

9-2010

Matching Markets with Mixed Ownership: The Case for A Real-life Assignment Mechanism

Pablo Guillen
University of Sydney

Onur Kesten
Carnegie Mellon University, okesten@andrew.cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/tepper>

 Part of the [Business Commons](#)

Matching Markets with Mixed Ownership: The Case for A Real-life Assignment Mechanism*

Pablo Guillen[†] Onur Kesten[‡]

September 2010

Abstract

We consider a common indivisible good allocation problem in which agents have both social and private endowments. Popular applications include student assignment to on-campus housing, kidney exchange, and particular school choice problems. In a series of experiments Chen and Sönmez (*American Economic Review* 92: 1669-1686, 2002) have shown that a popular mechanism from recent theory, the Top Trading Cycles (TTC) mechanism, induces a significantly higher participation rate by agents with private endowments and leads to significantly more efficient outcomes than the most commonly used real-life mechanism, the Random Serial Dictatorship with Squatting Rights.

We first show that a particular mechanism, the so-called New House 4 (NH4) mechanism, which has been in use at MIT since the 1980s, is in fact outcome-equivalent to a natural

*We would like to thank Harvard Business School and, in particular, Al Roth for their hospitality and financial support of this project. We are grateful to Jan Eeckhout (Editorial Board), and two anonymous referees of the journal for very useful suggestions that significantly improved the paper. We also thank Al Roth, Yan Chen, Tayfun Sönmez, Thomas Palfrey, Guillaume Frechette, Fuhito Kojima, Yosuke Yasuda, Utku Ünver, Jordi Massó, Vasilis Sarafidis, and seminar/conference participants at Harvard, Pittsburgh, Rochester, Boston College, Columbia, Sydney, SCW 2006 (Istanbul), RED 2006 (Bodrum), and Matching Workshop 2007 (Barcelona) for helpful comments, and Siu Kho for software assistance. This paper supersedes a working paper entitled “On Campus Housing: Theory and Experiment” (Guillen and Kesten, 2009). All remaining errors are our own.

[†]Room 340, Discipline of Economics, Merewether Building (H04), The University of Sydney, NSW 2006, Sydney, Australia. E-mail: pablo.guillenalvarez@sydney.edu.au. Tel: +61 2 9036 9188.

[‡]Corresponding author. Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA. E-mail: okesten@andrew.cmu.edu. Tel:+1-412 268 9823. Fax: +1-412 268 7064.

adaptation of the well-known Gale-Shapley mechanism of two-sided matching theory. This implies that the NH4 mechanism is the most efficient mechanism within the class of fair and individually rational mechanisms, and that it is essentially the only incentive compatible mechanism satisfying the two properties. We then experimentally compare NH4 and TTC. We find that under NH4, the participation rate is significantly higher than under TTC. We also propose a new efficiency test based on ordinal preference information and show that NH4 also outperforms TTC in terms of efficiency.

Keywords: Matching; House allocation; Priority; Deferred acceptance

JEL classification: C78, C79, D61, D78, I20

1 Introduction

In this paper we study the problem of assigning a set of indivisible goods, without explicitly using monetary transfers, to a set of agents some of whom may also have prior claims to some of the goods. More specifically, there exist two sets of objects: those that are socially owned and those that are privately owned. Agents have strict preferences over the objects. Three canonical examples¹ of this problem are *on-campus housing*, *kidney exchange*, and *school choice* for particular settings.

Many colleges offer on-campus housing opportunities to incoming as well as to already enrolled students. In *on-campus housing*, the goal is to allocate dormitory rooms to students at college campuses.² Students consist of incoming freshmen and more senior students. Incoming freshmen do not initially occupy any rooms, and each senior student is the occupant of a room from the previous year. There are also vacant rooms which used to be occupied by the previous year's graduating class and have now become available for reallocation. Senior students are entitled to keep their room but may also apply for a new one. Another ingredient in this application

¹Other examples of this problem include the assignment of tasks to workers, offices to faculty, and parking spaces to commuters.

²Clearly, there are several other examples in different contexts, where houses are allocated to prospective tenants using non-market mechanisms. Subsidized public housing like "Council Houses" in the U.K. or "Housing Commission" dwellings in Australia are a few of such examples.

is an exogenously given priority ordering of agents, which is usually determined according to the assignment policies of the central clearinghouse of the particular college, and could be, for example, based on seniority, GPA, the result of a lottery, or a combination of these.

The preferred treatment for the most serious forms of kidney disease is transplantation. Since there is a significant shortage of deceased-donor kidneys compared to demand,³ and because a healthy person can remain healthy on only one kidney, transplantation from a live-donor is also quite common.⁴ Nevertheless, a willing donor may not always be able to donate to her intended patient due to blood-type or immunological incompatibilities. Rapaport (1986) proposed the idea of exchanging donors between two incompatible pairs if the donor of one pair can feasibly donate to the patient of the other pair. Until recently, however, feasible exchanges were sought in an unorganized and decentralized way in most parts of the United States. In a series of influential papers,⁵ economists Alvin E. Roth, Tayfun Sönmez, and Utku Ünver have proposed innovative ideas to implement kidney exchanges in an organized way through the lens of mechanism design: In *kidney exchange*, there are patients initially paired with incompatible donors (the analogues in this context of “senior students with occupied rooms” in on-campus housing) who wish to receive a compatible kidney, and patients without any donors (the analogues in this context of “incoming freshmen” in on-campus housing). There are also options such as altruistic donors⁶ and priority on the deceased donor waiting list (the analogues in this context of “vacant rooms” in on-campus housing). In this application, a priority ordering of patients is determined based on the seriousness of their medical conditions or their waiting-times. The mechanisms and their properties which we discuss in this paper have immediate counterparts in the context of kidney exchange.

Another common problem in the United States and elsewhere arises when a school district offers students the option to attend public schools other than their neighborhood schools. In *school choice*, each student submits a rank-ordered-list of schools to a centralized clearing-house

³For example, in 2002 there were about 3,400 patients who died while on the waiting list. In the same year, there were another 900 patients who became too ill to be eligible for transplantation.

⁴For example, in 2004 there were 6,086 live-donor transplants in the U.S.

⁵See Roth, Sönmez, and Ünver (2004, 2005ab, 2007). Also see Al Roth’s Market Design blog for more recent developments on the current kidney exchange practices.

⁶An altruistic donor is a non-directed living donor, also known as a *Good Samaritan donor*.

(e.g., school district), which then determines students’ assignments based also on school-specific priorities. While other criteria may also play a role in determining school-specific priorities in general,⁷ some places use residence exclusively as the sole determinant of priority.⁸ For example, the Tokyo school district first gives each student the option to attend her neighborhood school (the analogue idea in this context of “a senior student keeping the room she initially occupies” of on-campus housing); if the student chooses otherwise, then she participates in the centralized assignment process.⁹

In the remainder of the paper, for convenience, we shall use the on-campus housing application as our running example. In keeping with this, we refer to the indivisible goods that are socially owned as *vacant houses* and to those that are privately owned as *occupied houses*. Each agent is either a *newcomer* who does not currently own a house, or an *existing tenant* who currently owns an occupied house and may be seeking a better one. Given the fixed resources, a *house allocation problem (with existing tenants)* is characterized by two pieces of information: (1) a *priority ordering* over all agents; and (2) a *list of preferences* of each agent over houses, typically a rank-ordered-list of houses that each agent decides on by comparing different housing types available. An assignment *mechanism* is a systematic procedure that chooses an assignment of agents to available houses based on the two pieces of information. The present paper takes a theoretical as well as an experimental approach to this problem.

