

School Choice and Information

An Experimental Study on Matching Mechanisms

Joana Pais*, Ágnes Pintér†

May, 2007

Abstract

We present an experimental study where we analyze three well-known matching mechanisms—the Boston, the Gale-Shapley, and the Top Trading Cycles mechanisms—in different informational settings. Our experimental results are consistent with the theory, suggesting that the TTC mechanism outperforms both the Boston and the Gale-Shapley mechanisms in terms of efficiency and it is slightly more successful than the Gale-Shapley mechanism regarding the proportion of truthful preference revelation, whereas manipulation is stronger under the Boston mechanism. In addition, even though agents are much more likely to revert to truth-telling in lack of information about the others' payoffs—ignorance may be beneficial in this context—the TTC mechanism results less sensitive to the amount of information that participants hold. These results therefore suggest that the use of the TTC mechanism in practice is more desirable than of the others.

*Corresponding author; E-mail: jpais@iseg.utl.pt; Instituto Superior de Economia e Gestão, Universidade Técnica de Lisboa and UECE; Rua Miguel Lupi, 20; 1249-078 Lisboa; Portugal.

†E-mail: apinter@eco.uc3m.es; Departamento de Economía, Universidad Carlos III de Madrid; C. Madrid, 126; 28903 Getafe (Madrid); Spain.

1 Introduction

There is now a vast literature on matching problems. Matching is a pervasive phenomenon arising in several economic and social settings. The assignment of civil servants to civil service positions, the admission of students to colleges, some entry-level labor markets—as the widely explored market for graduating physicians—, or the school choice problem are among the matching situations that have gained attention in the last decades. The working of some matching mechanisms, along with strategic issues that confront individuals in these contexts, have been explored theoretically under the assumption of complete information.

Very briefly, in a two-sided matching market, agents belong to one of two disjoint sets, say colleges and students, and each agent—college and student—has preferences over the other side of the market—students and colleges, respectively—and the prospect of being unmatched. The matching problem then reduces to assigning students to colleges by means of a matching mechanism. Stability, strategy-proofness, and Pareto efficiency of such mechanisms are highly valued properties. A mechanism is stable if it always selects stable matchings; by definition, under a stable matching every agent in the market prefers his partner to being alone and, moreover, no pair of agents—consisting of a college and a student—who are not matched to each other would rather prefer to be so matched. A mechanism is strategy-proof if it is immune to preference manipulation, i.e., truth is a dominant strategy. A mechanism is Pareto efficient if it always selects Pareto efficient matchings.

The perhaps most famous matching mechanism relies on the Gale-Shapley deferred-acceptance algorithm (Gale and Shapley, 1962). Gale and Shapley were motivated by the problem of the admission of students to colleges and the Gale-Shapley algorithm was written as a means to show that a stable matching always exists in such a two-sided matching market. The Gale-Shapley deferred-acceptance algorithm transforms a matching where all agents are unmatched into a stable matching, thus proving existence. Besides guaranteeing stability, the Gale-Shapley (GS) mechanism has other appealing properties. Namely, truth is a dominant strategy for one side of the market (Dubins and Freedman, 1981, Roth, 1982a). Moreover, it is Pareto efficient when the welfare of both sides of the market is considered (Roth, 1982a).

Most theoretical studies on matching mechanisms rely on the assumption of complete information, however implausible: knowing the true preferences of every agent in the market is more than we may reasonably expect in most matching markets. Only a few papers have relaxed this assumption and are thus worth mentioning. Roth (1989) is a first attempt to deal with the incom-

plete information case. Under incomplete information, even though truth remains a dominant strategy for one side of the market when the Gale-Shapley mechanism is employed, the equilibrium characterization for the complete information case is not robust. Ehlers and Massó (2003) study Bayesian Nash equilibria for mechanisms producing stable matchings—as the Gale-Shapley mechanism—and find a necessary and sufficient condition for truth-telling to be an equilibrium: truth-telling is a Bayesian Nash equilibrium in the revelation game induced by a stable mechanism and a common belief if and only if all profiles in its support have a singleton core. Finally, Roth and Rothblum (1999) and Ehlers (2003, 2004) are less ambitious and do not aim at characterizing equilibria, but give advice to individuals on how to participate in matching markets when there is uncertainty about the others' strategies.

Still, many questions regarding the strategic incentives agents face under incomplete information remain to be answered on theoretical grounds. How the amount of information held by individuals on the elements of the game actually influences individuals' decision making, affecting the performance of matching mechanisms, is thus a question to be explored. For instance, Barberà and Dutta (1995) consider truth-telling as a form of "protective" behavior, claiming that risk averse agents may revert to faithfully revealing their true preferences when they are poorly informed. Moreover, it is clear that in mechanisms for which truth is not a dominant strategy, computing the optimal strategies requires a lot of information on others' preferences. In this paper we present an experimental study to investigate these and other issues, providing a direction into which the role of information on decision making may be ascertained.

We investigate a particular class of matching problems: the assignment of individuals to indivisible items. In these problems, individuals—let us call them teachers—have strict preferences over the indivisible items—henceforth, schools—and, on the other hand, schools have a maximum capacity and a strict priority ordering of all teachers. This problem has been referred to as the school choice problem (Abdulkadiroglu and Sönmez, 2003) and is closely related to the college admissions problem explored by Gale and Shapley (Gale and Shapley, 1962), the main difference being that, in contrast to the college admissions model, here schools are not strategic agents, but mere objects to be assigned to teachers. Hence, while teachers may not straightforwardly reveal their true preferences, schools have no chance of manipulating priorities.

The influence of information is assessed for the GS mechanism and for another well-known matching mechanism, the Top Trading Cycles (TTC) mechanism, as well as for the Boston mechanism, which has been widely

used in the context of school choice problems.¹ The TTC (Shapley and Scarf, 1974) fulfills two appealing properties—it is both strategy-proof (Roth, 1982b) and Pareto efficient—but it is not stable. The GS mechanism is both strategy-proof and stable, but not efficient (Roth, 1982a), since we only consider teachers’ welfare in this setup. Finally, we have included the Boston mechanism for reference, as it fails to meet all three properties: it is not strategy-proof, neither stable, nor Pareto efficient. Note that, whereas the GS mechanism fails to be efficient even when agents act straightforwardly, under the Boston mechanism efficiency losses result from the fact that there is room for profitable manipulation and Pareto efficiency would result if agents revealed their true preferences.²

Besides providing yet another test of theoretical results on matching mechanisms with boundedly rational individuals, we address two main questions. First, we compare the three above mentioned mechanisms under four informational scenarios, ranging from complete ignorance about the other participants’ preferences and schools’ priorities to complete information on all elements of the game. In particular, we are concerned in comparing the incentives agents face under different mechanisms, as well as in comparing efficiency levels and stability of the outcomes, for different information levels. These comparisons may have important policy implications. The results in this paper suggest that, just as predicted by theory, the TTC mechanism prevails over the GS in what efficiency is concerned. Moreover, when agents hold very little or full information on the elements of the game, the TTC may be more successful than the GS in inducing truth-telling. On the other hand, the Boston mechanism performs surprisingly well, delivering an exceptionally high proportion of stable matchings and efficiency levels that are close to those obtained under any of the two alternative mechanisms. Second, within each mechanism, we evaluate the influence of the amount of information held by individuals on decision making. Namely, we are concerned in testing whether truth-telling emerges as a very salient form of behavior when information is low. Moreover, if we are able to determine that information significantly affects individuals’ behavior, we can immediately conclude that the existing theoretical results, which rely on the fundamental assumption

¹School choice programs have become increasingly popular in the US. The best known of these programs rely on the so-called Boston mechanism, which was used to assign students to schools in Boston, having been recently replaced by a mechanism based on Gale-Shapley deferred-acceptance algorithm. The Boston mechanism remains in use in Cambridge and Seattle, among others.

