.

Until recently, neither of these two theoretically appealing mechanisms had been in use at any school district. In contrast the *Boston mechanism* has been very influential in practice. School districts which uses this mechanism and its variants include Boston, Cambridge, Denver, Minneapolis, Seattle and St. Petersburg-Tampa.

In the Boston mechanism, students submit preference rankings, which are then used in a series of rounds to assign students to schools, following the priorities of schools within each round but otherwise giving priority to students according to their stated ranking of schools. More precisely, in the first round, only the first choices of the students are considered, and they are sorted by their priorities at their first choice schools. Assignments are then made until each school is full, or until all first choices are accommodated, whichever comes first. In round two, the second choices of students still unassigned are considered, and they are now sorted by their priorities at their second choice schools. Assignments are then made for remaining slots (if any) until each school is full or until all second choices are accommodated. This process is repeated with the third choices and so on, until all students have been assigned a seat.

While it has been influential in practice, Abdulkadiroğlu and Sönmez point out that the Boston mechanism has the following deficiencies [2]:

- it does not eliminate justified envy and hence may produce outcomes that may result in lawsuits from dissatisfied parents;
- it is not efficient; and
- it gives students and their parents strong incentives to misrepresent their preferences by improving the ranking of schools for which they have a high priority.³

In fact, such preference misrepresentation is advocated by the Central Placement and Assessment Center (CPAC) in Minneapolis:

The Minneapolis algorithm places a very high weight on the first choice, with second and third choices being strictly backup options. This is reflected in the advice CPAC gives out to parents, which is to make the first choice a true favorite and the other two “realistic,” that is, strategic choices [14].

Similarly, preference misrepresentation is often advocated in the local press for school districts which rely on the Boston mechanism and its variants. Consider the following statement from Mas in the *Seattle Press* [24]:

The method the school district uses to sort the school choice requests gives first priority to students who are already enrolled at that school. Next in line come those students with siblings at the school. Both of these factors are beyond your control. These students are sure things. Enrollment at these schools is theirs for the asking. No amount of

³ The Boston mechanism is a special case of a priority matching mechanism. Roth shows that these mechanisms are vulnerable to preference manipulation in general [34]. Collins and Krishna report similar incentives under an on-campus housing mechanism used at Harvard for the period 1977–1989 [6].

strategizing, short of polling all existing students to determine how many have younger siblings about to enter the school, can help you here. Third in line, and the first effect of any real choice, are those students who live in the school's reference area. This is why you have such an excellent chance of getting into your reference school if you make it your top choice. Choosing another neighborhood's reference school, however, puts a lot of kids in line ahead of yours. That reduces your chances of getting in, particularly if the school has small classes.

Although both the CPAC and the *Seattle Press* advocate misrepresentation, the type of misrepresentation advocated by each is different. Finally, in St. Petersburg, Tobin (*St. Petersburg Times*, [43]) gives the parents similar advice as that in the *Seattle Press*:

Make a realistic, informed selection on the school you list as your first choice. It's the cleanest shot you will get at a school, but if you aim too high you might miss.

Here's why: If the random computer selection rejects your first choice, your chances of getting your second choice school are greatly diminished. That's because you then fall in line behind everyone who wanted your second choice school as their first choice. You can fall even farther back in line as you get bumped down to your third, fourth and fifth choices.

In spite of the strong incentives and the evidence for preference misrepresentation by the Boston mechanism, practitioners in education have relied on stated preferences when evaluating this mechanism. For example, Cookson indicates the following results [8]:

91% of all students entering the Cambridge public school system at the K-8 levels have gained admission to a school of their choice, 75% to the school of their first choice, and 16% to either their second or third choice.

Given the incentives under the Boston mechanism, field data that rely on stated preferences cannot adequately assess the effectiveness of the Boston mechanism. In this paper, we turn to controlled laboratory experiments to compare the performance of the Boston mechanism with the two theoretically superior mechanisms, GS and TTC. As Roth emphasizes [35], experimental economics is one of the basic tools of market design. Laboratory experiments make it possible to evaluate allocations according to induced, true preferences, and to assess the performance of alternative school choice mechanisms that have not yet been used in the field.

Overall, GS and TTC have three advantages over the Boston mechanism. First, unlike the Boston mechanism, truthful preference revelation is a dominant strategy under either mechanism. Second, TTC mechanism is Pareto efficient whereas the GS mechanism is constrained-efficient among mechanisms that eliminate justified-envy. Finally, the GS mechanism eliminates justified envy. Despite their theoretical benefits, it is important to evaluate their performance in the laboratory before introducing them into the field. In particular, we are interested in testing

- (1) the extent of preference manipulation and the resulting efficiency loss under the Boston mechanism;

- (2) the extent to which subjects recognize and use their dominant strategies without prompting under the two competing mechanisms; and
- (3) the robustness of theoretical efficiency comparisons when the mechanisms are implemented among boundedly rational subjects and across different environments.

The school choice problem belongs to the general class of matching problems. While we are not aware of any previous experimental studies of school choice mechanisms, there have been experimental studies of other one- and two-sided matching problems, motivated by various real-world applications. These include Olson and Porter [29], Nalbantian and Schotter [26], Harrison and McCabe [16], Kagel and Roth [22], Ünver [44], Haruvy et al. [17], Chen and Sönmez [5], and McKinney et al. [25]. Our experiment differs from the above experiments in both the particular mechanisms compared and the potential field applications of the mechanisms.

The paper is organized as follows. Section 2 introduces the school choice problem. Section 3 reviews the theoretical properties of the three mechanisms. Section 4 contains a description of the experimental design. Section 5 summarizes the main results of the experiments. Section 6 concludes the paper.

2. The school choice problem

In a *school choice problem* [2], there are a number of students, each of whom should be assigned a seat at one of a number of schools. In addition, each student has strict preferences over all schools. Furthermore, each school has a maximum capacity and a strict priority ordering of all students. Here, priorities do not represent school preferences. They are imposed by the school district based on state and local laws. For example, in several school districts, students who already have siblings at a school are given priority at that school. Similarly, in most school districts, students are given priority at their district schools. The fact that the law mandates school priorities also means that they are publicly known. However, schools might have discretion in the interpretation of the law, which might be interesting to study for future work.

As we have emphasized in the Introduction, the school choice problem is closely related to the well-known *college admissions problem* [13]. The college admissions problem has been well-analyzed and successfully applied in both the U.S. and British entry-level labor markets (see [33,34]). The main difference between the two models is that in college admissions, schools are strategic agents who have preferences over students, whereas in school choice, schools are “objects” to be consumed by the students.

The outcome of a school choice problem is referred to as a *matching*, and it is an assignment of school seats to students such that each student is assigned one seat and no more seats than its capacity are filled for any school. A matching is *Pareto efficient* if there is no matching which assigns each student a weakly better school and at least one student a strictly better school. A matching μ eliminates *justified envy* if there is no unmatched student–school pair (i, s) such that:

- student i prefers school s to her assignment under μ and
- student i has higher priority at school s than some other student who is assigned a seat at school s under μ .

Elimination of justified envy in the present context corresponds to the well-known *pair-wise stability* condition in the context of college admissions.

Finally, a *student assignment mechanism* is a systematic procedure that selects a matching for each school choice problem. A mechanism is Pareto efficient if it always selects a Pareto efficient matching; it eliminates justified envy if it always selects a matching which eliminates justified envy; and it is *strategy-proof* (i.e., *dominant strategy incentive-compatible*) if no student can possibly benefit by unilaterally misrepresenting her preferences.

3. Theoretical properties of the mechanisms

In this section, we introduce the three mechanisms and their theoretical properties. We also informally state our hypotheses.

3.1. Boston mechanism

One of the most widely used student assignment mechanisms is the *Boston mechanism* (BOS) which works as follows:

- (1) For each school, a priority ordering of students is determined based on state and local laws/policies.⁴
- (2) Each student submits a preference ranking of the schools.
- (3) The final phase is student assignment based on submitted preferences and priorities. The outcome is obtained in several rounds.

Round 1: In Round 1, only the first choices of the students are considered. For each school, consider the students who have listed it as their first choice and assign seats of the school to these students one at a time following their priority order until either there are no seats left or there is no student left who has listed it as her first choice.

In general at

Round k: Consider the remaining students. In Round k only the k th choices of students are considered. For each school with available seats, consider the students who have listed it as their k th choice and assign the remaining seats to these students one at a time following their priority order until either there are no seats left or there is no student left who has listed it as her k th choice.

The major difficulty with the Boston student assignment mechanism is that it is not strategy-proof; that is, students may benefit from misrepresenting their preferences. In particular, Abdulkadiroğlu and Sönmez point out that the Boston mechanism gives students

⁴ For example at Boston the priority ordering is determined based on the following hierarchy:

- First priority: sibling and walk zone.
- Second priority: sibling.
- Third priority: walk zone.
- Fourth priority: others

Students in the same priority group are ordered based on a previously announced lottery. In the experimental design we consider the simpler case where the priority is based on “walk zone” together with a random lottery.

and their parents a strong incentive to misrepresent preferences by improving the ranking of schools which they have a high priority [2].⁵ As we have indicated in the introduction, such preference manipulation is even advocated by the central agencies in charge of school choice programs as well as by the local press.