By and large, the performance of a mechanism is evaluated along four merits: (1) *individual rationality* (i.e., an existing tenant should be encouraged to participate by giving her the guarantee of obtaining a house that is no worse than her occupied house); (2) *efficiency* (i.e., resources should be optimally allocated according to the preferences of agents); (3) *fairness* (i.e., the assignment should respect the priority order); and (4) *incentive compatibility* (i.e., each agent should be

⁷In Boston, for example, the priorities are assigned based on walk zone and whether the student has a sibling enrolled at the particular school. A random lottery draw is also commonly used to break ties within priority groups.

⁸In most places in the U.S., the highest priority for a particular school is given to those students who reside within the walk zone of the school. However, even in places where “proximity” is not the only determinant of priority, a student often has a high chance of admission at her neighborhood school should she decide to list it as her first choice. In this sense, most students are typically entitled to attend their neighborhood schools if they wish to (unless perhaps if these schools are extremely popular). Hence, even more general school choice problems can be viewed as similar to the present problem.

⁹We thank Yosuke Yasuda and Fuhito Kojima for bringing this example to our attention.

induced to act straightforwardly and reveal her true preferences). Among these criteria, *individual rationality* is of critical importance. This is because lack of participation by existing tenants can entail a loss of potentially large gains from trade. Furthermore, ensuring high participation by existing tenants not only is appealing in the context of on-campus housing but is also vital for the sustenance and proper-functioning of a kidney transplant center. Indeed, it may not be reasonable to expect a kidney patient to commit to a transplant that has a smaller chance of success than the one through her own intended donor.¹⁰

A widely used mechanism is the so-called *random serial dictatorship with squatting rights (RSD)*.¹¹ RSD works as follows: First, each existing tenant decides whether she will enter the house assignment process, or keep her occupied house. The houses of those who choose to enter become available for allocation together with the vacant houses. Second, the centralized clearing house, e.g., the housing office, randomly determines the priority ordering of all participating agents from a given distribution (which may simply be uniform, or may favor some students because of a specific university policy). Finally, available houses are allocated to agents based on the priority ordering: The first agent is assigned her top choice, the second agent is assigned her top choice among the remaining houses, and so on.

While quite popular, RSD has a serious shortcoming: It is not *individually rational*. Since it cannot guarantee an existing tenant a house at least as good as her current one, existing tenants may show reluctance to participate under such a mechanism, which in turn may result in efficiency losses. Abdulkadiroglu and Sönmez (1999), arguing that some of these efficiency losses can be recovered if existing tenants are instead allowed to trade their houses through a market-like procedure, proposed the *top trading cycles (TTC) mechanism*.¹² TTC is *individually rational*, *Pareto efficient*, and *incentive compatible* (Abdulkadiroglu and Sönmez, 1999). Remarkably, a follow-up experimental study by Chen and Sönmez (2002) found TTC to lead to higher participation rates

¹⁰Given the obvious importance of the size of the central kidney pool, participation is probably one of the most important considerations in the kidney exchange context with many lives at stake.

¹¹Some examples include undergraduate housing at the University of Pennsylvania, Carnegie Mellon, Duke, Northwestern, and the University of Michigan.

¹²The idea of top trading cycles was originally proposed by David Gale (Shapley and Scarf, 1974). Also see Sönmez and Ünver (2005) for a study on the relationship between between TTC and the core; and Sönmez and Ünver (2010) for an interesting characterization of TTC.

and to more efficient outcomes than the widely used RSD.

After being advocated as a promising school choice mechanism by economists Atila Abdulkadiroğlu, Parag Pathak, Alvin E. Roth, and Tayfun Sönmez, the well-known *Gale-Shapley deferred acceptance mechanism* of two-sided matching theory has gained increasing popularity among school districts in the United States, and has recently replaced two deficient mechanisms in New York City (Abdulkadiroğlu, Pathak, and Roth, 2005) and Boston (Abdulkadiroğlu, Pathak, Roth, and Sönmez, 2005). Motivated by the success of the Gale-Shapley mechanism in numerous¹³ other matching markets, this paper offers a natural and intuitive adaptation of the Gale-Shapley mechanism to house allocation. Our main result shows that this adaptation of the Gale-Shapley (GS) mechanism is indeed equivalent to a real-life mechanism, the so-called *New House 4 (NH4)*, which has already been in use at Massachusetts Institute of Technology (MIT) for about three decades (Theorem 1). This result has important implications. First, it shows that NH4 fulfills important desirable properties such as (*dominant strategy*) *incentive compatibility*, *individual rationality*, and *fairness* (Corollary 1). Second and more critically, it shows that NH4 is Pareto superior to any *individually rational* and *fair* mechanism (Corollary 2). Third, it shows that NH4 is *essentially* the unique *incentive compatible*, *individually rational*, and *fair* mechanism (Corollary 3).

We show that no mechanism can be *individually rational*, *Pareto efficient*, and *fair* (Proposition 1). The three mechanisms we discussed, however, serve to show that any other three properties can be satisfied simultaneously.¹⁴ From a theoretical point of view, TTC has the edge in terms of efficiency and NH4 in terms of fairness. Despite their theoretical appeal, it is important to know how the two mechanisms perform in controlled laboratory experiments in order to make better-informed policy recommendations. We designed a laboratory experiment to compare the performance of NH4 with that of the leading theory mechanism TTC, with specific attention to the participation decisions of existing tenants and the efficiency of the outcomes. In terms of existing

¹³Probably one of the best-known such markets is the redesigned National Resident Matching Program (Roth and Peranson, 1999) for assigning new physicians to hospital positions. Also see Roth and Rothblum (1999) for an extensive list of these markets.

¹⁴Hence, RSD is Pareto efficient, fair, and incentive compatible. The RSD allocation is Pareto efficient only when restricted to the set of agents who participate in the assignment process, which is the efficiency criterion we adopt throughout the paper. It may, however, be inefficient without such a restriction.

tenant participation, we found that under NH4 the participation rate is significantly higher than under TTC.

When making efficiency comparisons, the matching literature has thus far relied on only cardinal measures of efficiency.¹⁵ However, this is not directly in line with the theory because most mechanisms elicit only ordinal preferences. Furthermore, the commonly used efficiency concept in the theoretical literature is *Pareto efficiency* with respect to ordinal preference information, whereas the common efficiency tests are based on the payoffs of the subjects. In this paper, we also make a methodological contribution to the experimental matching literature by proposing an *ordinal efficiency test (OET)*.¹⁶ The idea is simply based on using the Pareto criterion to compare each observed NH4 outcome with each observed TTC outcome. Observe that for any given two outcomes, only one of the following can hold: (a) the NH4 outcome Pareto dominates the TTC outcome; (b) the TTC outcome Pareto dominates the NH4 outcome; and (c) the two cannot be Pareto ranked. To make the ordinal efficiency comparisons, we first considered 10,000 randomly generated priority orderings and then computed the number of cases where one mechanism's outcome Pareto dominates that of the other. Finally, we ran statistical tests to determine if there are any significant differences between the frequencies of NH4 and TTC dominations. Based on this test, we found NH4 to be significantly more likely to Pareto dominate TTC than the other way around. This result can be explained by the significantly higher participation rates under NH4 than under TTC. On the other hand, we did not discover any significant differences between truth-telling rates under NH4 and those under TTC.

Finally, we conducted a second experiment to compare the two theoretically equivalent mechanisms NH4 and GS in a laboratory environment. Behavior was far from optimal as many existing tenants decided not to participate and a substantial proportion of individuals did not reveal their true preferences. However, NH4 and GS induced very similar behavior in the lab. Indeed, we did not find any significant differences between the two mechanisms in terms of participation, truthful

¹⁵See, for example, Chen and Sönmez (2002, 2006), Pais and Pintér (2008), and Calsamiglia, Haeringer, and Klijn (forthcoming).