²On the functioning and strategic properties of the Boston mechanism, check Abdulkadiroglu and Sönmez (2003), Abdulkadiroglu et al. (2005), Chen and Sönmez (2006), and Ergin and Sönmez (2004).

of complete information, are insufficient to deal with markets where agents know little about others. Otherwise, if the effect of information is not relevant, theory may be considered apt to deal with the incomplete information case. Our results partially support the first conjecture: in general, there is a large difference between scenarios where information levels are extremely low and those where agents hold substantial levels of information, while the differences between partial and full information scenarios are not significant. In particular, in a very low information environment, acting straightforwardly is a very salient form of behavior and there is a significant drop in the proportion of agents who play truthfully once agents have some information on the elements of the game. Furthermore, in low information environments, significantly higher levels of efficiency are achieved under every mechanism except for TTC, which appears to be less sensitive to the amount of information held by participants. Finally, our results do not disclose a significant effect of information on the proportion of stable matchings achieved under any of the mechanisms under study.

We are aware of several experimental studies of matching problems, some of which aim at testing the above mentioned mechanisms. The main difference with respect to this paper derives from our main objective: to test the role of information in evaluating matching mechanisms. These studies include Harrison and McCabe (1996) that explores the GS mechanism and shows that profitable manipulation of agents' preferences becomes more difficult as markets get larger; Chen and Sönmez (2002a) that compares a random serial dictatorship mechanism used to allocate dormitory rooms in American universities with a variant of the TTC in an incomplete information environment, concluding that the TTC produces significantly more efficient allocations; in a companion paper, Chen and Sönmez (2002b) evaluate the performance of these mechanisms under complete information, reaching the same qualitative results. Finally, Chen and Sönmez (2006), consider the school choice problem and analyze the TTC, the GS, and the Boston mechanisms under incomplete information, concluding that, in what efficiency is concerned, the TTC outperforms the Boston mechanism and in turn the GS improves upon the TTC. This efficiency reversal result contrasts with the results obtained in this paper. Other experimental studies, dealing with other matching mechanisms, are: Olson and Porter (1994), Nalbantian and Schotter (1995), Kagel and Roth (2000), Ünver (2001), Haruvy, Roth, and Ünver (2001), and McKinney, Niederle, and Roth (2005).

We proceed as follows. In Section 2 we present the theoretical properties of the three matching mechanisms under study. We describe the experimental design in Section 3. Section 4 summarizes the main results of the experiments. Some concluding remarks follow in Section 5.

2 The Theoretical Model

We first introduce the model and then describe the three matching mechanisms and their theoretical properties.

In this assignment problem there are a number of teachers to fill a number of vacancies or teaching positions across different schools. Each teacher has strict preferences over all schools, while each school has a strict priority ranking of all teachers, as well as a maximum number of teachers to employ. Priorities are exogenous and not subject to manipulation by schools. The fact that only teachers can act strategically is what distinguishes this problem from the college admissions model. It is not difficult to justify the use of the school choice model. Besides the fact that this model is easier to be implemented in the laboratory, we can find plenty of real-life situations that can be described as one side of the market being inactive. We have already mentioned the use of the school choice problem in the admission of children to public schools in the US, which also applies to other countries as Spain, but we can also think about the admission of students to universities (which is, in most countries, based on students' grades), the MIR system ("Médicos Internos Residentes," a residence training system for physicians in public hospitals based on their performance) in Spain, or in general the assignment of civil servants to civil service positions (for example, teachers, judges, or tax inspectors in Spain), which is, in several countries, based on an objective scoring system.

The outcome of the school choice problem is a matching, an assignment of teachers to teaching positions such that each teacher is assigned one vacancy and each vacancy is filled by one teacher only. A matching is Pareto efficient if there is no matching that assigns at least one teacher a strictly better school and every other teacher a weakly better school, and it is stable if every agent in the market prefers his partner to being alone and, moreover, no pair of agents—consisting of a school and a teacher—who are not matched to each other would rather prefer to be so matched. A matching mechanism consists of a systematic procedure that selects a matching for each school choice problem. A matching mechanism is efficient if it always chooses Pareto efficient matchings; it is stable if it always selects stable matchings; and it is strategy-proof if truth is a dominant strategy, i.e., no teacher can profitably manipulate her preferences, independently of the other agents' strategies.

2.1 The Top Trading Cycles Mechanism

In this context, the TTC works as follows:

1. Each school gives priority to a number of teachers up to its capacity; in this setting, for simplicity, each teacher has priority in one school only.
2. Each teacher reports her preferences over the schools.
3. An ordering of teachers is randomly chosen.
4. For any submitted teachers' preferences, schools' priorities, and ordering of teachers, the outcome is obtained after undergoing the following steps:
 - (a) Assign each teacher to a school (tentative assignment); in this setting, each teacher is tentatively assigned to her priority school.
 - (b) The first teacher in the ordering proposes to her top ranked school. If she has priority at this school, the assignment is finalized and both the teacher and teaching position are removed from the system; the procedure continues with the second teacher in the ordering. Otherwise, the first teacher in the ordering that is tentatively assigned to the proposed school is inserted at the top of the ordering, in front of the requester.
 - (c) When the ordering is modified, this procedure is repeated, so that the teacher who just became first in the ordering sends an application to her highest-ranked school. If she has priority at this school, the assignment is finalized and the procedure continues with the next teacher in line. Otherwise, the first teacher in the ordering tentatively assigned to the proposed school is inserted at the top of the ordering, in front of the requester.
 - (d) If a cycle forms, it consists of a sequence of proposals of the kind: A proposes to B's tentative assignment, B applies to C's tentative assignment, and C proposes to A's tentative assignment. In such cases, all teachers in the cycle are assigned to the schools they proposed to and teachers, as well as their respective assignments, are removed from the system.
 - (e) The procedure stops when all teachers are assigned to a position.

The TTC mechanism satisfies two appealing properties: it is strategy-proof, i.e., truth is a dominant strategy for every teacher, and Pareto efficient, but it is not stable.³ We thus expect that individuals reveal their preferences

³Note also that, in this setting, the resulting final assignment is independent of the random ordering of teachers defined in step 3, as proved in Theorem 3 in Abdulkadiroglu and Sönmez (1999).

in a straightforward manner, independently of the amount of information they hold on the elements of the game.⁴ As a result, we expect to observe a high efficiency level but not necessarily a high frequency of stable outcomes.

2.2 The Gale-Shapley Mechanism

The GS mechanism is certainly one of the best known mechanisms in the matching literature. Its theoretical properties and the incentives it gives to agents have been scrutinized and its applications encompass a significant number of markets. In what follows we describe the functioning of the GS mechanism:

1. A priority ordering of teachers is determined for each school.
2. Each teacher reports her preferences over the schools.
3. Given the submitted preferences of the teachers and schools' priority orderings, positions are allocated after undergoing the following steps:
 - (a) Each teacher proposes to her first ranked school. Each school keeps the applicants with higher priority order on hold until positions are filled, while rejecting the lowest priority teachers in excess of its capacity.
 - (b) In general:

Every teacher who got rejected in the previous step proposes to the next school on her list of preferences. Each school considers the teachers it holds from the previous step together with the new applications. The lowest priority teachers in excess of the school's capacity are rejected, while remaining applications are kept on hold.
 - (c) This process is repeated until no applications are rejected. Each participant is then assigned the position at the school that keeps her on hold.

As the TTC mechanism, the GS mechanism is strategy-proof. Again, we expect individuals to faithfully reveal their true preferences over schools

⁴Still, Chen and Sönmez (2002a) find, in an experiment about on-campus housing, that about one-third of the subjects manipulate their preferences under a variant of the TTC mechanism and Chen and Sönmez (2006) report preference misrepresentation levels that amount to around one-half of the subjects.

in every informational treatment.⁵ This mechanism is efficient when the welfare of both sides of the market is taken into account. Nevertheless, in this assignment problem, schools are mere objects to be allocated among teachers and only teachers' welfare is taken into consideration in the determination of the efficiency level. Since there may exist a matching that Pareto dominates the outcome of the GS mechanism for teachers, the mechanism is not efficient in this setup. It follows that, if theory is to be confirmed, the TTC should outperform the GS in efficiency terms. On the other hand, when considering stability, theory predicts that the occurrence of stable matchings under the GS mechanism should be more frequent than under either the TTC or the Boston mechanisms.