Despite this widespread preference manipulation, practitioners in education have been using stated preferences to evaluate the Boston mechanism. In addition to Cookson's evaluation of the Cambridge public school system [8], cited in the introduction, Glenn also uses stated preferences to evaluate the choice program of the Massachusetts public school system [15]:

A majority of the students are accepted into their first choice schools. For example, in 1991, 74% of sixth graders were assigned to their first choice school, 10% to their second choice, and 15% to schools they did not select.

Furthermore, school districts that rely on the Boston mechanism often emphasize in their marketing efforts that they are able to accommodate most students' top choices. Based on these field results and the incentives under the Boston mechanism, we expect that under the Boston mechanism a significant proportion of students will misrepresent their preferences. Furthermore, we expect that most participants will be assigned their stated top choices in the experiment.

In our experimental design each participant has high priority at her district school and low priority at other schools. Based on this priority structure in our experimental design, we expect that, under the Boston mechanism, a significant proportion of students will improve the ranking of their district schools.

Recall that Pareto efficiency refers only to student preferences in our school choice model. The outcome of the Boston mechanism is Pareto efficient provided that students submit their true preferences. However, given the strong incentive to misrepresent preferences, a substantial efficiency loss can be expected. Therefore, under the Boston mechanism, we expect that efficiency will be lower than that under the two competing mechanisms.

3.2. Gale–Shapley mechanism

In the context of college admissions, the counterpart of the elimination of justified envy is pairwise stability. The GS mechanism plays a key role in the pairwise stability literature⁶ and it is adapted to the school choice setting as follows:

- (1) For each school, a priority ordering of students is determined based on state and local laws/policies.
- (2) Each student submits a preference ranking of the schools.

⁵ While students are exogenously priority ordered at each school, the “effective priorities” are endogenous under the Boston mechanism in the sense that each student who ranks a school as her k th choice is considered before each student who ranks it $(k + 1)$ th for each k . The exogenous priorities are only utilized to tie-break among students who have ranked a school at the same rank order. It is this ability of students to influence the effective priorities that results in the vulnerability of the Boston mechanism to preference manipulation.

⁶ See, for example, [36].

(3) The final and key phase is student assignment based on submitted preferences and priorities. The outcome is obtained at the end of several steps.

Step 1: Each student proposes to her first choice. Each school rejects the lowest priority students in excess of its capacity and keeps the remaining students on hold.

In general, at

Step k: Each student who has been rejected in the previous step proposes to her next choice. Each school considers the students it has been holding together with its new proposers; it rejects the lowest priority students in excess of its capacity and keeps the remaining students on hold.

The algorithm terminates when no student proposal is rejected and each student is assigned a seat at her final tentative assignment.

The Gale–Shapley mechanism is strategy-proof. Therefore, we expect that, under the Gale–Shapley mechanism, participants will reveal their preferences truthfully.

Moreover, the Gale–Shapley mechanism eliminates justified envy, and is thus legally appealing. Re-interpreting an example from [31], Abdulkadiroğlu and Sönmez show that Pareto efficiency is incompatible with the elimination of justified envy [2]. Therefore, elimination of justified envy has an efficiency cost and hence the Gale–Shapley mechanism is not Pareto efficient. Nevertheless, it Pareto dominates any other mechanism that eliminates justified envy.

3.3. Top trading cycles mechanism

In the trade-off between elimination of justified envy and Pareto efficiency, the Gale–Shapley mechanism gives up Pareto efficiency. The *top trading cycles mechanism*, on the other hand, gives up elimination of justified envy, but is Pareto efficient.

While the description of the TTC mechanism is somewhat involved in general, it takes a relatively simple form for the priority structure we consider in our experimental design. Suppose that:

- each student has a high priority at her district school and low priority at other schools and
- at each school, the priority among high-priority students, as well as the priority among low-priority students, is determined with a single tie-breaking lottery.

Under this special case, the TTC mechanism works as follows:

- (1) For each school, a priority ordering of students is determined.
- (2) Each student submits a preference ranking of the schools.
- (3) Based on the submitted preferences and priorities, the student assignment is determined as follows:
 - (a) Each participant is tentatively assigned a seat at her district school.
 - (b) All participants are lined up in an initial queue based on the tie-breaking lottery.
 - (c) An application to the highest-ranked school is made on behalf of the participant at the top of the queue.
 - If the application is made to her district school, then her tentative assignment is finalized. The participant and her assignment are removed from the system. The process continues with the next participant in line.

- If the application is made to another school, say school S , then the first participant in the queue who tentatively holds a seat at school S is moved to the top of the queue directly in front of the requester.
- (d) Whenever the queue is modified, the process continues similarly: an application to the highest ranked school with still available seats is made on behalf of the participant at the top of the queue.
 - If the application is made to her district school, then her tentative assignment is finalized. The participant and her assignment are removed from the system. The process continues with the next participant in line.
 - If the application is made to another school, say school S , then the first participant in the queue who tentatively holds a seat at school S is moved to the top of the queue directly in front of the requester.
- (e) A mutually beneficial exchange is obtained when a cycle of applications is made in sequence (e.g., I apply to John's district school, John applies to your district school and you apply to my district school). In this case, the Pareto-improving exchange is carried out and the participants, as well as their assignments, are removed from the system.
- (f) The process continues until all participants are assigned a school seat.

The TTC mechanism is not only Pareto efficient, but also strategy-proof. Therefore, we expect that, under the TTC mechanism, participants will reveal their preferences truthfully.

In theory, the TTC mechanism has an efficiency advantage over the Boston mechanism as well as the Gale–Shapley mechanism: The outcome of the Boston mechanism is Pareto efficient provided that participants reveal their preferences truthfully. So any efficiency loss in the Boston mechanism is a consequence of preference manipulation. The Gale–Shapley mechanism, on the other hand, is strategy-proof, but elimination of justified envy and Pareto efficiency are not compatible. Since Gale–Shapley mechanism Pareto dominates any other mechanism that eliminates justified envy, any efficiency loss in the Gale–Shapley mechanism is a consequence of this incompatibility. Based on theory, we expect that the TTC mechanism is more efficient than both the Boston mechanism and the Gale–Shapley mechanism.

Nonetheless, contrary to theory, there is room for efficiency loss in the TTC mechanism and additional efficiency loss in the Gale–Shapley mechanism if agents misrepresent their preferences. Chen and Sönmez find that about a third of the participants misrepresent their preferences under a variant of the TTC mechanism in an experiment about on-campus housing [5]. Therefore, at least in that context, the strategy-proofness property of the TTC mechanism is not completely transparent. Thus, it is important to investigate whether the appealing theoretical properties of the Gale–Shapley mechanism and the TTC mechanism can be realized when they are implemented among real people with real incentives.

4. Experimental design

We design our experiment to compare the performance of the Boston, Gale–Shapley and TTC mechanisms, with particular attention to the questions of truthful preference revelation

and efficiency comparisons. The environment is designed to capture the key aspects of the school choice problem and to simulate the complexity inherent in potential applications of the mechanisms. We implement a 3×2 design: for each of the three mechanisms, we examine it in a designed environment as well as a random environment to test for robustness.

To simulate the complexity inherent in real-world applications, we use relatively large sessions. For all treatments, in each session, there are 36 students and 36 school slots across seven schools.⁷ These schools differ in size, geographic location, specialty and quality of instruction in each specialty. Each school slot is allocated to one participant. There are three slots each at schools A and B, and six slots each at schools C, D, E, F and G. Students 1–3 live within the school district (or walk zone) of school A, students 4–6 live within the school district of school B, students 7–12 live within the school district of school C, students 13–18 live within the school district of school D, students 19–24 live within the school district of school E, students 25–30 live within the school district of school F and students 31–36 live within the school district of school G.

4.1. Designed environment

In real life, student preferences are influenced by the proximity of the school and the school quality. In the designed environment, we construct a “realistic” environment by correlating student preferences with both school proximity and quality. In this environment, schools A and B are higher quality schools, while C–G are lower quality schools. A is stronger in Arts and B is stronger in Sciences: A is first tier in Arts and second tier in Sciences, whereas B is second tier in Arts and first tier in Sciences. Lower quality schools C–G are third tier in Arts as well as in Sciences. There are two types of students: odd-labelled students are gifted in Sciences, while even-labelled students are gifted in Arts.