¹⁶We thank Jan Eeckhout (Editorial Board) and an anonymous referee for encouraging us to think in this direction.

preference revelation, or efficiency.

2 Related Literature

Since the seminal work of Gale and Shapley (1962), matching markets have been the focus of a growing theoretical and experimental mechanism design literature. Some popular applications in the literature include the design of the National Resident Matching Program (cf. Roth, 1984; and Roth and Peranson, 1999), the design of central student assignment mechanisms for U.S. public schools (cf. Abdulkadiroglu and Sönmez, 2003), the design of central kidney exchange clearing houses for kidney patients (cf. Roth, Sönmez, and Ünver 2004, 2007), and the design of course-bidding mechanisms at business schools (cf. Sönmez and Ünver, forthcoming).

Our paper lies at the intersection of the theoretical and the experimental literature on matching markets. On the theory side, our contribution is the reported equivalence of a real-life assignment mechanism with a celebrated mechanism from the theory of two-sided matching. To the best of our knowledge, there have only been two similar reports on the coincidence of a real-life mechanism with a Gale-Shapley deferred acceptance mechanism so far in the literature. The first is due to Roth (1984) who shows that the mechanism used by the National Resident Matching Program in the United States since 1951 until the 1990s to assign medical interns to hospital positions is actually an exact equivalent of the hospital-proposing Gale-Shapley mechanism. More recently, Balinski and Sönmez (1999) show that the multi-category serial dictatorship, a mechanism used for student placement to colleges in Turkey, is actually equivalent to the college-proposing Gale-Shapley mechanism. Interestingly, as opposed to these two equivalences, ours is the first report on an equivalent *agent-proposing* Gale-Shapley mechanism. Most notably, this equivalence enables us to single out NH4 as the most efficient individually rational and fair mechanism; and as essentially the only individually rational, fair, and incentive compatible mechanism.

Recent theory has explored a wide range of applications for top trading cycles mechanisms as a result of their attractive efficiency and incentives features. Several variations of top trading cycles mechanisms have been studied extensively in a number matching problems. Two well-known

applications for which this kind of mechanism has proved promising are *school choice* and *kidney exchange*, discussed earlier.

Despite their theoretical virtues, little is known about the performances of trade-based mechanisms when agents are boundedly rational. Our contribution on the experimental side lies in this direction. There is a tradition of using laboratory experiments for testing matching problems related to different real-world applications (cf. Olson and Porter 1994, Nalbantian and Schotter 1995, Harrison and McCabe 1996, Kagel and Roth 2000, and Ünver 2001). Two papers closely related to ours are those by Chen and Sönmez (2002, 2006). The former is the only experimental study of the present problem thus far, and we have already discussed it. The latter paper is an experimental study on three school choice mechanisms, which also includes a more subtle and complex counterpart of the top trading cycles mechanism.¹⁷ Our experimental findings are also consistent with those of Chen and Sönmez (2006), who observe the Gale-Shapley mechanism to generate more efficient outcomes than the complex counterpart of the top trading cycles mechanism introduced for the school choice context.¹⁸

Finally, two more recent and important contributions on experimental evaluation of school choice mechanisms are due to Pais and Pintér (2008) and Calsamiglia, Haeringer, and Klijn (forthcoming). In the former, the authors highlight the significance of the informational environment the subjects face. In contrast with Chen and Sönmez (2006), Pais and Pintér (2008) find the complex counterpart of the top trading cycles mechanism to be empirically more efficient than the Gale-Shapley mechanism. Calsamiglia, Haeringer, and Klijn (forthcoming) test the effect of limiting the number of schools that can be ranked and find that the incentives of the subjects are drastically affected by such a constraint.

¹⁷See Kesten (2006) for a theoretical comparison of the two mechanisms. Also see Erdil and Ergin (2008) and Kesten (2010) for two new mechanisms that Pareto-improve upon the Gale-Shapley mechanism in the context of school choice.

¹⁸In addition to the conceptual difference of house allocation from school choice due to the presence of existing tenants, an important difference of the model considered in this paper from a school choice model is that it allows for a simpler and more intuitive equivalent version of the top trading cycles mechanism. See Section 4.2.

3 The Model

Prior to the centralized assignment process, each existing tenant chooses whether to participate or not. Then a *house allocation problem (with existing tenants)* (Abdulkadiroğlu and Sönmez, 1999), or a *problem* for short, is given by¹⁹

1. a finite set of existing tenants I_{E+} who have chosen to participate,
2. a finite set of newcomers I_N ,
3. a finite set of occupied houses $H_O = \{h_i\}_{i \in I_{E+}}$,
5. a finite set of vacant houses H_V ,
6. a priority ordering f over all agents, and
7. a list of preferences $P = (P_i)_{I_{E+} \cup I_N}$ of all agents where P_i is the preference relation of agent i over all houses including the no-house option.

Let $I = I_{E+} \cup I_N$ denote the set of all agents. Each existing tenant $i \in I_{E+}$ is endowed with (i.e., currently lives in) a corresponding occupied house $h_i \in H_O$. Let h_0 denote the no-house option, or the so-called null house. Let $H = H_O \cup H_V \cup \{h_0\}$ denote the set of all houses. The exogenous priority ordering f may be determined based on certain criteria of the clearing house and is assumed to be strict. Formally, it is a one-to-one function $f : \{1, 2, \dots, |I_{E+} \cup I_N|\} \rightarrow I_{E+} \cup I_N$. Thus, agent $f(1)$ has the highest priority; agent $f(2)$ has the second highest priority; and so on. Let \mathcal{F} denote the set of all priority orderings.

For each agent $i \in I$, preference relation P_i is assumed to be strict. Let R_i denote the weak preference relation associated with P_i . Formally, we assume that R_i is a linear order, i.e., a complete, transitive, and anti-symmetric binary relation on H . That is, for any $h, h' \in H$, $h R_i h'$ if and only if $h = h'$ or $h P_i h'$. For expositional simplicity, we assume throughout the paper that the null house is the last option for each agent. Our analysis straightforwardly extends to the more general case. Let \mathcal{P}_i denote the set of all preferences for agent i , and let $\mathcal{B} = \times_{i \in I} \mathcal{P}_i$.

¹⁹Since our focus will be on individually rational mechanisms, we use a simpler version of the model proposed in Abdulkadiroğlu and Sönmez (1999) to facilitate exposition. In particular, our analysis neglects the effect of the choice of the assignment mechanism on the formation of the set of non-participating existing tenants.

Throughout, we shall suppress the first five components of a problem assuming that they are exogenously given and fixed. Hence, a problem is a pair (f, P) consisting of a priority ordering and a list of preferences. An *allocation* μ is a list of assignments such that (1) every agent is assigned one house; and (2) no house other than the null house is assigned to more than one agent. Formally, it is a function $\mu : I \rightarrow H$ such that for each $h \in H_O \cup H_V$, $|\mu^{-1}(h)| \in \{0, 1\}$. Let $\mu(i)$ denote the assignment of agent i at μ . Let \mathcal{A} denote the set of all allocations.

A *mechanism* φ is a systematic procedure that chooses an allocation for each problem. Formally, it is a function $\varphi : \mathcal{F} \times \mathcal{B} \rightarrow \mathcal{A}$. Let $\varphi(f, P)$ denote the allocation chosen by φ for the problem (f, P) and let $\varphi_i(f, P)$ denote the assignment of agent i at this allocation.

An allocation $\alpha \in \mathcal{A}$ is **individually rational** if no participating existing tenant prefers the house she has been occupying to the house she is assigned at α . Formally, for every $i \in I_{E+}$, $\alpha(i) R_i h_i$. A mechanism φ is individually rational if for every $(f, P) \in \mathcal{F} \times \mathcal{B}$, $\varphi(f, P)$ is individually rational.