2.3 The Boston Mechanism

The Boston mechanism has been the most widely used assignment mechanism in real-life applications of school choice problems. It works as follows:

1. A priority ordering of teachers is determined for each school.
2. Each teacher reports her preferences over the schools.
3. Given the submitted preferences of the teachers and schools' priority orderings, positions are allocated after several rounds:
 - (a) Each teacher proposes to her top ranked school. Each school accepts the proposals from the teachers with higher priority order until positions are filled (or no teachers proposing to the school remain). These applicants and their positions are removed from the system. All other applications are rejected by the schools.
 - (b) In general at round k :
Each teacher remaining in the system proposes to its k^{th} school. Each school with vacant positions accepts the proposals from the teachers with higher priority order until positions are filled (or no teachers proposing to the school remain in this round). These applicants and their positions are removed from the system. All other applications are rejected by the schools.
 - (c) The procedure terminates when each teacher is assigned a position.⁶

⁵Even though Chen and Sönmez (2006) report an experiment where more than one-third of the participants manipulate their preferences.

⁶Hence, if there are n teachers and l schools, the process ends in a maximum number of l rounds.

A major handicap of the Boston mechanism is that it leads to preference manipulation.⁷ In fact, teachers are given incentives to rank high on their submitted preferences the schools where they have good chances of getting in. This has two important consequences. First, evaluating the performance of this mechanism according to the revealed preferences is clearly inadequate. Moreover, even though the outcome of the Boston mechanism is Pareto efficient when teachers submit their true preferences, preference manipulation may lead to a substantial efficiency loss. Hence, we expect high rates of preference manipulation and a low level of achieved efficiency. Moreover, as the mechanism is not stable, the GS mechanism should outperform the Boston mechanism in this aspect as well.

3 Experimental Design

These experiments were designed to analyze participants' decision making in different informational settings under each of the above described matching mechanisms: the Boston, the GS, and the TTC mechanisms. We use a 3x4 design: for each mechanism we construct four treatments differing in the amount of information held by participants about the elements of the game. We then compare decision making throughout the treatments, concentrating on the role of information in truthful preference revelation, in the achieved level of efficiency, and in stability. The environment is designed to capture the key aspects and difficulties of each mechanism, under a controlled environment, with relatively small groups of participants.

Participants were randomly and anonymously sorted into groups of five. Each participant plays the role of a teacher to be assigned to a teaching position. For each group of five teachers, there are five vacancies—or teaching positions—across three schools that differ in capacity (number of opening positions) and desirability. Each position should be assigned to one teacher only. Preferences over schools are induced by the monetary payoff a teacher obtains depending on the school where she fills a vacancy at the end of the experiment. The payoffs obtained are symmetric: every teacher gets 15 experimental currency units (ECU) for her top choice, 9 ECU for the second choice, and 3 ECU for the last choice, but different teachers need not agree on which school is either her top, second, or last choice. In the experiment, 1 ECU equals 0.5 Euro. The payoffs of different outcomes are sufficiently dispersed so as to have a monetarily salient difference (12 ECU, which equals 6 Euro) between getting one's best and one's worst choice.

⁷In Chen and Sönmez (2006) approximately 80% of the subjects manipulate their true preferences.

Finally, schools have priorities over teachers. This means that schools may prefer some applicants to others and are able to rank all the participants in a list of priorities. Moreover, as the priorities of the schools are given, schools are not real strategic agents (i.e., they "play" truthfully) and all the participants know this.

3.1 Informational Settings

In each experimental session one of the different informational treatments is implemented for one of the three mechanisms. The four informational settings are the following:

- **Zero information setting:** Each participant knows her possible payoff amounts depending on the school where she holds a position (i.e., her own induced preferences), but not the other participants' preferences. She is only told that different participants might have different payoff tables.
Participants have no information about the schools' priority ordering in this treatment. They are only told each school's capacity (i.e., its number of vacancies).
- **Low information setting:** Besides her own induced preferences and the capacity of each school, each participant knows for which school she is the favorite candidate.
- **Partial information setting:** Each participant knows her own induced preferences, capacities, and the favorite candidates of all schools, up to their capacities.
- **Full information setting:** Each participant has complete information on both the induced preferences of all participants and the full priority ordering of schools over candidates.

In the case of the TTC mechanism, as the schools' priority orderings are reflected in the tentative assignments, under the low information treatment each participant is told—besides her own induced preferences and schools' capacities—her own tentative assignment, while in the partial information treatment each participant is aware of all the participants' tentative assignments and in the full information treatment she also knows the induced preferences of all the participants.

We conducted 12 sessions with undergraduate students from the Universitat Autònoma de Barcelona, recruited via e-mail using the web-based online recruitment system for Economic Experiments (ORSEE, Greiner, 2004),

where the experimental sessions (on paper/by hand) took place. In total, 435 subjects have participated in the experiment. Each treatment was implemented in a different session, therefore each subject was allowed to participate in one session only. In each session we had 45 subjects (9 groups) participating, except for the session where treatment B0 was implemented, where only 30 subjects (6 groups) participated. Subjects were informed that they would participate in a decision making task. At the beginning of each session, subjects were randomly seated at the tables and printed instructions were given to them. Once everybody was seated, instructions were read aloud and questions were answered privately. In each session subjects were asked to submit a school ordering, from their top to their last choice. Once everybody made her decision and the answer sheets were collected, subjects were asked to recall their submitted ordering and give a brief explanation of why they chose the given ranking.⁸ In the meanwhile, each participant’s final matching was determined and finally earnings were paid. Sessions lasted about 45-60 minutes and average net payments—including a 2 Euro show-up fee—were around 7,5 Euro.

The instructions and Decision Sheets in English for the GS mechanism can be found in the Appendix.⁹

4 Experimental Results

In this section we present our experimental results. Our aim is to analyze how the level of information participants hold affects the decision making process and consequently the properties of the three matching mechanisms. The first keypoint is related to whether individuals report their preferences truthfully. We investigate whether the amount of information given to participants influences the rate of truthful preference revelation (keeping the mechanism under analysis fixed) and, additionally, whether under the same informational setting truth-telling changes with the implemented mechanism. Second, given the reported preferences, we compare efficiency levels under each mechanism across informational settings, as well as across mechanisms for each information level. Finally, we examine the performance of each mechanism and the influence of the amount of information on the proportion of stable outcomes

⁸Subjects were not informed about this additional task in advance not to interfere with their decisions. On the other hand, this additional task was not compulsory (even though around 90% of the participants decided to do it).

⁹The instructions for the other mechanisms only differ in the description of the allocation method. The Decision Sheets for players in different roles look similar to the ones shown in the Appendix.

obtained.

4.1 Truthful Preference Revelation

Table 1 presents the proportion of participants who played truthfully (regarding induced preferences) and of those who used three possible kinds of preference manipulation.

Mechanism	Info	Not.	Truth	PSB	SSB	PSB&SSB	Other
Boston	Zero	B0	86,7%	-	10,0%	-	3,3%
	Low	B0.5	62,2%	11,1%	0,0%	15,6%	11,1%
	Partial	B1	46,7%	15,6%	13,3%	15,6%	8,9%
	Full	B2	46,7%	22,2%	6,7%	15,6%	8,9%
GS	Zero	GS0	82,2%	-	4,4%	-	13,3%
	Low	GS0.5	75,6%	4,4%	0,0%	6,7%	13,3%
	Partial	GS1	66,7%	13,3%	0,0%	20,0%	0,0%
	Full	GS2	66,7%	17,8%	2,2%	8,9%	4,4%
TTC	Zero	TTC0	95,6%	-	2,2%	-	2,2%
	Low	TTC0.5	82,2%	6,7%	0,0%	6,7%	4,4%
	Partial	TTC1	75,6%	4,4%	0,0%	11,1%	8,9%
	Full	TTC2	86,7%	6,7%	2,2%	2,2%	2,2%

Table 1: Proportion of truthful play and preference manipulations

The proportion of players who played truthfully varies between 46,7% and 95,6%, depending on the treatment being implemented. Note that, even though the GS and the TTC mechanisms are strategy-proof, there is some misrepresentation of preferences under these mechanisms. It is thus important to examine who manipulates the preferences and in which manner. Based on the explanations given by the subjects, we have identified three possible ways of preference manipulation. First, a substantial proportion of the participants has ranked the school where they have priority higher in the submitted ranking than it would be according to the induced preferences; this is what we call the "Priority School Bias" (PSB).¹⁰ The second identified way of manipulating the true preferences is to underrank the most competitive school (i.e., the school with only one vacancy); following Chen and Sönmez (2006) we name this form of behavior "Small School Bias" (SSB). The third manipulation method (PSB&SSB) is simply the simultaneous use of both previously described ways.¹¹ We can see that a relatively small proportion of

¹⁰In the case of the zero information setting it does not make sense to check for this kind of preference manipulation, as participants in these settings have no information about the priorities of the schools.