Each student’s ranking of the schools is generated by a utility function, which depends on the school’s quality, proximity of the school and a random factor. The utility function of each student has three components:

$$u^i(S) = u_p^i(S) + u_q^i(S) + u_r^i(S),$$

where

- $u_p^i(S)$ represents the proximity utility for student i at school S . This utility is 10 if student i lives within the walk zone of school S and zero otherwise.
- $u_q^i(S)$ represents the quality utility for student i at school S . For odd-labelled students (i.e., for students who are gifted in sciences), $u_q^i(A) = 20$, $u_q^i(B) = 40$, and $u_q^i(S) = 10$ for $S \in \{C, D, E, F, G\}$. For even-labelled students (i.e., for students who are gifted in arts), $u_q^i(A) = 40$, $u_q^i(B) = 20$, and $u_q^i(S) = 10$ for $S \in \{C, D, E, F, G\}$.
- $u_r^i(S)$ represents a random utility (uniform in the range 0–40) which captures diversity in tastes.

⁷ While 36 per session is relatively large in the experimental literature, it is still far from the relevant population size in real-world school choice problems. A further increase in size should not affect the incentive constraints. Nor do we expect computational problems from such increase, based on our simulation analysis.

Based on the resulting ranking, the monetary payoff to each student is \$16 for her first choice, \$13 for her second choice, \$11 for her third choice, \$9 for her fourth choice, \$7 for her fifth choice, \$5 for her sixth choice and \$2 for her last choice. The payoffs between different outcomes are sufficiently dispersed so that there is a monetarily salient difference (\$14) between getting one's top choice and getting one's last choice.

Table 1 presents the monetary payoff (i.e., the induced preferences over schools, [42]) for each participant as a result of the type of school she holds at the end of the experiment. Boldfaced numbers indicate that the participant lives within the school district of that school. For example, participant #1 lives within the school district of school A. She will be paid \$13 if she holds a slot at school A at the end of the experiment, \$16 if she holds a slot at school B, \$9 if she holds a slot at school C, etc.

4.2. *Random environment*

To check the robustness of the results with respect to changes in the environment, we also implement a random environment, as presented in Table 2.

In the random environment, for each participant the payoff for attending each school is a distinct integer in the range 1–16, chosen without replacement. In the case where the district school has the highest payoff, the entire row is discarded for the student and a new payoff vector is constructed. Otherwise, the decision for the participant would be trivial under all three mechanisms.

All three mechanisms are implemented as one-shot games of incomplete information. Each subject knows his own payoff table, but not the other participants' payoff tables. He also knows that different participants might have different payoff tables. This information condition is a good approximation of reality. Following Chen and Sönmez, we use one-shot games to evaluate the mechanisms [5], as in real-world applications, the mechanisms will likely be used in a one-shot setting. Without any practice rounds or opportunities to learn over time, one-shot implementation presents the most realistic and (theoretically) cleanest test for the mechanisms. From the analysis perspective, the advantage of one-shot games is that each observation is an independent observation; therefore, we can run relatively fewer sessions for each treatment and use the recombinant estimation technique [23] to get an efficient estimation of summary statistics. While learning and experience may matter in practice, we leave them as interesting extensions for future work.

Table 3 summarizes features of experimental sessions. For each mechanism in each environment, we conducted by hand two independent sessions between September and November 2002 at the University of Michigan. The subjects are students at the University of Michigan. This gives us a total of 12 independent sessions and 432 subjects.⁸ Each session consists of one round only. The sessions last approximately 45 minutes, with the first 20–25 minutes being used for instructions. The conversion rate is \$1 for all sessions. Each subject also receives a participation fee of \$3 in addition to her earnings from the experiment. The average earning (including participation fee) is \$15.23.

⁸ Despite our explicit announcement in the advertisement that subjects could not participate in this set of experiments more than once and our screening before each session, two subjects participated twice.

Table 1
Payoff table in the designed environment

Student ID	Schools						
	A	B	C	D	E	F	G
1	13	16	9	2	5	11	7
2	16	13	11	7	2	5	9
3	11	13	7	16	2	9	5
4	16	13	11	5	2	7	9
5	11	16	2	5	13	7	9
6	16	13	7	9	11	2	5
7	13	16	9	5	11	7	2
8	16	9	11	2	13	7	5
9	16	13	2	5	9	7	11
10	16	7	9	5	2	11	13
11	7	16	11	9	5	2	13
12	13	16	9	11	2	7	5
13	9	16	2	13	11	5	7
14	16	5	2	9	7	13	11
15	13	16	9	11	2	7	5
16	16	13	11	5	9	7	2
17	13	16	5	7	2	9	11
18	16	13	5	9	7	11	2
19	11	16	7	5	13	9	2
20	16	13	7	9	5	2	11
21	13	16	2	7	9	11	5
22	16	11	7	2	9	5	13
23	16	13	7	2	5	11	9
24	16	13	11	5	9	2	7
25	13	16	2	5	11	9	7
26	16	13	5	9	7	2	11
27	7	11	5	2	13	9	16
28	16	13	7	2	11	5	9
29	7	11	16	13	2	9	5
30	16	9	7	2	5	11	13
31	11	16	7	2	5	9	13
32	13	9	16	2	5	7	11
33	13	16	11	9	7	5	2
34	16	11	2	7	5	13	9
35	7	16	2	5	11	13	9
36	16	13	5	7	9	2	11

In the experiment, each subject is randomly assigned an ID number and is seated in a chair in a classroom. The experimenter reads the instructions aloud. Subjects ask questions, which are answered in public. Subjects are then given 15 minutes to read the instructions again at their own pace and to make their decisions. At the end of 15 minutes, the experimenter collects the decisions and asks a volunteer to draw ping-pong balls out of an urn, which

Table 2
Payoff table in the random environment

Student ID	Schools						
	A	B	C	D	E	F	G
1	14	8	10	13	16	12	3
2	7	15	1	6	14	3	4
3	4	10	5	1	12	8	15
4	8	7	14	9	4	11	1
5	11	9	14	7	12	5	4
6	4	11	5	9	3	16	7
7	14	10	8	15	11	6	5
8	6	9	12	5	14	10	8
9	12	10	13	16	9	15	3
10	16	5	13	12	3	1	4
11	5	11	8	2	16	10	7
12	9	6	7	4	10	13	11
13	4	7	13	11	8	10	1
14	10	1	7	5	14	13	16
15	16	13	5	9	8	3	6
16	7	12	5	8	15	9	4
17	14	11	3	4	10	6	8
18	10	1	6	11	15	2	8
19	8	12	16	5	14	4	13
20	15	4	1	2	11	14	3
21	3	16	8	6	7	10	2
22	1	8	14	15	5	3	4
23	11	16	12	1	3	7	15
24	12	11	6	3	9	4	14
25	8	14	7	15	1	5	12
26	11	16	9	7	4	12	1
27	7	13	15	14	3	10	12
28	4	15	10	11	9	6	7
29	11	12	7	14	6	10	9
30	13	8	3	12	16	2	11
31	6	3	10	14	11	16	1
32	13	12	7	11	2	16	14
33	2	5	7	8	15	10	6
34	8	10	3	14	16	1	12
35	10	12	2	7	3	14	8
36	16	3	14	13	8	10	11

generates the lottery. Because of the complexity involved in computing the results by hand, the experimental sessions end after the lottery is generated and publicly announced. The experimenter then puts the subject decisions and the lottery into a computer to generate the allocations, announces the allocations by email and pays the subjects after the email announcement.

Table 3
Features of experimental sessions

Treatment	Mechanism	Environment	Subjects per session	Total # of subjects
BOS _d	Boston	Designed	36	72
GS _d	Gale–Shapley	Designed	36	72
TTC _d	Top Trading Cycle	Designed	36	72
BOS _r	Boston	Random	36	72
GS _r	Gale–Shapley	Random	36	72
TTC _r	Top Trading Cycle	Random	36	72

Experimental instructions and data can be found at the *Journal of Economic Theory* web site.

5. Results

Three questions are important in evaluating the mechanisms. The first is whether individuals report their preferences truthfully. The second is whether one mechanism is more efficient given the reported preferences. The third is whether the experimental results are robust to changes in the environment.

We first examine whether individuals reveal their preferences truthfully, and if not, how they manipulate their preferences under each of the three mechanisms. In presenting the results, we introduce some shorthand notations. Let $x > y$ denote that a measure under mechanism x is greater than the corresponding measure under mechanism y at the 5% significance level or less. Let $x \sim y$ denote that a measure under mechanism x is not significantly different from the corresponding measure under mechanism y at the 5% significance level.

Table 4 presents the proportion of truthful preference revelations, as well as various patterns of misrepresentation for each treatment. Note that, under the Boston mechanism, truthful preference revelation requires that the entire reported ranking is identical to the agent's true preference ranking. The only exception is when an agent's district school is her top choice. In this case, truthful preference revelation entails stating the top choice. Under GS or TTC, however, truthful preference revelation requires that the reported ranking is identical to the true preference ranking from the first choice up to the participant's district school. The remaining rankings, from the district school to the last choice, are irrelevant under both GS and TTC.