An allocation $\alpha \in \mathcal{A}$ is **Pareto efficient** if its outcome cannot be improved by an allocation at which all agents are at least as well off and at least one agent is strictly better off. Formally, there is no $\mu \in \mathcal{A}$ such that $\mu(i) R_i \alpha(i)$ for all $i \in I$ and $\mu(j) P_j \alpha(j)$ for some $j \in I$. A mechanism φ is Pareto efficient if for every $(f, P) \in \mathcal{F} \times \mathcal{B}$, $\varphi(f, P)$ is Pareto efficient.

In the present context, “fully” respecting the priority order f may conflict with individual rationality.²⁰ Therefore, we propose and adopt the following intuitive “fairness” concept. An allocation $\alpha \in \mathcal{A}$ is **fair** if whenever an agent prefers another agent’s assignment to her own assignment at α , then either (1) the other agent has higher priority than herself; or (2) the other

²⁰For example, consider a simple problem with one newcomer and one existing tenant where the occupied house is the favorite house of both agents. Then, even though the newcomer may have higher priority than the existing tenant, any individually rational allocation must “favor” the existing tenant.

agent is an existing tenant who is assigned her own house. Formally, for every $i, j \in I$, if $\alpha(j) = h_j$ and $\alpha(i) = h_i$, then either $f(j) < f(i)$, or $j \in I_{E+}$ with $\alpha(j) = h_j$. A mechanism φ is fair if for every $(f, P) \in \mathcal{F} \times \mathcal{B}$, $\varphi(f, P)$ is fair.

A mechanism φ is **(dominant strategy) incentive compatible** (or, strategy-proof) if it is a dominant strategy for each agent to truthfully report her preferences. Formally, for every $(f, P) \in \mathcal{F} \times \mathcal{B}$, every $i \in I$, and every $P'_i \in \mathcal{P}_i$, $\varphi_i(f, P) R_i \varphi_i(f, P'_i, P_{-i})$.²¹

4 Three Mechanisms

We next study three mechanisms of interest: the most common on-campus housing mechanism in the United States, an attractive theory mechanism, and a particular mechanism currently in use at MIT.

4.1 Random Serial Dictatorship with Squatting Rights

Several U.S. universities including Carnegie Mellon, Duke, Northwestern, the University of Michigan, and the University of Pennsylvania employ the *random serial dictatorship with squatting rights (RSD)* mechanism. The mechanism works as follows:

Initially, each existing tenant decides whether she will enter the central assignment process, or keep her current house and stay out of the assignment process. In the former case, she gives up her occupied house, and it becomes available for allocation.

Consider a given house allocation problem (P, f) . Agents are successively asked to choose their top choice houses from the remaining ones: The first agent is assigned her top choice house; the second agent is assigned her top choice house among the remaining ones; and so on.

The main drawback of RSD is that since it cannot guarantee an existing tenant a house no worse than her current house, it is not individually rational. In fact, even though RSD is Pareto efficient according to our model, since some existing tenants may choose not to participate, it may

²¹Here, P_{-i} denotes the restriction of profile P to the set $I \setminus \{i\}$.

not be Pareto efficient within the set of all agents. Chen and Sönmez (2002) suggest that this aspect of the mechanism is one of its main deficiencies in practice.

4.2 Top Trading Cycles

Abdulkadiroğlu and Sönmez (1999) proposed the *top trading cycles (TTC)* mechanism, which is based on Gale’s top trading cycles idea (Shapley and Scarf, 1974). TTC works as follows:²²

Initially, each existing tenant decides whether she will enter the central assignment process, or keep her current house and stay out of the assignment process.

Consider a given house allocation problem (P, f) . Assign the first agent (according to f) her top choice, the second agent her top choice among the remaining houses, and so on, until someone demands the house of an existing tenant. If at that point the existing tenant whose house is demanded is already assigned a house, then do not disturb the procedure. Otherwise, modify the remainder of the ordering by inserting her at the top and proceed. Similarly, insert any existing tenant who is not already served at the top of the line once her house is demanded. If at any point a loop forms, it is formed by exclusively existing tenants each of whom demands the house of the tenant next in the loop. (A loop is an ordered list of agents (i_1, i_2, \dots, i_k) where agent i_1 demands the house of agent i_2 , agent i_2 demands the house of agent i_3, \dots , agent i_k demands the house of agent i_1 .) In such cases, remove all agents in the loop by assigning them the houses they demand and proceed.

TTC is Pareto efficient, individually rational, and incentive compatible (Abdulkadiroğlu and Sönmez, 1999). It achieves fairness only in a weak sense.²³ In their experimental study, Chen and Sönmez (2002) found that TTC induces higher participation rates and leads to more efficient

²²The description of TTC we give below differs from that of a direct adaptation of Gale’s original idea to house allocation with existing tenants. Abdulkadiroğlu and Sönmez (1999) have shown that the two alternative descriptions are nonetheless equivalent. Chen and Sönmez (2002) also used the version of TTC given below in their experiments.

²³Precisely speaking, in the following sense. Under TTC, it is possible that an agent, say i , may be assigned a house worse for her than a house an agent with lower priority, say j , is assigned. In fact, when this is the case, agent j is either an existing tenant who has been assigned her own house, or an existing tenant whose house has been assigned either to another existing tenant or to a newcomer with higher priority than j .

outcomes than RSD.

4.3 New House 4

The next mechanism has been in use at residence New House 4 of MIT for about three decades.²⁴

The New House 4 (NH4) mechanism works as follows:²⁵

Initially, each existing tenant decides whether she will enter the central assignment process, or keep her current house and stay out of the assignment process. Consider a given house allocation problem (P, f) .

1. The first agent (according to f) is tentatively assigned her top choice among all houses, the next agent is tentatively assigned her top choice among the remaining houses, and so on, until a squatting conflict occurs.

2. A squatting conflict occurs if it is the turn of an existing tenant but every remaining house is worse than her current house. That means someone else, the conflicting agent, is tentatively assigned the existing tenant's current house. When this happens:

(a) the existing tenant is assigned her current house and removed from the process, and

(b) all tentative assignments starting with the conflicting agent and up to the existing tenant are erased.

At this point the squatting conflict is resolved and the process starts over with the conflicting agent. Every squatting conflict that occurs afterward is resolved in a similar way.

The process is over when there are no houses or agents left. At this point all tentative assignments are finalized.

²⁴We are not aware of any other universities using a similar mechanism. This mechanism is first reported by Abdulkadiroğlu and Sönmez (1999).

²⁵See http://scripts.mit.edu/~nh4/wiki/index.php?title=Housing_Rules for an online description of this mechanism. The housing system at residence New House 4 was crafted through a joint effort between the president and the housing chair in the late 70's. (Personal communication with Sean Collins, the current president of New House 4.)

5 A Popular Mechanism from a Related Problem

The *Gale-Shapley (deferred acceptance)* mechanism (Gale and Shapley, 1962) has long dominated two-sided matching theory²⁶ because of its attractive stability and incentive properties. It has also been adopted by a number of real-life matching markets as a much more satisfactory alternative to the deficient mechanisms it replaced. The most recent success of the Gale-Shapley mechanism has been in “school choice” applications in the United States. In a *school choice problem* there are a finite set of students, a finite set of schools, a list of student preferences over the schools including a no-school option, and a list of school-specific priority orderings²⁷ over all students.²⁸ The main objective in this application is to assign students to schools while accommodating the school-specific priority orderings. Shortly after being advocated for school choice, the Gale-Shapley mechanism has replaced two controversial school choice mechanisms in New York City (Abdulkadiroğlu, Pathak, and Roth, 2005) and Boston (Abdulkadiroğlu, Pathak, Roth, and Sönmez, 2005).