¹¹For the reasons above, in the zero information setting it does not make sense to check for this kind of manipulation.

the participants used the SSB method, although when comparing the average net payoff of the subjects who manipulate their preferences, we find that it yields a higher payoff (11 ECU) than the other two methods. The average payoffs obtained with the PSB and the PSB&SSB methods (8,7 ECU and 8,5 ECU respectively) suggest that these methods are used as means to ensure their second best payoff (9 ECU). This form of behavior can be considered as a kind of risk aversion, since it allows the participants to avoid their worst payoff (3 ECU). Our data also shows that preference manipulation yields higher payoffs under the Boston mechanism (9,9 ECU) than under any of the other two, strategy-proof, mechanisms (8,5 ECU).

In order to analyze whether the quantity of information participants hold affects behavior, we compare the proportion of participants playing truthfully in different informational settings under each mechanism.

Result 1: *Under each mechanism, having no information about the other parties' preferences results in a significantly higher proportion of subjects revealing their true preferences than having some additional information, while there is no statistically significant difference among the treatments where agents hold low, partial, and full information.*

Statistical evidence. Under each mechanism, the null hypothesis of equal proportions of truthful preference revelation across the four informational settings can be rejected at the 5% significance level. As the null hypotheses are rejected, multiple comparisons are made. Some differences in the proportion of truth-telling across information settings and the p -values of the significance t -tests of differences can be found in Table 2. For conciseness, we omit all tests involving the low information treatment, noting that the values obtained are very close to those of the partial information treatment.¹²

Mechanism	Boston			GS			TTC		
	B0-B1	B0-B2	B2-B1	GS0-GS1	GS0-GS2	GS1-GS2	TTC0-TTC1	TTC0-TTC2	TTC2-TTC1
Difference	40,0%	40,0%	0,0%	15,6%	15,6%	0,0%	20,0%	8,9%	11,1%
	(0,00)	(0,00)	(1,00)	(0,05)	(0,05)	(1,0)	(0,01)	(0,14)	(0,18)

Table 2: Difference in the proportion of truth-telling across information settings and p -values of the significance t -tests of differences

Result 2: *Across mechanisms there is a significant difference in the proportion of truth-telling. In particular, in the zero and the full information settings, the TTC mechanism performs better than the GS and Boston mechanisms and, in the low, partial, and full information settings, the proportion of truth-telling under both GS or TTC is higher than under the Boston mechanism.* ■

¹²This information is available upon request.

Statistical evidence. Under each informational settings, the null hypothesis of equal proportions of truthful preference revelation across the three matching mechanisms can be rejected at the 5% significance level. As the null hypotheses are rejected, multiple comparisons are made. Some differences in the proportion of truth-telling across mechanisms and the p -values of the significance t -tests of differences can be found in Table 3.¹³

Info	Zero			Partial			Full		
	B0-GS0	TTC0-B0	TTC0-GS0	GS1-B1	TTC1-B1	TTC1-GS1	GS2-B2	TTC2-B2	TTC2-GS2
Difference	4,4%	8,9%	13%	20,0%	28,9%	8,9%	20,0%	40,0%	20,0%
	(0,3)	(0,08)	(0,047)	(0,03)	(0,00)	(0,355)	(0,03)	(0,00)	(0,03)

Table 3: Difference in the proportion of truth-telling across mechanisms and p -values of the significance t -test of differences

To summarize our results regarding truth-telling, in lack of any information about the other participants' payoffs and preferences agents are much more likely to revert to truth-telling. This suggests that complete ignorance can be very convenient in this setting. Still, this does not make the TTC any less desirable: even under complete ignorance, the TTC mechanism outperforms the GS mechanism, which in turn results as successful as the Boston mechanism in what playing truthfully is concerned. In the settings where agents have additional information about the elements of the game, the TTC mechanism performs clearly better in revealing preferences straightforwardly than the Boston mechanism, and at least as well as the GS mechanism. ■

4.2 Efficiency

We now investigate mechanism efficiency in the different information scenarios. As we know, there is a strong link between preference manipulation and efficiency: even when the mechanism used is Pareto efficient—in terms of the revealed preferences—strategic behavior may lead to inefficient allocations.

In calculating efficiency levels we use the following definitions. The efficiency of a group of participants is calculated as the ratio of the sum of the actual earnings of the members of the group and the Pareto-optimal earnings of the group. The efficiency of a treatment is simply the average of the efficiency of all the groups. Table 4 shows the average efficiency of each treatment.

¹³Once more, results concerning the low information treatment, which resembles the partial information setting, were omitted but are available upon request.

Mechanism	Boston				GS				TTC			
Info setting	B0	B0.5	B1	B2	GS0	GS0.5	GS1	GS2	TTC0	TTC0.5	TTC1	TTC2
Efficiency	92.8%	87.4%	73.9%	80.7%	86.5%	88.4%	69.1%	78.7%	95.2%	87.4%	86.5%	91.3%

Table 4: Average efficiency

Our first question regarding efficiency is whether some informational settings yield higher efficiency levels than others. After this, we will turn our attention to the efficiency differences that may arise as a result of the different characteristics of the implemented mechanisms.

Result 3: *Under the Boston and GS mechanisms the amount of information has a significant effect on the average efficiency achieved by participants, while under the TTC mechanism average efficiency does not depend on the implemented information setting. In particular, under the Boston and GS mechanisms, having either no or a low level of information about the other parties’ preferences results in a significantly higher average efficiency than when participants hold partial information. On the other hand, there is no significant difference in the efficiency achieved under any mechanism between the partial and full information treatments.*

Statistical evidence. The pairwise differences in average efficiency across information settings and the p -values of the significance (permutation) tests of differences can be found in Table 5:¹⁴

Mechanism	Boston			GS			TTC		
	B0-B1	B0-B2	B2-B1	GS0-GS1	GS0-GS2	GS2-GS1	TTC0-TTC1	TTC0-TTC2	TTC2-TTC1
Difference	18.9%	12.1%	6.8%	17.4%	8%	9.7%	8.7%	3.9%	4.8%
	(0,005)	(0,05)	(0,288)	(0,044)	(0,317)	(0,203)	(0,194)	(0,524)	(0,536)

Table 5: Difference in the proportion of average efficiency across information settings and p -values of the significance (permutation) tests of differences

Under the TTC mechanism the null hypothesis of equal average efficiency across information settings can not be rejected. So, in this case—although there is a significant difference in the truthful preference revelation when comparing truthfulness between the zero and the partial information settings—there is no significant difference in the average efficiency of the four treatments at any reasonable significance level. ■

Result 4: *In the partial and in the full information treatments, the average efficiency under the TTC mechanism is significantly higher than under*

¹⁴The results concerning the low information treatment were again omitted, as they are similar to those obtained for the zero information treatment. All tests are available upon request.

any of the other mechanisms. On the other hand, there is no significant difference in average efficiency between the mechanisms in either the zero or low information settings.

Statistical evidence. Table 6 shows the differences in average efficiency across mechanisms and the p -values of the significance (permutation) tests of differences.¹⁵

Info	Zero			Partial			Full		
	B0-GS0	TTC0-B0	TTC0-GS0	B1-GS1	TTC1-B1	TTC1-GS1	B2-GS2	TTC2-B2	TTC2-GS2
Difference	6,3%	2,4%	9%	4,8%	12,6%	17,4%	1,9%	10,6%	12,6%
	(0,514)	(0,722)	(0,155)	(0,552)	(0,09)	(0,044)	(0,858)	(0,07)	(0,046)

Table 6: Difference in the proportion of average efficiency across mechanisms and p -values of the significance (permutation) tests of differences

Another way to evaluate the performance of a mechanism is to consider the proportion of participants who obtain their top choices. For this purpose, it might be tempting to use the reported preferences, so that a high fraction of participants receiving their submitted top choice would suggest that the mechanism performs well.¹⁶ However, as we have seen above, a substantial proportion of participants manipulate their preferences, so that in order to get a more accurate depiction of the performance of each mechanism in each information setting, it is worth to examine the proportion of participants who get their true top choices.