Result 1 (*Truthful preference revelation*). *In both the designed and random environments, the proportion of truthful preference revelation under BOS is significantly lower than that under either GS or TTC. The proportion of truthful preference revelation under TTC is significantly lower than that of GS in the designed environment, and is weakly lower in the random environment.*

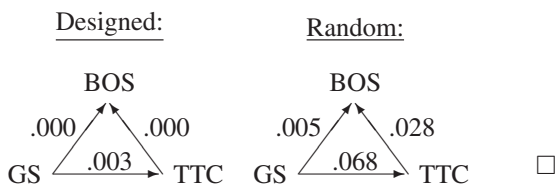
Support. The second column of Table 4 presents the proportions of truthful preference revelation for each mechanism. P -values for pairwise one-sided t -tests of proportions are

Table 4
Proportions of truthful preference revelation and misrepresentation

Treatment	Truth	Misrepresentation of preferences							
		1. DSB	2. SSB	3. SPB	1 & 2	1 & 3	2 & 3	1 & 2 & 3	Other
BOS _d	0.139	0.153	0.083	0.000	0.153	0.014	0.028	0.431	0.000
GS _d	0.722	0.014	0.014	0.014	0.000	0.000	0.153	0.069	0.014
TTC _d	0.500	0.014	0.014	0.000	0.042	0.000	0.167	0.236	0.028
BOS _r	0.278	0.083	0.042	0.042	0.097	0.208	0.028	0.208	0.014
GS _r	0.556	0.028	0.069	0.069	0.028	0.042	0.125	0.042	0.042
TTC _r	0.431	0.056	0.028	0.083	0.028	0.139	0.111	0.097	0.028

Notes: (1) DSB = district school bias; (2) SSB = small school bias; (3) SPB = similar preference bias.

as follows:



Result 1 has important efficiency implications, which will be explored later in this section.

In the education literature, the performance of the Boston mechanism is evaluated through the proportion of students who get their reported top choices. Cookson writes that [8], in the Cambridge public school system, 75% of the students got “into the school of their first choice”. We compare the proportion of participants receiving their reported top choices vs. the proportion who actually receive their true top choices. We find a comparable proportion of students get into their reported top choices as in the field data. However, only 40% of the students who receive their reported top choices also receive their true top choices.

Result 2 (*Top choices under BOS*). Under the Boston mechanism, 70.8% of the participants receive their reported top choices, while only 28.5% of the participants receive their true top choices. The difference is statistically highly significant.

Support. In the designed environment, 54.2% of the participants receive their reported top choices, while 11.1% receive their true top choices. A *t*-test of proportions with the null hypothesis of equal proportion is rejected in favor of H_1 : reported top choices% > true top choices%, yielding $z = 5.509$ (p -value < 0.001). In the random environment, 87.5% of the participants receive their reported top choices, while 45.8% receive their true top choices. The same *t*-test of proportions with the null hypothesis of equal proportion is rejected in favor of H_1 : reported top choices% > true top choices%, yielding $z = 5.303$ (p -value < 0.001). □

Results 1 and 2 indicate that, under the Boston mechanism a significant proportion of participants misrepresent their preferences and most participants are assigned their reported top choices. However, with the induced value method, we uncover the fact that only 28.5% of the participants receive their true top choices.⁹ It is important to emphasize that more subjects receive their true top choices in the random environment. This is due to the fact that preferences are not correlated in the random environment, thus reducing competition and increasing a student's chance of receiving her top choice.

However, even though both GS and TTC are strategy-proof, we have some misrepresentation of preferences. Since roughly 80% of the participants under Boston, 53% under TTC and 36% under Gale–Shapley misrepresent their preferences, it is important to analyze who manipulates preferences in each mechanism and how.

Table 5 reports logit models of truthful preference revelation under each of the three mechanisms. In each of three specifications, the dependent variable is truth-telling, which equals one if a participant reports her preference ordering truthfully and zero otherwise. The independent variables are: district school position dummies, D_{dsi} , where $i = 1, \dots, 6$, which equals one if a participant's district school is her i th choice and zero otherwise; small school dummy, D_{ab1} , which equals one if schools A or B is a participant's first choice and zero otherwise; and a constant. We omit the dummy D_{ds7} . Therefore, the estimated coefficient of D_{dsi} captures the difference between district school positions i and 7 on the likelihood of a participant reporting her true preferences. The coefficient of dummy D_{ab1} captures the effect of having a small school as a top choice on the same likelihood. When including D_{abi} , where $i = 2, \dots, 6$, we cannot reject the hypothesis that the coefficients for these dummies are jointly zero,¹⁰ therefore, we do not include them in our models.

Result 3 (*Effects of district school and small school*). Under BOS, having a small school (A or B) as the top choice significantly reduces the likelihood of a participant reporting her true preference, while under GS or TTC, having a small school as the top choice has no significant effect on truthful preference revelation. Under all three mechanisms, having a district school as the top choice predicts truthful preference revelation perfectly.

Support. Table 5 reports logit estimations of the effects of small school and district school rankings on truthful preference revelation under each of the three mechanisms. □

Result 3 indicates that strategic behavior under the Boston mechanism is most likely to be found when true first choices are capacity constrained. This result has important practical implications, e.g., if one is going to continue to use the Boston mechanism, it is worth knowing when it is better and when it is worse. Meanwhile, we do not observe this effect under either GS or TTC. This result indicates that most participants understand the basic incentives of each mechanism: since the Boston mechanism puts a lot of weight

⁹ In comparison, under GS, 51.4% of the participants receive their reported top choices, while 35.4% receive their true top choices. The corresponding statistics under TTC are 56.3% and 31.3%.

¹⁰ In each of the three specifications, $\chi^2(5)$ tests of the null hypotheses that coefficients of $D_{ab2}–D_{ab6}$ are jointly zero yield p -values of 0.9099 (BOS), 0.8325 (GS) and 0.5447 (TTC), respectively.

Table 5
Logit regression: probability of truthful preference revelation

	Dependent Variable: Truth-telling		
	(1) BOS	(2) GS	(3) TTC
D_{ds2}	-0.075 (1.234)	1.530 (0.849)*	-0.273 (0.742)
D_{ds3}	1.228 (1.177)	-0.111 (0.775)	-0.915 (0.771)
D_{ds4}	1.004 (1.132)	-0.033 (0.722)	-0.497 (0.713)
D_{ds5}	0.173 (1.331)	0.544 (0.913)	-1.105 (0.904)
D_{ds6}	0.966 (1.216)	-0.886 (0.800)	-1.023 (0.799)
D_{ab1}	-1.117 (0.471)**	0.589 (0.392)	-0.014 (0.361)
Constant	-1.6491 (1.084)	0.0599 (0.690)	0.0599 (0.681)
Observations	140	140	140

Notes: (1) Standard errors are in parentheses. (2) D_{dsi} is a dummy variable which equals one if the district school is the i th choice, and zero otherwise. (3) D_{ab1} is a dummy variable which equals one if school A or B is the top choice, and zero otherwise. (4) In all three specifications, D_{ds1} predicts success perfectly and is hence not reported. (5) Significant at: * 10% level; ** 5% level.

on the top choice and a small school is competitive, we expect more manipulation from those whose top choice is a small school. By contrast, under either GS or TTC, a student is guaranteed a school at least as desirable as her district school; therefore, when a district school is a student's top choice, it is easier for subjects to figure out that truthful preference revelation is a dominant strategy, and hence, we expect less manipulation from these subjects.

We now analyze the manipulation patterns under each mechanism. Columns 3–8 in Table 4 represent eight exhaustive and mutually exclusive categories of misrepresentation under each treatment.

- *District school bias* (1): A participant puts her district school into a higher position than that in the true preference order. This category contains observations which exhibit only district school bias.
- *Small school bias* (2): A participant puts school A or B (or both) into lower positions than those in the true preference ordering. This type of manipulation might be caused by the fact that schools A and B have fewer slots than schools C–G. This category contains observations which exhibit only small school bias.
- *Similar preference bias* (3): A participant puts schools with the highest payoffs into lower positions. This might be due to the belief that participants have similar preferences, therefore, my most desirable school might also be desired by others. This category contains observations which exhibit only similar preference bias.

- 1 and 2: District school bias often implies small school bias. When both are present, we put them in a combined category. Note that, in the designed environment, as A and B are ranked highly by all agents, upgrading the district school in most cases implies a downgrading of school A or B.
- 1 and 3: As district school bias is sometimes also correlated with similar preference bias, we put observations that exhibit both types of manipulation into a combined category.
- 2 and 3: When school A or B is an agent's top choice, small school bias is perfectly correlated with similar preference bias. Therefore, we put them into a combined category.
- 1, 2 and 3: When school A or B is an agent's top choice, and an agent upgrades her district school to her top choice, this type of manipulation exhibits all three biases.
- *Other*: This category contains manipulations which cannot be put into the above categories, possibly due to confusion, or misunderstanding of the allocation mechanisms.