Given the growing popularity and the success of the Gale-Shapley mechanism in matching markets, it is tempting to consider an adaptation of this mechanism for house allocation problems. Nevertheless, despite the mathematical similarities between the two indivisible good allocation problems, there are two important differences between house allocation and school choice that prevent us from immediately adapting the Gale-Shapley mechanism to house allocation. First, in school choice, but not in house allocation, for each school there is a separate (often different) priority ordering of students; and second, in school choice but not in house allocation, the individual rationality property is irrelevant since there is no conceptual analogue of existing tenants and their occupied houses. Therefore, we first provide a formal link between the two problems by transforming a given house allocation problem into an “associated school choice problem,” and then propose a direct adaptation of the Gale-Shapley mechanism to our setting.

²⁶See Roth and Sotomayor (1990) for a comprehensive survey on two-sided matching.

²⁷In the U.S., these priority orderings are typically based on a combination of a random lottery draw and specific policies of the particular school district.

²⁸A standard school choice problem also specifies a capacity vector that denotes the capacity of each school. To facilitate the exposition here we consider a “special” school choice problem where each school has unit capacity.

Given a house allocation problem $(I_{E+}, I_N, H_O, H_V, f, P)$, in the *associated school choice problem* $(I = I_{E+} \cup I_N, H = H_O \cup H_V \cup \{h_0\}, F = (F_h)_{h \in H}, P)$, the set of students is the set of all agents, the set of schools is the set of all houses, and the priority ordering for each house is constructed from f as follows: (1) for each $h \in H_V \cup \{h_0\}$, let $F_h \equiv f$; and (2) for each $h_i \in H_O$, let $F_{h_i}(1) \equiv i$ and for each $j, k \in I \setminus \{i\}$, $F_{h_i}^{-1}(j) < F_{h_i}^{-1}(k) \iff f^{-1}(j) < f^{-1}(k)$. In words, under this transformation, in the associated school choice problem, (1) the priority ordering for each vacant house coincides with the priority ordering f of the given house allocation problem; and (2) the priority ordering for each occupied house assigns the highest priority for this house to the corresponding existing tenant and exactly follows f in assigning the remaining priorities.

As with the house allocation problems, in the remainder of the paper, we shall suppress the set of agents I and the set of houses H and define an associated school choice problem as a pair (F, P) consisting of a list of house-specific priority orderings and a list of preferences. The definitions of an *allocation* and a *mechanism* identically apply in associated school choice problems. In school choice, a central consideration is “stability.” Given a school choice problem (F, P) , an allocation μ is **stable** if there is no unmatched agent-house pair (i, h) such that either (i) agent i prefers h to her assignment and h is unassigned at μ , or (ii) agent i prefers h to her assignment and has higher priority for h than the agent assigned to h .²⁹ Formally, there exists no $(i, h) \in I \times H$ such that either (i) $h P_i \mu(i)$ and $\mu^{-1}(h) = \emptyset$, or (ii) $h P_i \mu(i)$ and $F_h^{-1}(i) < F_h^{-1}(\mu^{-1}(h))$.

Given a house allocation problem (f, P) , the outcome of the *Gale-Shapley (GS) mechanism* is obtained by applying the following *deferred acceptance (DA) algorithm* (Gale and Shapley, 1962) to the associated school choice problem (F, P) :

Step 1: *Each agent applies to her top choice house. If an agent applies to the null house, then she is permanently assigned to it. For each remaining house, consider its applicants. The agent with the highest priority according to the priority ordering for that house is tentatively placed at that house. The rest are rejected.*

²⁹The traditional stability notion also requires each agent to be assigned a house that is at least as good as the null house. (This is also referred as *individual rationality*.) Since the null house is the last option for each agent (and since each agent is acceptable to each house) in our model, this requirement is trivially satisfied.

In general:

Step k: *Each rejected agent applies to her next top choice house. If an agent applies to the null house, then she is permanently assigned to it. For each remaining house, consider its applicants at this step together with the agent (if any) who is currently tentatively placed to it. Among these, the agent with the highest priority according to the priority ordering for that house is tentatively placed at that house. The rest are rejected.*

The algorithm terminates when no agent is rejected any more. For any given house allocation problem, GS always leads to a stable allocation for the associated school choice problem (Gale and Shapley, 1962). Moreover, this is the most favorable stable allocation to agents (i.e., it Pareto dominates any other stable allocation). Another remarkable feature of GS is that it is also incentive compatible (Dubins and Freedman, 1981, and Roth, 1982).

6 Main Results

We start with a negative result on the trade-offs among the properties of house allocation mechanisms. It turns out that any three of the four desirable properties we discussed in Section 2 are compatible except for individual rationality, Pareto efficiency, and fairness.

Proposition 1: *No mechanism is individually rational, Pareto efficient, and fair. However, there exist mechanisms satisfying any other three properties simultaneously.*

The first statement in Proposition 1 is in the same spirit as the classical impossibility between stability and Pareto efficiency in two-sided matching due to Roth (1982). For the second statement, consider the three mechanisms we previously discussed. RSD is Pareto efficient, fair, and incentive compatible, but not individually rational. TTC is individually rational, Pareto efficient, and incentive compatible, but not fair.³⁰ As will be shown shortly in Corollary 1, NH4 is individually

³⁰Abdulkadiroglu and Sönmez (1999) showed that TTC achieves the first three properties. Then Proposition 1 implies that it is not fair.

rational, fair, and incentive compatible, but not Pareto efficient.³¹

A much weaker notion than Pareto efficiency is “non-wastefulness.” An allocation is *non-wasteful*³² if no agent prefers an unassigned house to her assignment at this allocation. The next result presents a formal connection between the three properties of individually rationality, fairness, and non-wastefulness for a house allocation problem and the stability property for the associated school choice problem.

Proposition 2: *Given a house allocation problem (f, P) , an allocation is individually rational, fair, and non-wasteful if and only if it is stable for the associated school choice problem (F, P) .*

Much to our surprise, the above natural adaptation of the Gale-Shapley deferred acceptance procedure in fact always yields exactly the same outcome as NH4.

Theorem 1: *NH4 and GS are equivalent.*

Theorem 1, to the best of our knowledge, is the third reported coincidence of the Gale-Shapley deferred acceptance procedure with a real-life mechanism. Interestingly, contrary to the previous two reports, NH4 is an *agent-proposing* deferred acceptance procedure.

By Proposition 2, the equivalence result in Theorem 1 enables NH4 to claim all the attractive properties of GS.

Corollary 1: *NH4 (as well as GS) is individually rational, fair, and incentive compatible.*

The next result states that among all fair and individually rational mechanisms, NH4 is simply the most efficient one. By Theorem 2 of Balinski and Sönmez (1999), the following corollary is now immediate.

³¹See Abdulkadiroglu and Sönmez (1999) for an example that shows that the NH4 outcome may not be Pareto efficient. The actual housing rules at MIT also allow two agents to swap rooms if they have the consent of every agent with priorities between them.

³²Formally, an allocation $\alpha \in \mathcal{A}$ is **non-wasteful** if for every $i \in I$ and every $h \in H$, $h P_i \alpha(i)$ implies that there is $j \in I$ such that $\alpha_j(P) = h$. A mechanism φ is non-wasteful if for every $(f, P) \in \mathcal{F} \times \mathcal{B}$, $\varphi(f, P)$ is non-wasteful.