Result 5a: *Under the Boston and the GS mechanisms, in the low, partial, and full information settings, there is a highly significant difference between the proportion of participants who receive their reported top choice and those who receive their true top choice. Under the TTC this difference is only significant in the low and partial information setting.*

Statistical evidence. Table 7 shows the difference between the proportion of subjects who got assigned their reported top choice and of those who actually got their true top choice, as well as the corresponding p -values of

¹⁵The low information treatment resembles the zero information treatment and was, thus, excluded. All values are available upon request.

¹⁶From this point of view, the results in this experiment suggest that all mechanisms are similar: there is no statistically significant difference in the proportion of reported top choices obtained across mechanisms. Moreover, under each mechanism, there is no significant difference across information settings.

the t -tests of proportions.

Mechanism	Boston				GS				TTC			
	B0	B0.5	B1	B2	GS0	GS0.5	GS1	GS2	TTC0	TTC0.5	TTC1	TTC2
Reported top choice	83,3%	88,9%	73,3%	77,8%	71,1%	75,6%	64,4%	75,6%	77,8%	71,1%	77,8%	71,1%
True top choice	76,7%	53,3%	44,4%	46,7%	55,6%	53,3%	35,6%	44,4%	73,3%	53,3%	57,8%	66,7%
Difference	6,7%	35,6%	28,9%	31,1%	15,6%	22,2%	28,9%	31,1%	4,4%	17,8%	20,0%	4,4%
	(0,52)	(0,00)	(0,01)	(0,003)	(0,129)	(0,03)	(0,01)	(0,003)	(0,625)	(0,085)	(0,045)	(0,65)

Table 7: Proportions of the reported and the true top choices, differences between them, and the corresponding p -values of the significance t -tests of differences

While in the zero information setting there is no significant difference between the proportion of subjects who get their reported top choices and the ones who get their true top choices, in the low, partial, and full information settings, under the Boston and the GS mechanisms, the null hypothesis of equal proportions can be rejected in favor of the alternative hypothesis of the percentage of reported top choices being higher than the percentage of true top choices. Under the TTC mechanism, the same null can be rejected in the low and partial information settings. ■

Result 5b: *The amount of information has an important effect on the proportion of subjects who receive their true top choices. In particular, under the Boston and the GS mechanisms, a significantly higher number of participants get their true top choices in the zero information treatment than when holding some information on the elements of the game. The TTC mechanism is not sensitive to the amount of information and prevails over GS in what the assignment of the true top choices is considered.*

Statistical evidence. Table 8 reports the differences in the proportion of subjects who receive their true top choices across information settings (upper half) and across mechanisms (lower half), as well as the corresponding p -values of the significance t -tests for differences.

Mechanism	Boston			GS			TTC		
	B0-B1	B0-B2	B2-B1	GS0-GS1	GS0-GS2	GS2-GS1	TTC0-TTC1	TTC0-TTC2	TTC2-TTC1
Difference	32,2%	30,0%	2,2%	20,0%	11,1%	8,9%	15,6%	6,7%	8,9%
	(0,003)	(0,006)	(0,833)	(0,055)	(0,292)	(0,39)	(0,12)	(0,49)	(0,385)

Info setting	Zero			Partial			Full		
	B0-GS0	B0-TTC0	TTC0-GS	B1-GS1	TTC1-B1	TTC1-GS1	B2-GS2	TTC2-B2	TTC2-GS2
Difference	21,1%	3,3%	17,8%	8,9%	13,3%	22,2%	2,2%	20,0%	22,2%
	(0,052)	(0,74)	(0,076)	(0,39)	(0,205)	(0,033)	(0,833)	(0,053)	(0,032)

Table 8: Difference in the proportion of participants who got their true top choice, across information settings and across mechanisms, and p -values of the significance t -tests of differences

The above results indicate that the TTC is more efficient than both the GS and the Boston mechanisms, which confirms theory and contrasts with

the efficiency reversal result obtained in Chen and Sönmez (2006). In fact, while in Chen and Sönmez (2006) a substantial proportion of agents manipulate their preferences under TTC and, as a consequence, the GS emerges as more efficient, in this experiment many participants have recognized the strategy of truth-telling as dominant under TTC, leading to comparable rates of manipulation under both mechanisms. Another advantage of the TTC mechanism unveiled above is that the proportion of subjects who get their true top choice under the TTC is significantly higher than under the GS mechanism. On the other hand, efficiency levels and the fraction of subjects obtaining their true top choices under the TTC appear to be less dependent on the level of information participants hold, when compared to what is achieved under either Boston or GS. Finally, the Boston mechanism is more efficient than we could reasonably expect; this is due to the fact that a relatively high number of participants reveal their true preferences under this mechanism.

4.3 Stability

In what stability is concerned, a couple of remarks is in order. First, recall that the GS mechanism generates outcomes that are stable with respect to the submitted preferences, while neither the Boston nor the TTC mechanisms fulfil this property. Second, the stability of an outcome is evaluated for the true preferences of the participants and schools' priorities.¹⁷ Last, the average stability of a treatment is calculated simply as the proportion of stable outcomes observed among all the realized matchings in that treatment.

Our experimental data regarding the average stability of the treatments is shown in Table 9.

Mechanism	Boston				GS				TTC			
Info setting	B0	B0.5	B1	B2	GS0	GS0.5	GS1	GS2	TTC0	TTC0.5	TTC1	TTC2
Stable	16,7%	55,6%	11,1%	44,4%	55,6%	88,9%	77,8%	44,4%	11,1%	44,4%	22,2%	11,1%

Table 9: Proportion of stable outcomes

An interesting feature of our results on stability is that although we could expect the GS mechanism to yield the teachers' optimal outcome, in the partial information setting, 71% of the stable matchings turn out to be the

¹⁷The set of stable matchings is composed of the following outcomes: teachers 1 and 2 assigned to school B, 3 assigned to C, and 4 and 5 to A (the teachers' optimal stable matching); teacher 1 matched to B, 2 to A, 3 to C, 4 to A, and 5 to B; and finally 1 assigned to B, 2 to A, 3 to B, 4 to A, and 5 to C (the schools' optimal stable matching).

schools’ optimal stable matching. This can be explained on account of the substantial rate of manipulation under this treatment. Still, the theoretical stability ranking between GS and TTC appears to hold empirically, as stated in the following result.

Result 6: *The amount of information has no significant effect on the proportion of stable outcomes under any of the mechanisms. On the other hand, the GS mechanism results to be more successful than the TTC mechanism, but it only performs significantly better than the Boston mechanism in the partial informational scenario.*

Proof. Table 10 reports the differences in the proportion of stable matchings across informational treatments (upper half) and across mechanisms (lower half) and the p -values of the significance (Fisher’s exact) tests of differences.

Mechanism	Boston			GS			TTC		
	B0-B1	B2-B0	B2-B1	GS1-GS0	GS0-GS2	GS1-GS2	TTC1-TTC0	TTC0-TTC2	TTC1-TTC2
Difference	5,6%	27,8%	33,3%	22,2%	11,1%	33,3%	11,1%	0,0%	11,1%
	(1,0)	(0,58)	(0,29)	(0,62)	(1,0)	(0,19)	(1,0)	(1,0)	(1,0)

Info setting	Zero			Partial			Full		
	GS0-B0	B0-TTC0	GS0-TTC0	GS1-B1	TTC1-B1	GS1-TTC1	B2-GS2	B2-TTC2	GS2-TTC2
Difference	38,9%	5,6%	44,4%	66,7%	11,1%	55,6%	0,0%	33,3%	33,3%
	(0,17)	(1,0)	(0,06)	(0,01)	(1,0)	(0,03)	(0,68)	(0,29)	(0,15)

Table 10: Difference in average stability across information settings and across mechanisms, and p -values of the significance (Fisher’s exact) tests of differences

■

5 Concluding Remarks

In this paper we examine a particular class of matching problems that is closely related to the college admissions problem: the assignment of individuals to indivisible items. We analyze three well-known matching mechanisms—the Boston, the GS, and the TTC mechanisms—under different informational settings. Our experimental design allows us to explore two main questions.