The manipulation analysis serves three purposes. First, for the Boston mechanism, the distribution of various types of manipulations help us understand when the mechanism performs well or poorly. Second, for GS and TTC, if parents do not understand the mechanisms, there can be potential efficiency loss. Therefore, manipulation data can be used to construct educational brochures about a new mechanism. They can also be used to provide an estimate of the value of education or the potential efficiency loss without education. Lastly, manipulation analysis for all three mechanisms will be used later in this section to simulate the operation of the mechanisms in intermediate settings.

Result 4 (*District school bias*). *A significant proportion (two-thirds) of the subjects under BOS misrepresent their preferences using district school bias, which is significantly higher than the proportion of misrepresentation under either GS or TTC.*

Support. Columns 3, 6, 7 and 9 in Table 4 present the proportions of misrepresentation related to District School Bias (1, 1 & 2, 1 & 3 and 1 & 2 & 3, respectively) under each treatment. Pooling all four categories together and using the one-sided proportion of t -tests, we get

- (1) $BOS_d > 0$ at p -value < 0.001 ; $BOS_r > 0$ at p -value < 0.001 ;
- (2) $BOS_d > GS_d$ at p -value < 0.001 ; $BOS_d > TTC_d$ at p -value < 0.001 ; and
- (3) $BOS_r > GS_r$ at p -value < 0.001 ; $BOS_r > TTC_r$ at p -value $= 0.004$. \square

Result 4 indicates that, under BOS, a significant proportion of the students improve the ranking of their district schools. As we report in the Introduction, practical advice from the *Seattle Press*, *St. Petersburg Times*, as well as the CPAC in Minneapolis, are all variants of the District School Bias.

We now turn to a particular type of District School Bias. Recall that the Central Placement and Assessment Center in Minneapolis actually instructs the parents on how to manipulate: “make the first choice a true favorite and the other two ‘realistic’” [14]. We call this strategy the “Minneapolis strategy,” which is a special kind of District School Bias and Small School Bias. The Minneapolis strategy advocates that the reported top choice should be the true top choice, while the reported second or third choice should be a child's district school or a school with a large capacity (schools C–G). We find that, in the Boston

mechanism treatments in our experiment, 18.75% of our subjects use the Minneapolis strategy.¹¹

Result 5 (Minneapolis strategy). *Under the Boston mechanism, the average payoff from the Minneapolis strategy is not significantly different from that of other misrepresentations in the designed environment; however, in the random environment, the Minneapolis strategy performs significantly better than other types of misrepresentations.*

Support. Under BOS, in the designed environment, the average payoff from the Minneapolis strategy is 10.967, while the average payoff from other misrepresentations is 11.235. A *t*-test under the null hypothesis of equal mean yields a *p*-value of 0.685. In the random environment, however, the corresponding average payoffs are 14 from the Minneapolis strategy and 12.5 from non-Minneapolis manipulations. A *t*-test under the null hypothesis of equal mean yields a *p*-value of 0.027 in favor of the alternative hypothesis that the Minneapolis strategy yields higher payoffs. □

Result 5 provides the first empirical evaluation of the Minneapolis strategy, which is offered for one of the leading inter-district open enrollment programs. As the Minneapolis strategy reports an agent's true top choice and manipulates the second and third choices, we expect it to do well in less competitive environments, where most first choices can be accommodated. In such environments, the Minneapolis strategy is essentially equivalent to truthful preference revelation in most cases. The random environment is one such environment, as agent preferences are uncorrelated. Therefore, the Boston mechanism need not be very inefficient when first choices can mostly be accommodated.

We next investigate the efficiency of the three mechanisms. To identify the efficiency gain of each of the three choice mechanisms, we first look at the expected per capita payoffs in some benchmark cases. In computing these benchmarks, we use 10,000 randomly generated lotteries.

Table 6 presents the expected per capita payoff for six benchmark cases. In the first three cases, we compute the expected per capita payoff for each of the three mechanisms under the designed and random environments, assuming 100% truthful preference revelation. This gives us an idea of the upper bounds for the expected payoffs for these mechanisms. The last three columns present three relevant hypothetical mechanisms. For each of the 10,000 lotteries, the *Serial Dictator* mechanism lets the students sequentially choose their best school among the remaining schools, from the first student in the lottery to the last one in the lottery. The average per capita payoff over all lotteries gives us a reference payoff for full efficiency.¹² The *No Choice* mechanism simply assigns each student to her district school without any choices. This is the traditional public school assignment mechanism.

¹¹ Furthermore, we observe the Minneapolis strategy being used in GS and TTC as well, although by a much smaller fraction of subjects (6.94% and 9.03%, respectively). This is due to subjects' misunderstandings of the mechanisms.

¹² Abdulkadiroğlu and Sönmez show that the Serial Dictator mechanism is Pareto efficient for any ordering of the students and, conversely, any Pareto efficient allocation can be obtained for some ordering of the students [1].

Table 6
Expected per capita payoff in benchmark cases

Environment	100% truth			Hypothetical mechanisms		
	BOS	GS	TTC	Serial dictator	No choice	Random Assignment
Designed	12.546	12.078	12.413	12.262	9.360	8.277
Random	14.133	13.703	14.099	14.002	8.530	8.666

The *Random Assignment* mechanism randomly assigns each student to a school, regardless of residency or preferences.

Since all three mechanisms are implemented as true one-shot games with a total of 432 independent observations, we can use the recombinant estimator [23] to compare the mean efficiencies. With recombinant estimation, we recombine the strategies of different players to compute what the outcomes would have been if players' groupings had been different. Under each mechanism in each environment, we have a total of $(2)^{36}$ different groups. In the actual estimation, we randomly generate 14,400 groups¹³ under each mechanism to estimate the mean ($\hat{\mu}$), variance (σ^2) and covariance (φ) of the expected per capita payoff. Results of the recombinant estimation are reported in Table 7.

Result 6 (Efficiency). *In the designed environment, the efficiency ranking is highly significant: GS > TTC > BOS. In the random environment, the efficiency ranking is: GS ~ BOS > TTC. In all six cases, efficiency is significantly higher than in the No Choice mechanism.*

Support. The upper three rows of Table 7 report results of the recombinant estimation of expected per capita payoffs for the designed environment. A *t*-test of $H_0 : \hat{\mu}_x = \hat{\mu}_y$ against $H_1 : \hat{\mu}_x > \hat{\mu}_y$ yields $t = 5.505$ ($p < 0.001$) for $x = \text{GS}$ and $y = \text{BOS}$, $t = 2.402$ ($p = 0.008$) for $x = \text{TTC}$ and $y = \text{BOS}$, and $t = 2.209$ ($p = 0.011$) for $x = \text{GS}$ and $y = \text{TTC}$. The lower three rows of Table 7 report results of recombinant estimation of expected per capita payoffs for the random environment. A *t*-test of $H_0 : \hat{\mu}_x = \hat{\mu}_y$ against $H_1 : \hat{\mu}_x > \hat{\mu}_y$ yields $t = 0.364$ ($p = 0.359$) for $x = \text{BOS}$ and $y = \text{GS}$, $t = 2.165$ ($p = 0.015$) for $x = \text{BOS}$ and $y = \text{TTC}$, and $t = 1.728$ ($p = 0.042$) for $x = \text{GS}$ and $y = \text{TTC}$. All comparisons with the No Choice mechanism yield *p*-values < 0.01 . \square

Result 6 is one of our main results. Contrary to theoretical predictions, the Gale–Shapley mechanism, rather than TTC, emerges as the most efficient mechanism. This is due to the fact that significantly more subjects misrepresent their preferences under TTC than under GS. We speculate that the dominant strategy feature of TTC is not transparent to subjects without education. Similar empirical results have been discovered in the experimental auction literature, where not many subjects can figure out that bidding one's

¹³ We generate 200 recombinations per subject for each of the 72 subjects in each treatment.

Table 7

Recombinant estimation of expected per capita payoffs

Treatment	Mean, $\hat{\mu}$	Variance, σ^2	Covariance, ϕ	Asym. variance, $\text{var}(\hat{\mu})$	Standard deviation
BOS _d	11.150	0.040	0.001	0.011	0.105
GS _d	11.713	0.115	0.001	0.010	0.098
TTC _d	11.412	0.059	0.001	0.023	0.151
BOS _r	12.835	0.133	0.0002	0.004	0.064
GS _r	12.787	0.234	0.002	0.037	0.193
TTC _r	12.351	0.228	0.004	0.066	0.256

true valuation is a dominant strategy under Vickrey auction [21]. Another point to emphasize is that the efficiency ranking of BOS improves under the random environment. The success of the Minneapolis strategy in the random environment accounts for much of this efficiency improvement.

Ideally, we would compare the performance of these three mechanisms under many different environments. However, budget constraints make such extensive testing infeasible. Therefore, to understand how the environment affects mechanism performance, we run simulations using behavior calibrated from the experimental data. The simulation algorithm is described below.