Corollary 2: *NH4 (as well as GS) Pareto dominates any other individually rational and fair mechanism.*

By Theorem 4 of Balinski and Sönmez (1999), NH4 turns out to be essentially the only incentive compatible, individually rational, and fair mechanism.

Corollary 3: *NH4 (as well as GS) is the unique individually rational, fair, incentive compatible, and non-wasteful mechanism.*

The leading theory mechanism for house allocation TTC and mechanism NH4 of MIT both satisfy three of the four properties in our desiderata. Theory suggests that TTC has the edge in terms of efficiency and NH4 in terms of fairness. Our next goal will be to experimentally test the two mechanisms. This is the subject of the next section.

7 Experiment 1: NH4 vs. TTC

7.1 Experimental design

Our experiment compares the performances of NH4 and TTC in terms of efficiency, participation of existing tenants, and truthful preference revelation. We implemented two treatments which differ only in the house allocation mechanism. For the sake of comparison, we tried to keep our design as close as possible to that of Chen and Sönmez (2002); participants were given the same description of TTC that was provided by Chen and Sönmez (2002). As for NH4, we used the description of this mechanism that was provided by MIT.³³

We ran five replications for each treatment (NH4-1, NH4-2,..., NH4-5; TTC-1, TTC-2,..., TTC-5). Each replica was run in a separate session at the CLER experimental lab, Harvard Business School during spring and early summer 2006. We used Urs Fischbacher's z-Tree package

³³See the experimental instructions provided in the Supplementary Materials, available online, for a precise description of the mechanisms.

Fischbacher (2007). Each group consists of 12 participants. Participants 1 through 8 are existing tenants. Participants 9 through 12 are newcomers. There are also 12 houses of 8 different types to be allocated. House types go from A to H. Participants 1 through 12 are existing tenants, each living in a house type A through H. There are four additional vacant houses of types A, B, C, and D. Table 1 shows the payoff for each participant as a result of the house type she gets at the end of the experiment. A square bracket, [], shows that the participant is an existing tenant of a house of the specified type. For instance, participant 2 lives in a type B house. She gets \$10 at the end of the experiment if she ends up in the same house. Note that our payments are a scaled-up version of the Chen and Sönmez (2002) setup, as we added \$5 on top of each payment in their design. This was done in order to meet the payment criteria of the CLER laboratory. Our payoff parameters have the following implications:

1. There are nine Pareto efficient house allocations. The aggregate payoff adds up to 231 for each Pareto efficient allocation.
2. Existing tenants' houses range from their second to the seventh choice. Otherwise the decision to participate becomes trivial.
3. There is a salient monetary difference of \$14 between the top and the last choice.

Both treatments, NH4 and TTC, were implemented as one-shot games of incomplete information. Each participant knew her own payoff table but not the others' payoff tables. Participants did know the number of existing tenants and newcomers and that payoff tables may differ. In both treatments, existing tenants were given an option to keep their houses and then not participate in the assignment mechanism. In line with Chen and Sönmez (2002), participants were explained in the instructions the workings of the mechanisms, and in particular, how an existing tenant always obtains a house that is at least as good as her current house if she decides to participate. Though the descriptions of the mechanisms make the individual rationality property easy to see, this is not true for the incentive compatibility property. In neither treatment were the subjects given any information about the incentive compatibility property of the mechanisms. If the instructions stated that the mechanisms are incentive compatible without explaining why, we would

have generated an undesirable demand effect.³⁴ That is, we would have revealed what constitutes appropriate behavior in our experiment and that could become the main force driving the results.

The experiment was conducted as follows. Once each participant was assigned to a computer, the experimenter read the instructions aloud, and questions were answered. Then, participants saw their own payoff table in the computer screen. Participants had 10 minutes to go over the instructions and make decisions. Existing tenants had the option to keep their current house (by choosing “out”) or to participate in the mechanism (by choosing “in”). Existing tenants who chose “in” and newcomers submitted their list of preferences. Their ID numbers were introduced in a bowl by the experimenter, and a randomly chosen participant drew them one by one in order to generate the initial priority ordering. At this point the assignment of the houses was computed manually. At the end of the experiment participants were informed about the resulting assignment and were paid accordingly.

TABLE 1. PAYOFF TABLE FOR ALL AGENTS (NH4 vs TTC)

Types of Houses		A	B	C	D	E	F	G	H
Existing Tenants	#1	[11]	8	13	14	20	10	6	17
	#2	11	[10]	14	13	8	17	20	6
	#3	6	8	[14]	20	10	11	17	13
	#4	10	14	20	[17]	8	11	13	6
	#5	10	6	17	14	[8]	20	13	11
	#6	20	11	14	13	6	[17]	8	10
	#7	8	10	11	17	6	13	[14]	20
	#8	14	20	10	17	11	8	6	[13]
Newcomers	#9	6	10	17	14	11	20	13	8
	#10	11	6	17	14	10	20	8	13
	#11	20	10	14	6	17	11	13	8
	#12	13	20	8	10	11	14	17	6

Our experimental design allows us to test the following three hypotheses based on the theo-

³⁴See Zizzo (2010) for a state-of-the-art review of the experimenter demand effect.

retical properties of the NH4 and TTC mechanisms:

Hypothesis 1: *TTC can Pareto dominate NH4, the opposite is not true.*³⁵

Hypothesis 2: *Existing tenants choose to participate under both TTC and NH4.*

Hypothesis 3: *Participants choose to reveal their preferences truthfully under both TTC and NH4.*

7.2 Results

To evaluate the aggregate performance of NH4 vs. TTC, we compare the outcomes generated by each mechanism. Following the efficiency concept used in the theoretical literature we use a novel way to test for efficiency based on Pareto comparisons, the *ordinal efficiency test (OET)*. As far as we are aware, experimental studies on matching have so far used the cardinal concept of utility when making efficiency comparisons. This is the case in Chen and Sönmez (2002, 2006), Pais and Pintér (2008) and Calsamiglia, Haeringer, and Klijn (forthcoming). Pais and Pintér (2008) test for the difference in proportion of average efficiency across mechanisms. The other three papers follow Chen and Sönmez (2002) in using the recombinant estimation technique proposed by Mullin and Reily (2006). It is worth noting that Chen and Sönmez (2002) also include simulations in which the environment changes randomly and both participation rates and preference revelation follow experimental findings in order to obtain more robust results. The OET proposed in this paper works as follows:

1. We construct the possible 25 pairs combining the 5 NH4 replicas and the 5 TTC replicas:
NH4-1 vs. TTC-1, NH4-1 vs. TTC-2,..., NH4-2 vs. TTC-1,..., NH4-5 vs. TTC-5.
2. For each pair, we randomly generate 10,000 different priority orderings.

³⁵Because TTC is Pareto efficient and NH4 is not, according to the theory, given a fixed set of participating agents, there exist problems for which the TTC outcome Pareto dominates the NH4 outcome but, for no problem can the NH4 outcome Pareto dominate the TTC outcome. This forms the basis for our efficiency hypothesis.

3. We perform a Pareto comparison for each of the 10,000 priority orderings for all 25 pairs.
4. For each pair, we count the number of times NH4 dominates TTC (NH4 dominations) and the number of times TTC dominates NH4 (TTC dominations).
5. Finally, we use a sign rank Wilcoxon test for the equality of the 25 matched pairs of dominations.

Result 1 (Efficiency): *NH4 is more likely to Pareto dominate TTC than the other way around.*

The OET rejects the hypothesis of equality in the number of NH4 and TTC dominations: $z = 2.600$ ($p = 0.0093$). Overall, NH4 dominates TTC 18,203 times, and TTC dominates NH4 321 times out of 250,000 Pareto comparisons. Therefore, we do not find support for Hypothesis 1. (See the Supplementary Materials, available online, for the frequency of dominations for each of the 25 pairs considered.)