First, we compare individuals’ decision making regarding truth-telling, efficiency, and stability across the three mechanisms, in each informational setting. These results may serve as a test of the theoretical characterization of the above mechanisms. Our results show that in both the zero and full information settings, under the TTC mechanism, a significantly higher proportion of participants plays truthfully than under either the Boston or the GS mechanisms. In case the participants have some information about the elements of the game, under the Boston mechanism a significantly higher number of participants manipulate their preferences than under either the

GS or the TTC mechanisms.¹⁸ That under complete ignorance we do not find a significant difference in truth-telling between the Boston and the GS mechanisms is specially interesting as under the latter mechanism straightforward behavior is dominant. Regarding efficiency, the experimental results are in accordance with the predictions of the theory in the partial and full information settings: here the TTC mechanism yields a significantly higher efficiency level than either the Boston or the GS mechanisms. On the other hand, when participants have low information about the others' preferences, there is no significant difference between the achieved efficiency across mechanisms.¹⁹ As for stability, the theoretical superiority of the GS mechanism is only confirmed in the partial information treatment. In low information treatments, even though the GS performs better than the TTC mechanism, it is as successful as the Boston mechanism, while under complete information there is no difference between the three mechanisms in what the proportion of stable outcomes is concerned.²⁰

Our second aim is to evaluate the influence of the amount of information held by individuals on the decision making process under the three matching mechanisms. The experimental results show that if participants have no information about the others' preferences they are more likely to play truthfully than when holding some information, while there is no significant difference in truth-telling under any of the mechanisms among settings where agents hold some information.²¹ The amount of information plays a role in the achieved efficiency level as well. Under the Boston and the GS mechanisms participants reach higher efficiency levels in low information settings than when holding partial information. While under the Boston mechanism the efficiency in the zero information setting is significantly higher than in the full information setting, the same difference under the GS mechanism is not significant. Under any of these two mechanisms, there is no significant difference between the partial and the full information case. On the other hand, under the TTC mechanism, the amount of information does not have a significant effect on the achieved efficiency level.²² Finally, in what stability

¹⁸Summarizing, the comparison of the proportions of truth-telling across mechanisms is the following: $B_0=GS_0<TTC_0$; $B_{0.5}<GS_{0.5}=TTC_{0.5}$; $B_1<GS_1=TTC_1$; $B_2<GS_2<TTC_2$.

¹⁹Summarizing, the efficiency levels across mechanisms are the following: $B_0=GS_0=TTC_0$; $B_{0.5}=GS_{0.5}=TTC_{0.5}$; $TTC_1>B_1=GS_1$; $TTC_2>B_2=GS_2$.

²⁰Summarizing, the results of stability across mechanisms are the following: $GS_0>TTC_0$, $GS_0=B_0$, $B_0=TTC_0$; $GS_{0.5}=B_{0.5}=TTC_{0.5}$; $GS_1>B_1=TTC_1$; $GS_2=B_2=TTC_2$.

²¹Summarizing, the proportion of truth-telling across informational settings is the following: $B_0>B_{0.5}=B_1=B_2$; $GS_0>GS_1=GS_2$, $GS_0=GS_{0.5}$, and $GS_{0.5}=GS_1$; $TTC_0=TTC_{0.5}$, $TTC_0>TTC_1$, $TTC_0=TTC_2$, and $TTC_{0.5}=TTC_1=TTC_2$.

²²Summarizing, the efficiency levels across informational settings are the following:

is concerned, the information does not determine the proportion of stable outcomes.²³

Summarizing, we can conclude that the comparison of the mechanisms points to the superiority of the TTC mechanism. Although regarding truthful preference revelation—depending on the implemented informational setting—it may give similar results to the GS mechanism, in the achieved efficiency level the TTC mechanism performs clearly better than either the Boston or the GS mechanisms. Moreover, we find that the amount of information may play a role in participants’ decision making. In general we can say that if the participants only know their own induced preferences, i.e., own payoffs, they are more likely to play truthfully than in case of having additional information. This results also in higher efficiency levels under all mechanisms, except for the TTC, where the amount of information has no significant effect on the achieved efficiency. Complete ignorance may thus be very convenient in this context.

6 Appendix

6.1 Instructions

6.1.1 Instructions for the Gale and Shapley Mechanism²⁴

Thank you for participating in this experiment on decision making. From now until the end of the session any communication with other participants is forbidden. If you have any question, feel free to ask at any point of the experiment. Please do so by raising your hand and one of us will come to your desk to answer your question.

You and every other participant play the role of a teacher and have to indicate a preference ordering over schools. We will form groups of five participants, so that you will be grouped with 4 other participants, whose identity you will not know. There are 3 schools (A, B, and C) and 5 teaching positions available across them: two positions at schools A and B, and one position at school C. Each of the five positions will be allocated to a participant, based on the preference orderings submitted by the 5 participants of the group. Besides differing in size (number of teaching positions), schools differ in location and quality. The desirability of schools in terms of location and

$B0 > B2 = B1$, $B0 = B0.5$, $B0.5 > B1$, and $B0.5 = B2$; $GS0 = GS0.5 > GS1$, $GS0.5 > = GS2$ and $GS0 = GS2$ and $GS1 = GS2$; $TTC0 = TTC0.5 = TTC1 = TTC2$.

²³Our results on stability across information settings are: $B0 = B0.5 = B1 = B2$; $GS0 = GS0.5 = GS1 = GS2$; $TTC0 = TTC0.5 = TTC1 = TTC2$.

²⁴These instructions correspond to the Full information condition.

quality is summarized in the amounts shown in the payoff table (see Decision Sheets), which contains the payoff amounts in experimental currency units (ECU) corresponding to each participant and school position. This matrix is known by all the participants.

Submitted school ranking: During the experiment you will be asked to complete the Decision Sheet by indicating the preference ordering over schools you wish to submit. You have to rank every school. Once all participants have completed their Decision Sheets, the experiment is finished.

Priority ordering of schools: Schools when offering positions consider the quality of each applicant and the experience they have. On this basis, they build a priority ordering where all candidates are ranked. The following table contains the priority ordering of each school (this is known by all participants):

	School A	School B	School C
1st choice	4	3	5
2nd choice	2	1	4
3rd choice	1	4	2
4th choice	5	5	3
5th choice	3	2	1

Payoffs: During the session you can earn money. You will receive 4 ECU for your participation, in addition to the amount you earn in the experiment. This amount is displayed in the payoff matrix, corresponding to the position you hold at the end of the session. Note that the position you hold at the end of the experiment depends on your submitted ordering and the submitted ordering of the other participants of your group (which you do not know at the moment of submitting your order).

Once the experiment has finished and the allocations of the participants are determined, each participant will get paid her total payoff in euros. One ECU equals 0,5 Euros.

Allocation Method: With this method, each participant is assigned a position at the best possible school reported in her Decision Sheet that is consistent with the priority ordering of schools. Given the submitted preferences of the participants and the priority ordering of each school, positions are allocated in the following way:

- An application to the first ranked school in the Decision Sheet is sent for each participant.
- Each school accepts the applicants with higher priority order until positions are filled, and keep them on hold, while rejects the lowest priority students in excess of its capacity. Throughout the allocation process, a school can hold no more applications than its number of positions!
- Whenever an applicant is rejected at a school, an application is sent to the next highest school on his Decision Sheet.
- Whenever a school receives a new application (from an applicant that has been rejected in a previous round by a better ranked school), these applications are considered together with the (previously) retained applications for that school. Among the retained and new applicants, the lowest priority ones in excess of the number of the positions are rejected, while remaining applications are retained.
- This process is repeated until no more applications can be rejected, and the allocation is finalized; and each participant is assigned the position at the school that holds her application at the end of the process.

An Example: We will go through a simple example to illustrate how the allocation method works.

Applicants and schools: In this example there are four applicants (1 – 4) and three schools (A, B, C).

Positions: There are two positions at school B, and one each at A and C.