1. We first extract the basic features of the existing payoff matrices and observed manipulation patterns.
 - For each student, in her true preference ordering, we find the following three numbers, which represent the most important characteristics of her payoff table:
 - (1) the ranking of her district school,
 - (2) the ranking of her more preferred small school, and
 - (3) the ranking of her less preferred small school.
 - For each student, in her reported preference ranking, we extract her manipulation patterns by finding
 - (4) the ranking of her true 1st choice,
 - (5) the ranking of her true 2nd choice,
 - ⋮
 - ⋮
 - (10) the ranking of her true 7th choice.
2. We then generate new payoff matrices from the designed and random environments. Let $\alpha \in [0, 1]$ be the probability of using the random environment data. For each α , we divide the students into seven categories based on their district schools: 1–3, 4–6, 7–12, 13–18, 19–24, 25–30 and 31–36. The process is outlined below:
 - For students 1–3:
 - Randomly draw a number x from the uniform distribution on $[0, 1]$. If $x \leq \alpha$, use the random environment data. Otherwise, use the designed environment data.
 - Randomly draw a number y from the discrete uniform distribution on $\{1, 2, \dots, 6\}$, as there are 6 students (3 in each session) whose district school is A in each environment.

- Use the y th student's data by preserving the true ranking of district school and small schools and randomly generating the ranking of other schools.
 - For students 4–6: follow similar steps as students 1–3.
 - For students 7–12:
 - Randomly draw a number x from the uniform distribution on $[0, 1]$. If $x \leq \alpha$, use the random environment data. Otherwise, use the designed environment data.
 - Randomly draw a number y from the discrete uniform distribution on $\{1, 2, \dots, 12\}$, as there are 12 students (6 in each session) whose district school is C.
 - Use the y th student's data by preserving the true ranking of district school and small schools, and randomly generating the ranking of other schools.
 - For students 13–18: follow similar steps as students 7–12.
 -
 -
 - For students 31–36: follow similar steps as students 7–12.
3. We imitate the manipulation pattern according to (4)–(10), i.e., the ranking of the true 1st choice, 2nd choice, \dots and the 7th choice in the reported preference ranking.

In the simulation, we use $\alpha \in [0, 1]$, with a step size of 0.05. For each given α , we generate 1000 payoff matrices. For each payoff matrix, we use 50 tie-breakers.

The environments generated can be characterized using two alternative indices. The first index looks at the proportion of first choices which can be accommodated.¹⁴ We call it the *first choice accommodation index*. This is a simple measure which characterizes the competitiveness of the environment. In the simulation, for each environment (α), when all students report their preferences truthfully, the first choice accommodation index is computed as the average number of students who receive their first choice under the Serial Dictator mechanism divided by 36.

A second measure, called the *competitiveness index*, looks at all choices. For each environment, we find all efficient allocations and the corresponding rank of the assigned school for each student. We then find the mean value of the sum of the ranks, and normalize the index by computing the max and min indices, where max is the competitive environment where everyone has identical preferences and min is the easy environment where every student gets her top choice. Thus, the competitiveness index for an environment, where the mean value of the sum of ranks of efficient allocation is x , equals $(x - \min)/(\max - \min)$.

Fig. 1 presents the first choice accommodation index (dashed line) and the competitiveness index (solid line) when the environment varies. These two indices are mirror images of each other. As α increases, preferences become increasingly uncorrelated. As a result, the proportion of first choices which can be accommodated increase and the environment becomes less competitive. From now on, we present our simulation results using the first choice accommodation index.

Fig. 2 reports the simulated per capita payoff for each first choice accommodation index for each mechanism. We use the Serial Dictator mechanism based on truthful preference revelation as a benchmark for full efficiency. As expected, as α increases, the per capita payoff of all three mechanisms increases. However, the rates of increase are different across mechanisms. In particular, the rankings of BOS and TTC switch when α is close to 0.4.

¹⁴ We thank Alvin Roth for suggesting this index.

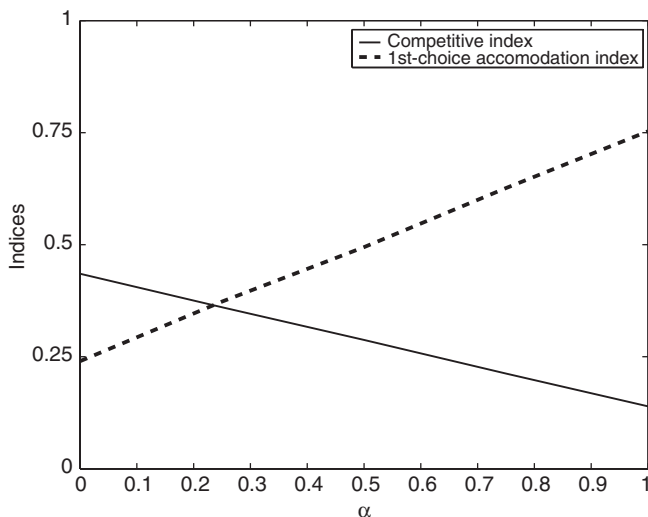


Fig. 1. First choice accommodation index and competitive index in different environments (α is the probability that data is drawn from the random environment).

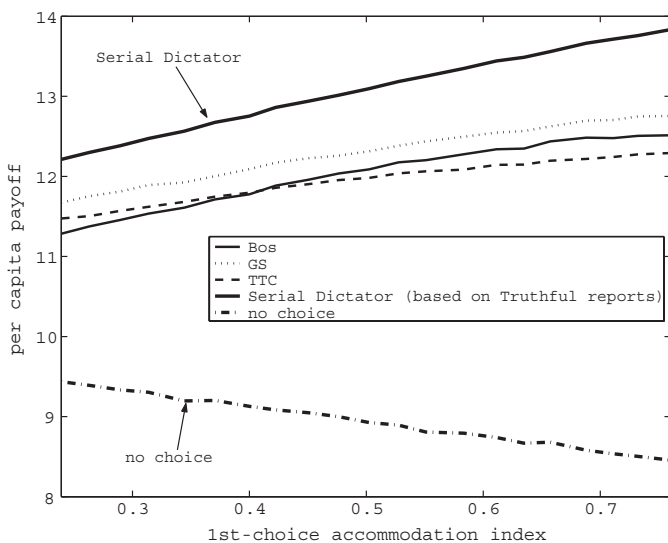


Fig. 2. Simulated per capita payoffs across mechanisms in different environments.

This is consistent with Result 6. Unlike Result 6, however, BOS never catches up with GS. Recall that the efficiency of mechanism x is computed as the ratio of per capita payoffs under mechanism x and the Serial Dictator. The simulated efficiency is reported in Fig. 3.

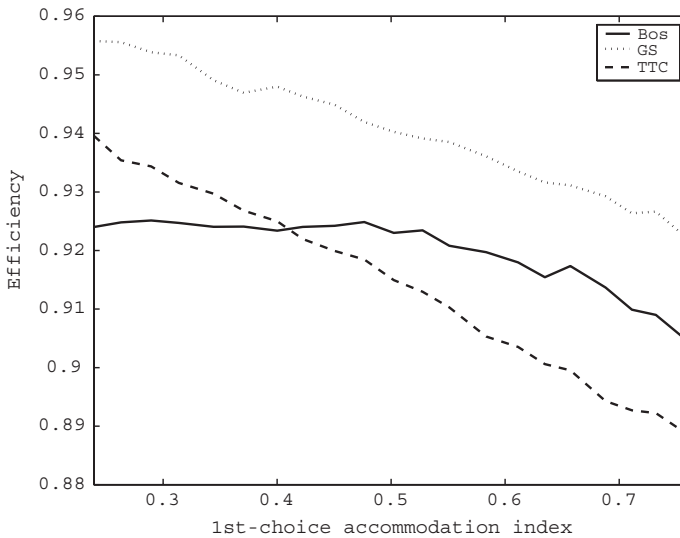


Fig. 3. Simulated efficiency across mechanisms in different environments.

When evaluating these mechanisms, it is important to examine each mechanism’s efficiency gain compared to the benchmark case of No Choice, where each student is simply assigned a seat in his district school. This benchmark is the traditional assignment mechanism for public schools in the U.S. Fig. 4 reports the simulated efficiency gain of each mechanism for a range of environments. We compute the efficiency gain of mechanism x as follows:

$$\text{Efficiency Gain of } x = \frac{\text{per capita payoff of mechanism } x - \text{per capita payoff of No Choice}}{\text{per capita payoff of Serial Dictator} - \text{per capita payoff of No Choice}}$$

Fig. 4 indicates that all three mechanisms have considerable efficiency gains compared to the No Choice mechanism. In particular, the efficiency gain of the Gale–Shapley mechanism is consistently above 80%, while that of the Boston mechanism varies depending on the environment. We now summarize our simulation results.