Result 2 (Participation): *Existing tenants under NH4 are significantly more likely to participate than those under TTC. The existing tenants' overall participation rate is 77.5% under NH4, but only 47.5% under TTC.*

A test of equality of proportions shows that the participation rate of existing tenants under NH4 is significantly higher than that under TTC: $z = 2.7713$ ($p = 0.0028$). Hence, we reject Hypothesis 2.

TABLE 3. PARTICIPATION AND TRUTHFUL PREFERENCE REVELATION (NH4 vs TTC)

Mechanisms	Group	Participation rate	Proportion of truthful preference revelation
NH4	NH4-1	8/8	10/12
	NH4-2	5/8	8/9
	NH4-3	6/8	7/10
	NH4-4	6/8	9/10
	NH4-5	6/8	7/10
TTC	TTC-1	4/8	6/8
	TTC-2	5/8	6/9
	TTC-3	3/8	5/7
	TTC-4	4/8	6/8
	TTC-5	3/8	4/7

Table 3 shows the participation rates in column 3 and the proportions of truthful preference revelation for each group in column 4.

Result 3 (Truthful Preference Revelation): *The overall proportion of truthful preference revelation is 80.4% under NH4, and 69.0% under TTC. The differences in proportions of truthful preference revelation under NH4 and TTC are not statistically significant.*

A test of equality of proportions shows that the proportion of truthful preference revelation under NH4 is not significantly different from that under TTC: $z = 1.2250$ ($p = 0.1103$). However, neither NH4 nor TTC induced truthful preference revelation for all participants. Therefore, we reject Hypothesis 3.

Results 1 to 3 show behavior deviating from theory in both the NH4 and the TTC treatments. NH4 is found to achieve greater efficiency than TTC. Since we do not find significant differences in truthful preference revelation (Result 3), we can conclude that participation is the key to why NH4 is outperforming TTC even though theory does not support this finding.

8 Experiment 2: NH4 vs. GS

8.1 Experimental design

Our second experiment tests the theoretical equivalence of NH4 and GS. The design of this experiment follows the design of NH4 vs. TTC, but there are a few differences. We again implemented two treatments, NH4 and GS, which differ only in the house allocation mechanism, or rather in the descriptions of the mechanisms given in the instructions.

We ran eight replications for each treatment (NH4v-1, NH4v-2,..., NH4v-5; Gsv-1, Gsv-2,..., Gsv-5) in two sessions for NH4 and three sessions for GS. Each session was run at the LINEEX experimental laboratory, University of Valencia, during spring 2010. Table 4 shows the payoff for each participant as a result of the house type she gets at the end of the experiment. Table 4 contains payoffs in Euros obtained by using the Euro/\$ exchange rate at the time and rounding to the closest integer. On top of the payments shown in Table 4, participants were paid a 3EUR show-up fee and 0.50EUR for each right answer in an 8-question post-experiment quiz.³⁶ Other than the different show-up fee and the reference to the post-experiment quiz, the NH4 instructions used in Valencia are a Spanish translation of the original instructions used for the Harvard NH4 vs TTC experiments. A new set of instructions for GS was written in English and also translated into Spanish. (See the Supplementary Materials, available online.)

³⁶The questionnaire is related to a different project. The questions are available upon request.

Table 4. Payoff Table for All Agents (NH4 vs GS)

		Types of Houses							
		A	B	C	D	E	F	G	H
Existing Tenants	#1	[9]	6	10	11	16	8	5	14
	#2	9	[8]	11	10	6	14	16	5
	#3	5	6	[11]	16	8	9	14	10
	#4	8	11	16	[14]	6	9	10	5
	#5	8	5	14	11	[6]	16	10	9
	#6	16	9	11	10	5	[14]	6	8
	#7	6	8	9	14	5	10	[11]	16
	#8	11	16	8	14	9	6	5	[10]
Newcomers	#9	5	8	14	11	9	16	10	6
	#10	9	5	14	11	8	16	6	10
	#11	16	8	11	5	14	9	10	6
	#12	10	16	6	8	9	11	14	5

Our experimental design allows us to test the following three hypotheses based on the theoretical equivalence of the NH4 and GS mechanisms:

Hypothesis 4: *NH4 and GS cannot Pareto dominate each other.*

Hypothesis 5: *Existing tenants choose to participate under both NH4 and GS.*

Hypothesis 6: *Participants choose to reveal their preferences truthfully under both NH4 and GS.*

8.2 Results

Result 4 (Efficiency): *NH4 and GS are equally likely to Pareto dominate each other.*

The OET does not reject the hypothesis of equality in the number of NH4 and GS dominations: $z = 0.375$ ($p = 0.7080$). Overall, NH4 dominates GS 6,790 times, and GS dominates NH4 3,896 times out of 250,000 Pareto comparisons. Hypothesis 4 is therefore rejected. (See the

Supplementary Materials, available online, for the frequency of dominations for each of the 25 pairs considered.)

Result 5 (Participation): *Existing tenants are equally likely to participate under both mechanisms. The existing tenants' overall participation rate is 67.0% under NH4, and 70.3% under GS.*

A test of equality of proportions shows that the participation rate of existing tenants under NH4 does not differ from that under TTC: $z = -0.38$ ($p = 0.65$). Existing tenants do not always choose to participate. Thus, Hypothesis 5 is rejected.

Table 5. Participation and Truthful Preference Revelation (NH4 vs GS)

Mechanisms	Group	Participation rate	Proportion of truthful preference revelation
NH4	NH4v-1	4/8	3/8
	NH4v-2	3/8	2/7
	NH4v-3	7/8	9/11
	NH4v-4	7/8	5/11
	NH4v-5	7/8	7/11
	NH4v-6	4/8	6/7
	NH4v-7	6/8	8/10
	NH4v-8	5/8	8/9
GS	GSv-1	5/8	4/9
	GSv-2	7/8	7/11
	GSv-3	5/8	7/9
	GSv-4	6/8	6/10
	GSv-5	4/8	4/7
	GSv-6	6/8	7/10
	GSv-7	7/8	6/9
	GSv-8	5/8	6/9

Table 5 shows participation rates in column 3 and proportions of truthful preference revelation for each group in column 4.

Result 6 (Truthful Preference Revelation): *The differences in proportions of truthful preference revelation under NH4 and GS are not statistically significant. The overall proportion of truthful preference revelation is 63.9% under NH4 and 63.3% under GS.*

A test of equality of proportions shows that the proportion of truthful preference revelation under NH4 is not significantly different from that under GS: $z = 0.17$ ($p = 0.43$). Preferences are often not truthfully revealed. Therefore, Hypothesis 6 is rejected.

Results 4 to 6 show behavior deviating from theory in both treatments of the experiment. We did not find significant differences in either participation or truthful preference revelation, and thus found no significant difference in efficiency. Our experiments suggest that, in expected terms, NH4 and GS generate similar results. Consequently, we do not find any empirical evidence to prefer NH4 over GS or the other way around.

9 Conclusion and Discussion

In this paper we have studied the problem of finding the “right” house allocation mechanism to allocate on-campus housing units to students from a market-design perspective. Abdulkadiroğlu and Sönmez (1999) and Chen and Sönmez (2002) advocated the prominent theory mechanism TTC as a serious candidate to replace the popular real-life mechanism RSD. In theory, TTC is Pareto efficient, individually rational, and incentive compatible, but not fair. We analyzed the MIT house allocation mechanism known as NH4. In theory, NH4 is individually rational, incentive compatible and fair, but not Pareto efficient. By Theorem 1, NH4 is however, the only constrained-efficient mechanism among all mechanisms that are individually rational and fair.