Submitted school ranking: Suppose the submitted school rankings of each participant are the following:

	Applicant 1	Applicant 2	Applicant 3	Applicant 4
1st choice	A	B	C	C
2nd choice	C	A	B	A
3rd choice	B	C	A	B

Schools' priority ordering: Suppose the priority orderings of the schools are:

	A	B	C
1st choice	2	1	1
2nd choice	4	2	3
3rd choice	1	4	4
4th choice	3	3	2

Allocation: This allocation method consists of the following rounds:

ROUND 1: Each applicant applies to her first choice:

- Applicant 1 applies to School A, 2 to School B, and Applicants 3 and 4 to School C.
- School A retains Applicant 1, School B retains Applicant 2; and School C retains Applicant 3;
- School C rejects Applicant 4, as it only has one position, and School C prefers Applicant 3 to 4.

ROUND 2: The applicant who is rejected in ROUND 1 (4) applies to her second choice:

- Applicant 4 applies to School A.
- School A compares Applicant 1 (retained in round 1) and 4, as it only has one position free; and retains 4 and rejects now 1 (as in School A's preference ordering 4 has priority over 1).

ROUND 3: The applicant who is rejected in ROUND 2 (1) applies to her second choice:

- Applicant 1 applies to School C.
- School C compares Applicant 3 (retained in round 1) and 1, as it only has one position free. School C retains 1 and rejects now 3 (as in School C's preference ordering 1 has priority over 3).

ROUND 4: The applicant who is rejected in ROUND 3 (3) applies to her second choice:

- Applicant 3 applies to School B.
- School B retains Applicant 2 since the first round, but still has a vacancy (as here there are two positions opening), therefore School B accepts Applicant 3.

Here the process finishes, as there are no more rejections; and the final allocations are:

Applicant	1	2	3	4
School	C	B	B	A

You will have 15 minutes to go over the instructions at your place, and make your decision. Are there any questions?

6.1.2 Instructions for the TTC Mechanism²⁵

(...)

Priority ordering of schools: Schools when offering positions consider the quality of each applicant and the experience they have. On this basis, each candidate is tentatively assigned to a school. This tentative assignment is the following:

Applicant	1	2	3	4	5
School	B	A	B	A	C

Allocation Method: In this process, initially each participant is tentatively assigned to one of the opening positions and all participants are ordered in a queue based on a fair lottery. This means that each participant has an equal chance of being first in the queue, second, . . . , as well as the last in the queue. Given the submitted preference orderings of the participants and the order in the queue determined by the lottery, the allocation process is the following:

- An application to the first ranked school in the Decision Sheet is sent for the participant at the top of the queue.
 - If the application is submitted to the school to which this participant was assigned initially, then her tentative assignment becomes her final position; and this participant and his position are removed from the subsequent process. The process continues with the next participant in the queue.

²⁵In this Appendix we only include those parts of the original instructions that differ from the instructions of the Gale-Shapley mechanism.

- If the application is submitted to another school, say school S, then the first participant who tentatively holds a position at school S is moved to the top of the queue, directly in front of the requester.
- Whenever the queue is modified, the process continues in the way described above. Now an application to the first ranked school in the Decision Sheet is sent for the (new) participant at the top of the queue.
 - If the application is submitted to the school to which this participant was assigned initially, etc. . .
 - If the application is submitted to another school, etc. . .
- A mutually-beneficial exchange is obtained when a cycle of applications is made in sequence, which benefits all affected participants; e.g., A applies to B’s tentative position, B applies to C’s tentative position, and C applies to A’s tentative position. In this case the exchange is completed and all three participants as well as their assignments are removed from the subsequent process.
- The process continues till all participants are assigned a position.

An Example: (…)

Priority queue of applicants: Suppose the lottery gave the following priority ordering: 1 – 2 – 3 – 4.

Tentative assignment: Suppose the initial (tentative) assignment of positions is the following:

Applicant	1	2	3	4
School	B	C	B	A

Allocation: The allocation method consists of the following process:

- Step 1: The first applicant in the queue (1) applies to her best choice, to School A, however, the only position here is tentatively held by participant 4. So participant 4 is moved to the top of the queue.
- Step 2: The new queue is now 4 – 1 – 2 – 3. Participant 4 ranked School C as her top choice, but the only position of this school is tentatively held by participant 2. Therefore 2 is moved to the top of the queue.

Step 3: The new queue is now 2 – 4 – 1 – 3. Participant 2 ranked School B as her top choice, but the two positions at school B are tentatively held by participant 1 and 3. As 1 has priority over 3 (as she is in front of 3 in the queue), participant 1 is moved to the top of the queue.

Step 4: The new queue is now 1 – 2 – 4 – 3. Remember, that applicant 1 has ranked School A as her best choice. A cycle of participants is now made in sequence in the last three steps: 1 applied to the tentative assignment of 4, 4 applied to the tentative assignment of 2, and 2 applied to the tentative assignment of 1. These mutually beneficiary changes are made: 1 gets the position in School A, 2 gets one of the two positions in School B, and 4 gets the position in School C. These participants and their assignments are removed from the process.

Step 5: The only participant left to be assigned is 3. As the only school with available position is School B and this position is tentatively assigned to her, it becomes her final assignment. The allocation process ends.

The final allocations are:

Applicant	1	2	3	4
School	A	B	B	C

(...)

6.1.3 Instructions for the Boston Mechanism²⁶

(...)

Allocation Method: Given the submitted preferences of the participants and the priority ordering of each school, positions are allocated in the following way (in max. 3 rounds):

- ROUND 1:
- An application to the first ranked school in the Decision Sheet is sent for each participant.
 - Each school accepts the participants with higher priority order until positions are filled. These

applicants and their positions are removed from the system. All other applications are rejected by the schools.

²⁶In this Appendix we only include those parts of the original instructions that differ from the instructions of the Gale-Shapley mechanism.

- ROUND 2: – The applicants remaining in the system send the application to their second ranked position in the Decision Sheet.
- If a school still has available positions remaining from Round 1, then it accepts the applicant with higher priority order until all positions are filled. The remaining applications are rejected.

ROUND 3: Each remaining participant is assigned a position at her last choice.

An Example: (...)

Submitted school ranking: Suppose the school rankings submitted by each participant are the following:

	Applicant 1	Applicant 2	Applicant 3	Applicant 4
1st choice	A	A	A	B
2nd choice	B	C	C	C
3rd choice	C	B	B	A

Schools' priority ordering: Suppose the priority orderings of the schools are:

	A	B	C
1st choice	1	1	2
2nd choice	2	3	3
3rd choice	3	4	4
4th choice	4	2	1

Allocation: This allocation method consists of the following rounds:

ROUND 1:

- Each applicant applies to her first choice:
 - Applicants 1, 2, and 3 apply to School A, Applicant 4 to School B.
 - School A accepts Applicant 1 (its first choice)
 - School B accepts Applicant 4
- Accepted applicants (1 and 4) and schools without remaining positions (School A) are removed from the subsequent process.

ROUND 2:

- Each applicant who is rejected in ROUND 1 (2 and 3) applies to her second choice:
 - Applicants 2 and 3 apply to School C.
 - School C accepts Applicant 2 (its first choice)
- Accepted applicants (2) and schools without remaining positions (School C) are removed from the subsequent process.

ROUND 3:

- Each remaining applicant who is rejected in the previous rounds (3) is assigned her last choice:
 - Applicant 3 gets the remaining position in School B.

Based on this method, the final allocations are:

Applicant	1	2	3	4
School	A	C	B	B

(...)

6.2 Decision Sheets

6.2.1 Decision Sheet for GS Mechanism Under the Zero Information Treatment²⁷

You are participant 2.

Recall: Your payoff amount depends on the school position you hold at the end of the experiment. Your possible payoff amounts are outlined in the following table:

Position received at school	A	B	C
Your payoff	9	15	3

This means, that if at the end of the experiment you hold a position:

- – at school A, you will be paid 9 ECU;

²⁷The same Decision Sheet was used in the corresponding information treatment of the Boston mechanism as well.

- at school B, you will be paid 15 ECU;
- at school C, you will be paid 3 ECU.

Recall: Different participants might have different payoff tables.

Recall: There are two positions opening at schools A and B, and one at school C.

Priority ordering of schools: You have no information about the preferences of the schools. This means, that if at the end of the experiment you hold a position:

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank EVERY school!

1st choice	2nd choice	3rd choice

This is the end of the experiment. Please wait until everybody finishes and you are told your result.