Result 7 (Simulated efficiency). *In competitive environments, the efficiency ranking of the mechanisms is: GS > TTC > BOS. In less competitive environments, the efficiency ranking is: GS > BOS > TTC.*

Support. Fig. 5 reports the z -statistics of pairwise t -tests comparing the mean payoffs of different mechanisms. The four horizontal lines indicate z -statistics exactly at the 1% and 5% significance levels. From these tests, GS is significantly more efficient than TTC at the 1% level for any environment, as GS – TTC always lies above the band. Similarly, BOS is always significantly less efficient than GS at the 1% level, as BOS – GS always lies below

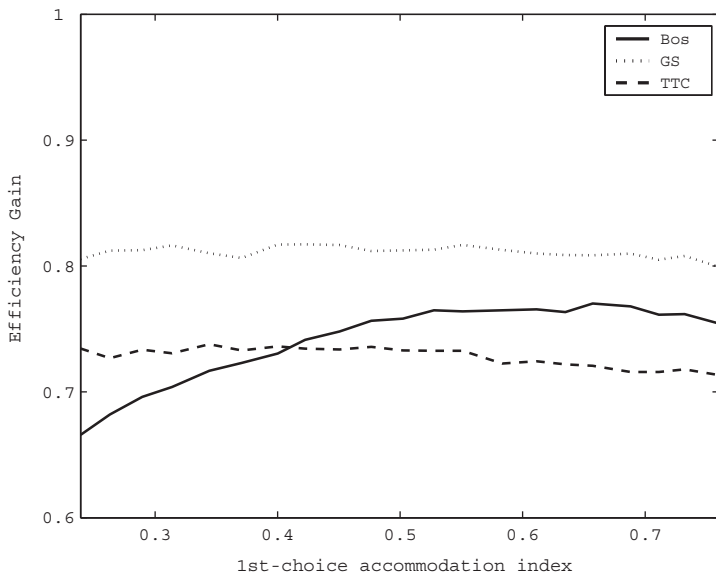


Fig. 4. Simulated efficiency gains across mechanisms in different environments.

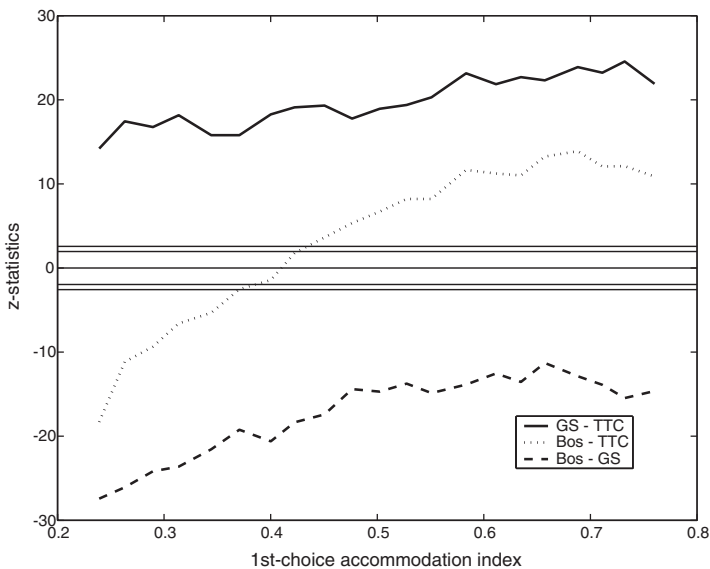


Fig. 5. Pairwise efficiency comparisons and z-statistics.

the band. TTC is significantly more efficient than BOS (at the 5% level) when the first choice accommodation index is less than 0.37. It is significantly less efficient than BOS (at the 5% level) when the first choice accommodation index is more than 0.42. □

Both our experimental and simulation results indicate that the Gale–Shapley mechanism performs better than BOS or TTC in terms of efficiency when implemented among boundedly rational agents.

6. Conclusion

For parents, there is no decision more important than the school choice decision. This experiment studies the school choice decision under three mechanisms. Our experimental evidence identifies deficiencies of the widely used Boston mechanism. It also offers insight into the effectiveness of two alternative mechanisms, the Gale–Shapley and Top Trading Cycles mechanisms.

Education practitioners have promoted the Boston mechanism by arguing that, under this mechanism, more than 70% of students are assigned their first choices [8,15]. On the other hand, theory suggests that, under the Boston mechanism, students can improve their assignments by improving the ranking of schools for which they have high priorities, such as their district schools [2]. There is strong evidence that parents are well aware of this weakness of the Boston mechanism. In a recent *Boston Globe* article, Cook states [7]:

Parents who learn how the system works begin to “game” it by lying about their top choices, ranking a less desirable school first in the hope that they’ll have a better chance of landing a spot.

Our experimental evidence is consistent with both the field data and the theory. Under the Boston mechanism, more than 70% of the participants are assigned a seat at their first choice based on their stated preference; however, because more than 80% of the participants misrepresent their preferences, less than 30% of them are assigned a seat at their first choice based on their true preferences. Therefore, we can conclude that an evaluation of the Boston mechanism based on stated preferences is not adequate. Indeed, often important policy decisions, such as which schools shall be terminated, depend on how popular schools are, based on stated preferences [28]. Experimental evidence supports the theory and suggests that such evaluation is not adequate under the Boston mechanism.

Furthermore, our experimental evidence on the efficiency of the Boston mechanism is mixed: While the efficiency of the Boston mechanism is significantly lower than that of the Gale–Shapley and TTC mechanisms under the designed environment, it is comparable to that of the Gale–Shapley mechanism under the random environment. In order to understand the difference, it is important to analyze the patterns of misrepresentation as well as the implications of these patterns within different environments.

First, the Boston mechanism puts a very high weight on top choices, especially the first choice. In both environments, there are two small schools and five large schools. Other things being equal, small schools are more likely to have stiff competition. Experimental evidence shows that participants whose (true) first choice is one of the large schools are more likely to truthfully reveal their preferences than those whose first choice is a small school. In the designed environment, preferences are correlated with school quality, with the small schools being high-quality schools. In the random environment, preferences are random and only a small fraction of the participants have one of the small schools as his

top choice. As a result, significantly more participants reveal their preferences truthfully under the random environment than the designed environment, considerably reducing the efficiency loss for the random environment.

Second, under the Boston mechanism, 19% of all participants adopt the manipulation strategy advocated by the Central Placement and Assessment Center (CPAC) in Minneapolis:

Make the first choice a true favorite and the other two “realistic,” that is strategic choices.

The Minneapolis strategy is adopted by these participants via keeping their true first choice and misrepresenting their second or third choices by either improving the ranking of their district school to the second or third place or by improving the ranking of a school with more seats to the second or third place. This strategy has different efficiency implications in the two environments. In the random environment, there is little conflict of interest and it is possible to assign most participants a seat at their first choice. Consequently, for most participants, the “relevant” component of the stated preference is merely the first choice. Hence, the Minneapolis strategy results in a very small efficiency loss in this environment. By contrast, in the designed environment, most participants’ first choices are the same and only a fraction of the first choices can be accommodated. Hence, the entire preference profile is “relevant” under the designed environment. Therefore, the Minneapolis misrepresentation causes significant efficiency loss in this environment.

These two factors are critical in explaining the efficiency difference for the Boston mechanism under the two environments. In reality, first choices are likely to be correlated, and therefore the efficiency loss can be expected to be high.

Under the Boston mechanism, students are forced to play a very complicated preference revelation game: Their best responses depend heavily on other students’ stated preferences and truthful preference revelation is rarely in their best interest. Cook states [7]:

Of course, no new system can create more seats at the most sought-after schools. But all parents interviewed by the *Globe* said that it would be a huge relief simply to write a truthful answer to the question: What school do you want?

“A lot of the alienation some parents have toward the choice system is solely attributable to the alienation of not making your first choice your first choice,” said Neil Sullivan, the father of four children who have attended Boston public schools.

The game induced by the Boston mechanism often has multiple equilibria [12]. This is yet another major deficiency of the Boston mechanism. By contrast, the games induced by the Gale–Shapley and the TTC mechanisms are considerably simpler. Under either mechanism, truthful preference revelation is a dominant strategy. However, it is not unusual for such strategy-proof mechanisms to underperform in laboratory experiments, as many participants do not recognize their dominant strategies. Well-known examples include the second price sealed-bid auction [21] and the pivotal mechanisms [4]. Naturally this possibility increases as the mechanism description becomes more complex. Experimental results show that 36% of participants misrepresent their preferences under the Gale–Shapley mechanism and 53% of participants misrepresent their preferences under TTC. The misrepresentation rate is especially high for TTC, reflecting its relatively more complex description. If either

mechanism is adopted for real-life implementation, participants should be clearly instructed that they can only hurt themselves by misrepresenting their preferences. Empirical patterns of misrepresentation identified in our experiment can be used for educational brochures for parents.

Perhaps the most surprising result of our experiments concerns the efficiency comparison of the three mechanisms, as our experimental results do not support theory. Specifically, we find that:

- In the designed environment, the Gale–Shapley mechanism is significantly more efficient than TTC, which is significantly more efficient than the Boston mechanism.
- In the random environment, the Boston mechanism has the same efficiency as the Gale–Shapley mechanism, and each is significantly more efficient than TTC.

Simulation shows that, in more competitive environments, GS dominates TTC, which in turn dominates BOS. In less competitive environments, GS dominates BOS, which in turn dominates TTC.

In extending these results to real-life implementation, we recommend that participants should be educated on the strategy-proofness of both the Gale–Shapley and the TTC mechanisms. Therefore, efficiency loss in these two mechanisms can be reduced. This will also reduce the efficiency advantage of GS over TTC or it may even give an efficiency advantage to TTC. Since we do not prompt or educate our subjects in the experiment, the efficiency loss of the two alternative mechanisms can also be used as a measure for the potential value of education.

Interestingly, in October 2003, New York City public schools decided to use a version of the Gale–Shapley mechanism to determine the placement of about 100,000 incoming ninth and tenth graders [18]. In addition, Boston Public Schools are currently working with matching theorists to discuss a possible change to either the Gale–Shapley or the Top Trading Cycles mechanism [37]. We hope that this paper will stimulate similar discussions at school districts that rely on variants of the Boston mechanism. We look forward to studying actual implementations of these two mechanisms in the field.

Acknowledgments

We thank Atila Abdulkadiroğlu, Tim Cason, Vincent Crawford, Eddie Dekel, Valerie Edwards, Carleton Jones, Peter Katuscak, Lutz Kilian, Aradhna Krishna, Parag Pathak, Alvin Roth, anonymous referees and seminar participants at 2003 Economic Science Association Meetings (Pittsburgh, PA), 2004 AEA (San Diego, CA), ESSET 2004 (Gerzensee), Games 2004 (Marseille), Stanford SITE 2004, Duke, Michigan, Ohio State, and Pittsburgh for helpful discussions, and Kan Takeuchi for excellent research assistance. Chen acknowledges the research support of the National Science Foundation through Grant No. SES-0079001 and Sönmez acknowledges the research support of KoçBank via the KoçBank scholar program and Turkish Academy of Sciences in the framework of the Young Scientist Award Program via Grant TS/TÜBA-GEBİP/2002-1-19. Any remaining errors are our own.

References

- [1] A. Abdulkadiroğlu, T. Sönmez, Random serial dictatorship and the core from random endowments in house allocation problems, *Econometrica* 66 (1998) 689–701.
- [2] A. Abdulkadiroğlu, T. Sönmez, School choice: a mechanism design approach, *Amer. Econ. Rev.* 93 (3) (2003) 729–747.
- [3] M. Balinski, T. Sönmez, A tale of two mechanisms: student placement, *J. Econ. Theory* 84 (1999) 73–94.
- [4] Y. Chen, Incentive-compatible mechanisms for pure public goods: a survey of experimental research, in: C. Plott, V. Smith (Eds.), *The Handbook of Experimental Economics Results*, Elsevier Press, Amsterdam, forthcoming.
- [5] Y. Chen, T. Sönmez, Improving efficiency of on-campus housing: an experimental study, *Amer. Econ. Rev.* 92 (5) (2002) 1669–1686.
- [6] S. Collins, K. Krishna, The harvard housing lottery: rationality and reform, Manuscript, Georgetown University and Penn State University, October 1995.
- [7] G. Cook, School assignment flaws detailed—two economists study problem, offer relief, *The Boston Globe*, September 12, 2003.
- [8] P.W. Cookson, *School Choice: The Struggle for the Soul of American Education*, Yale University Press, New Haven, 1994.
- [9] J.B. Cullen, S. Levitt, The effect of school choice on student outcomes: evidence from randomized lotteries, Manuscript, University of Michigan, October 2003.
- [10] L. Dubins, D. Freedman, Machiavelli and the Gale–Shapley Algorithm, *Amer. Math. Monthly* 88 (1981) 485–494.
- [11] D. Epple, R. Romano, Competition between private and public schools: vouchers and peer group effects, *Amer. Econ. Rev.* 88 (1998) 33–62.
- [12] H. Ergin, T. Sönmez, Games of school choice among boston students, Manuscript, MIT and KoçUniversity, 2003.
- [13] D. Gale, L. Shapley, College admissions and the stability of marriage, *Amer. Math. Monthly* 69 (1962) 9–15.
- [14] S. Glazerman, R.H. Meyer, Public school choice in Minneapolis, midwest approaches to school reform, in: *Proceedings of a Conference Held at the Federal Reserve Bank of Chicago*, 1994.
- [15] C. Glenn, Controlled choice in Massachusetts public schools, *Public Interest* 103 (1991) 88–105.
- [16] G. Harrison, K. McCabe, Stability and preference distortion in resource matching: an experimental study of the marriage problem, in: V. Smith (Ed.), *Research in Experimental Economics*, vol. 6, JAI Press, 1996.
- [17] E. Haruvy, A.E. Roth, M.U. Ünver, Dynamics of law clerk matching: an experimental and computational investigation of proposals for reform of the market, Harvard NOM Research Paper No. 01-08 (2001).
- [18] D. Herszenhorn, Revised admission for high schools, *New York Times*, October 3, 2003.
- [19] C. Hoxby, Do private schools provide competition for public schools?, NBER Working Paper 4978 (1994).
- [20] C. Hoxby, Does competition among public schools benefit students and taxpayers?, *Amer. Econ. Rev.* 90 (2000) 1209–1238.
- [21] J. Kagel, Auctions: a survey of experimental research, in: J. Kagel, A. Roth (Eds.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ, 1995, pp. 131–148.
- [22] J.H. Kagel, A.E. Roth, The dynamics of reorganization in matching markets: a laboratory experiment motivated by a natural experiment, *Quart. J. Econ.* 115 (1) (2000) 201–235.
- [23] D. Lucking-Reiley, C. Mullin, Recombinant estimation for normal-form games, with applications to auctions and bargaining, *Games Econ. Behav.*, in press.
- [24] C. Mas, *Navigating the School Choice Maze*, The Seattle Press, December 30, 1998.
- [25] M.N. McKinney, M. Niederle, A.E. Roth, The collapse of a medical labor clearinghouse (and why such failures are rare), *Amer. Econ. Rev.*, forthcoming.
- [26] H. Nalbantian, A. Schotter, Matching and efficiency in the baseball free-agent system: an experimental examination, *J. Lab. Econ.* 13 (1995) 1–31.
- [27] T.J. Nechyba, Mobility, targeting and private school vouchers, *Amer. Econ. Rev.* 90 (1) (2000) 130–146.
- [28] A.P. Abraham (Ed.), Office of Educational Research and Improvement, *Getting Started: How Choice Can Renew Your Public Schools*. U.S. Government Printing Office, Washington, DC, 1992.

- [29] M. Olson, D. Porter, An experimental examination into design of decentralized methods to solve the assignment problem with and without money, *Econ. Theory* 4 (1994) 11–40.
- [30] S. Papai, Strategy proof assignment by hierarchical exchange, *Econometrica* 68 (2000) 1403–1433.
- [31] A.E. Roth, The economics of matching: stability and incentives, *Math. Oper. Res.* 7 (1982) 617–628.
- [32] A.E. Roth, Incentive compatibility in a market with indivisibilities, *Econ. Lett.* 9 (1982) 127–132.
- [33] A.E. Roth, The evolution of labor market for medical interns and residents: a case study in game theory, *J. Polit. Economy* 92 (1984) 991–1016.
- [34] A.E. Roth, A natural experiment in the organization of entry-level labor markets: regional markets for new physicians and surgeons in the united kingdom, *Amer. Econ. Rev.* 81 (1991) 414–440.
- [35] A.E. Roth, The economist as engineer: game theory, experimentation and computation as tools for design economics, *Econometrica* 70 (2002) 1341–1378.
- [36] A.E. Roth, M. Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press, Cambridge, 1990.
- [37] K. Rothstein, Plans eyed for school assignment lottery, *Boston Herald*, October 14, 2004.
- [38] C.E. Rouse, Private school vouchers and student achievement: an evaluation of the milwaukee parental choice program, *Quart. J. Econ.* 113(2) 553–602.
- [40] M. Schneider, P. Teske, M. Marschall, *Choosing schools: Consumer Choice and the Quality of American Schools*, Princeton University Press, Princeton, NJ, 2000.
- [41] L. Shapley, H. Scarf, On cores and indivisibility, *J. Math. Econ.* 1 (1974) 23–28.
- [42] V. Smith, Microeconomic systems as an experimental science, *Amer. Econ. Rev.* 72 (1982) 923–955.
- [43] T. Tobin, Yep, it's complicated, *St. Petersburg Times*, September 14, 2003.
- [44] M.U. Ünver, On the survival of some two-sided matching mechanisms: an experimental and computational investigation of the stability hypothesis, *Int. J. Game Theory*, forthcoming.

Further reading

- [39] *School Choice: A Special Report*, The Carnegie Foundation for the Advancement of Teaching, 1992.