6.2.2 Decision Sheet for GS Mechanism Under the Partial Information Treatment²⁸

You are participant 2.

Recall: Your payoff amount depends on the school position you hold at the end of the experiment. Your possible payoff amounts are outlined in the following table:

Position received at school	A	B	C
Your payoff	9	15	3

This means, that if at the end of the experiment you hold a position:

- – at school A, you will be paid 9 ECU;
- at school B, you will be paid 15 ECU;

²⁸The same Decision Sheet was used in the corresponding information treatment of the Boston mechanism as well.

– at school C, you will be paid 3 ECU.

Recall: Different participants might have different payoff tables.

Recall: There are two positions opening at schools A and B, and one at school C.

Priority ordering of schools: Information about the priorities:

- – at school A participant 2 (you) and 4 have priority;
- at school B participant 1 and 3 have priority;
- at school C participant 5 has priority.

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank **EVERY** school!

1st choice	2nd choice	3rd choice

This is the end of the experiment. Please wait until everybody finishes and you are told your result.

6.2.3 Decision Sheet for GS Mechanism Under the Full Information Treatment

You are participant 2.

Recall: Each participant’s payoff amount depends on the school position she holds at the end of the experiment. The possible payoff amounts for each participant are known by everybody and these are outlined in the following table:

Position received at school	A	B	C
Payoff of participant 1	3	9	15
Payoff of participant 2 (YOU)	9	15	3
Payoff of participant 3	15	3	9
Payoff of participant 4	15	3	9
Payoff of participant 5	15	9	3

This means, that for example, if at the end of the experiment:

- participant 1 and participant 5 hold a position at school A, participant 3 and 4 hold a position at school B, and you hold a position at school C, the payoffs would be the following:
 - participant 1 would be paid 3 ECU; participant 5 would get 15 ECU; participant 3 would get 3 ECU; participant 4 would get 3 ECU; and you would get 3 ECU.

Recall: There are two positions opening at schools A and B, and one at school C.

Priority ordering of the schools: The complete priority ordering of the schools is known by each participant, and is shown in the following table:

	School A	School B	School C
1st choice	4	3	5
2nd choice	2	1	4
3rd choice	1	4	2
4th choice	5	5	3
5th choice	3	2	1

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank **EVERY** school!

1st choice	2nd choice	3rd choice

This is the end of the experiment. Please wait until everybody finishes and you are told your result.

7 Acknowledgements

We wish to thank Jordi Brandts, Pablo Guillen, Flip Klijn, Jordi Massó, Klara Pinter, Fernanda Rivas, Utku Ünver, Róbert F. Veszteg, and two anonymous referees for their insightful comments. Pais acknowledges the financial support from the Fundação para a Ciência e a Tecnologia under grant SFRH/BD/5308/2001 and from the spanish Ministerio de Ciencia y Tecnología under research grant BEC2002-02130. Pintér acknowledges the

financial support from the spanish Ministerio de Educación, Cultura y Deporte under the grant Beca de Postgrado para la Formación de Profesorado Universitario (Ref.: AP2001-0923) and from the Universidad Autònoma de Barcelona under the Beca de Formació i Support a la Recerca (Ref: BEC2003-01132).

References

- [1] Abdulkadiroglu, Atila and Sönmez, Tayfun (1999): "House Allocation with Existing Tenants" *Journal of Economic Theory* **88.**, 2, pp.233-260
- [2] Abdulkadiroglu, Atila and Sönmez, Tayfun (2003): "School Choice: A Mechanism Design Approach" *American Economic Review* **93.**, 3, pp.729-747
- [3] Abdulkadiroglu, Atila, Pathak, P., Roth. A., and Sönmez, Tayfun (2005): "The Boston Public School Match" *American Economic Review, Papers and Proceedings*, forthcoming
- [4] Barberà, S. and Dutta, B. (1995): "Protective Behavior in Matching Models" *Games and Economic Behavior* **8.**, pp.281-296
- [5] Chen, Yan and Sönmez, Tayfun (2002a): "Improving Efficiency of On-Campus Housing: An Experimental Study" *American Economic Review* **92.**, 5, pp.1669-1686
- [6] Chen, Yan and Sönmez, Tayfun (2002b): "An Experimental Study of House Allocation Mechanisms" *Economics Letters* **83.**, 1, pp.137-140
- [7] Chen, Yan and Sönmez, Tayfun (2006): "School Choice: An Experimental Study" *Journal of Economic Theory* **127**, pp. 202-231
- [8] Conover, W.J. (1980): *Practical Nonparametric Statistics*, New York: John Wiley & Sons
- [9] Dubins, L. and Freedman, D. (1981): "Machiavelli and the Gale-Shapley Algorithm" *American Mathematical Monthly* **88.**, pp.485-494
- [10] Ehlers, L. (2003): "In Search of Advice for Physicians in Entry-Level Medical Markets" *Mimeo, Université de Montréal*
- [11] Ehlers, L. (2004): "In Search of Advice for Participants in Matching Markets Which Use the Deferred-Acceptance Algorithm" *Games and Economic Behavior* **48.**, pp.249-270

- [12] Ehlers, L. and Massó, J. (2003): "Incomplete Information and Small Cores in Matching Markets" *Mimeo, Universitat Autònoma de Barcelona*
- [13] Ergin, Haluk and Sönmez, Tayfun (2004): "Games of School Choice under the Boston Mechanism" *Journal of Public Economics* **90.**, pp.215-237
- [14] Gale, D. and Shapley, L. (1962): "College Admissions and the Stability of Marriage" *American Mathematical Monthly* **69.**, pp.9-15
- [15] Greiner, Ben (2004): "An Online Recruitment System for Economic Experiments" In Kurt Kremer, Volker Macho (Eds.): *Forschung und wissenschaftliches Rechnen 2003.*, pp. 79-93, GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung
- [16] Harrison, G. and McCabe, K. (1996): "Stability and Preference Distortion in Resource Matching: An Experimental Study of the Marriage Problem" In V. Smith (ed.) *Research in Experimental Economics* **6.**, pp.53-129, JAI Press
- [17] Haruvy, E., Roth, A., and Ünver, U. (2001): "Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market" *Journal of Economic Dynamics and Control*, forthcoming
- [18] Kagel, J and Roth, A. (2000): "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment" *Quarterly Journal of Economics* **115(1).**, pp.201-235
- [19] McKinney, C., Niederle, M., and Roth, A. (2005): "The Collapse of a Medical Labour Clearinghouse (and Why Such Failures Are Rare)" *American Economic Review* **95.**, pp.878-889
- [20] Nalbantian, H. and Schotter, A. (1995): "Matching and Efficiency in the Baseball Free-Agent System: An Experimental Examination" *Journal of Labour Economics* **13.**, pp.1-31
- [21] Olson, M. and Porter, D. (1994): "An Experimental Examination Into Design of Decentralized Methods to Solve the Assignment Problem With and Without Money" *Economic Theory* **4.**, pp.11-40
- [22] Roth, A. (1982a): "The Economics of Matching: Stability and Incentives" *Mathematics of Operations Research* **7.**, pp.617-628

- [23] Roth, A. (1982b): “Incentive Compatibility in a Market With Indivisibilities” *Economics Letters* **9.**, pp.127-132
- [24] Roth, A. (1989): “Two-Sided Matching With Incomplete Information about Others’ Preferences” *Games and Economic Behavior* **1.**, pp.191-209
- [25] Roth, A. and Rothblum, U. (1999): “Truncation Strategies in Matching Markets - In Search of Advice for Participants” *Econometrica* **67.**, pp.21-43
- [26] Roth, A. and Sotomayor, M. (1990): *Two Sided Matching: a Study in Game Theoretic Modeling and Analysis*, Econometric Society Monograph Series. New York: Cambridge University Press
- [27] Shapley, L. and Scarf, H. (1974): “On Cores and Indivisibility” *Journal of Mathematical Economics* **1.**, pp.23-28
- [28] Ünver, U. (2001): “On the Survival of Some Two-Sided Matching Mechanisms: An Experimental and Computational Investigation of the Stability Hypothesis” *Mimeo, Koç University*
- [29] Ünver, U. (2005): “On the Survival of Some Two-Sided Matching Mechanisms” *International Journal of Game Theory* **33.**, pp.239-254