We designed a laboratory experiment in which NH4 and TTC went head to head. Notwithstanding the theoretical advantage of TTC, NH4 turned out to be superior to TTC in terms of both participation rates and efficiency. We designed a second experiment to test the theoretical

equivalence of NH4 and GS in the lab. Not all the existing tenants chose to participate and preferences were not always revealed truthfully. Furthermore, we did not find any significant differences between the two mechanisms in terms of participation, truthful preference revelation, or efficiency.

A second reason to be optimistic about the efficiency performance of NH4 comes from a result due to Ergin (2002): Loosely speaking, Ergin (2002) showed that GS tends to be more efficient as the priority orderings for each house tend to be more “correlated.” One feature of the particular adaptation of GS we have considered that might contribute to this possibility is that all the house-specific priority orderings for GS (the equivalent of NH4) are, by construction, generated from the same priority ordering.

10 Appendix

Proof of Proposition 1: For the first statement consider the following example. Suppose $I_N = \{1, 2\}$, $I_{E+} = \{3\}$, $H_V = \{a, b\}$, and $H_O = \{h_3\}$. Suppose the priority ordering f is 1-2-3. Agents’ preferences are as follows:

R_1	R_2	R_3
h_3	a	a
a	h_3	h_3
b	b	b

Any Pareto efficient mechanism has to assign either agent 2 or agent 3 to house a , for otherwise agent 1 gets house a and is made better off when she swaps it with the agent who gets house h_3 (who is also made better off by this swap). Then, since agent 2 has higher priority, by fairness she should be assigned house a . This means, by individual rationality, agent 3 should be assigned house h_3 . Then agent 1 is assigned house b . But this clearly violates fairness. **Q.E.D.**

Proof of Proposition 2: Take an allocation $\alpha \in \mathcal{A}$. Clearly, condition (i) of stability of α for (F, P) is equivalent to the non-wastefulness of α for (f, P) . Suppose that α is individually rational, fair, and non-wasteful for (f, P) but not stable for (F, P) . Then there exists $(i, h) \in I \times H$

such that $h P_i \alpha(i)$ and $F_h^{-1}(i) < F_h^{-1}(\mu^{-1}(h))$. If $h = h_i$, then α is not individually rational. If $h \neq h_i$, then α is not fair.

Suppose that α is stable for (F, P) but not individually rational for (f, P) . Then there exists $j \in I$ with $h_j P_j \alpha(j)$. Since $F_{h_j}(j) = 1$ by construction, α cannot be stable for (F, P) . Suppose that α is stable for (F, P) but not fair for (f, P) . Then there exist $k, l \in I$ with $\alpha(k) P_l \alpha(l)$ such that $f^{-1}(l) < f^{-1}(k)$ and either $k \in I_N$, or $k \in I_{E+}$ and $\alpha(k) \neq h_k$. For both cases, by construction, $F_{\alpha(k)}^{-1}(l) < F_{\alpha(k)}^{-1}(k)$. Thus, α cannot be stable for (F, P) . **Q.E.D.**

Proof of Theorem 1: Fix a problem (f, P) . Let (F, P) be the associated school choice problem. We first introduce a useful notion of a “partial allocation” to denote the pseudo-allocations that form during the execution of the NH4 algorithm as well as the DA algorithm. Let $I' \subset I$. A *partial allocation* is a function $\mu : I' \rightarrow H$ such that for each $h \in H_O \cup H_V$, $|\mu^{-1}(h)| \in \{0, 1\}$. In this case, we say that *partial allocation* μ is defined on I' . Clearly, any allocation is also a partial allocation (defined on I). A partial allocation μ defined on I' is fair for a set of agents $I'' \subset I'$ if for all $i, j \in I''$, whenever $\mu(j) P_i \mu(i)$, then either $f(j) < f(i)$, or $j \in I_{E+}$ with $\mu(j) = h_j$.

McVitie and Wilson (1970) showed that under the DA algorithm, the ordering according to which agents make proposals to houses has no effect on the outcome, and proposed an equivalent version of the DA algorithm, where agents make their applications according to any given ordering. To prove Theorem 1 we use the McVitie-Wilson version of the DA algorithm in which agents apply to houses in turn according to the ordering f . More precisely, it will be convenient to think of the DA algorithm in the following way:

First, agent $f(1)$ applies to her favorite house, say, to house h^1 . Since there are no previous applicants, she is tentatively assigned to her favorite house at the partial allocation that forms. Second, agent $f(2)$ applies to her favorite house, say, to house h^2 . If $h^2 = h^1$ and $h^2 \neq h_{f(2)}$, then [since $F_{h^2}^{-1}(f(1)) < F_{h^2}^{-1}(f(2))$] agent $f(2)$ is rejected from h^2 . Otherwise, agent $f(2)$ is tentatively assigned to h^2 at the partial allocation that forms. If some agent is rejected from house h^2 in either case, then she applies to her next favorite house as in the usual DA algorithm... In general, agent $f(k)$ applies to her favorite house, say, to house h^k . If h^k received a previous application and $h^k \neq h_{f(k)}$, then agent $f(k)$ is rejected from h^k . Otherwise, agent $f(k)$ is tentatively assigned

to h^k at the partial matching that forms. If some agent is rejected from house h^k in either case, then she applies to her next favorite house. Similarly, any subsequently rejected agent applies to her next favorite house, and the process continues as in the usual DA algorithm until no agent is rejected any more; and so on.

The following observations will be useful in proving Theorem 1.

Lemma 1: *Any partial (and tentative) allocation obtained throughout the NH4 algorithm is fair for the agent set it is defined on. Thus, NH4 is fair. Similarly, any partial (and tentative) allocation obtained throughout the DA algorithm is fair for the agent set it is defined on. Thus, GS is fair.*³⁷

Proof. Take any partial allocation $\mu : I' \rightarrow H$ that forms at some instance of the NH4 algorithm. Let $i \in I'$ and $h = \mu(i)$. There are two cases to consider:

Case 1. i is not a squatting agent: Consider the instance of the algorithm when i is assigned h . When it is the turn of agent i , she is assigned her best choice among the remaining houses. This implies that any house that she prefers to h must be previously assigned to a higher priority agent than i , or to its owner because of a squatting conflict. Then we conclude that for all $j \in I'$, if $\mu(j) P_i \mu(i)$, then either $f(j) < f(i)$, or $j \in I_{E+}$ with $\mu(j) = h_j$.

Case 2. i is a squatting agent: Clearly, $h = h_i$. Consider the instance of the algorithm when i is assigned h . When it is the turn of agent i , all remaining houses are worse for her than h_i . Thus agent i squats and all the (tentative) assignments starting with the conflicting agent and up to agent i are erased. Next, the conflicting agent is to be assigned her top choice among the remaining houses. Note that the set of remaining houses at this point is a strict subset of the set of remaining houses when the conflicting agent was (tentatively) assigned h_i . Since the preferences of the agents who have higher priority than i but weakly lower priority than the conflicting agent are fixed, when it is the turn of agent $f(f^{-1}(i) + 1)$, the set of remaining houses is a strict subset of the set of remaining houses at the point agent i squatted. Consequently, agent $f(f^{-1}(i) + 1)$ cannot be assigned a house that i prefers to h_i . Combining this with Case 1 implies that for all

³⁷Clearly, Proposition 2 also implies that GS is fair.

$j \in I'$, if $\mu(j) P_i \mu(i)$, then either $f(j) < f(i)$, or $j \in I_{E+}$ with $\mu(j) = h_j$.

Take any partial allocation $\mu' : I' \rightarrow H$ that forms at some instance of the DA algorithm. Let $i' \in I'$ and $h' = \mu(i')$. Then until this point agent i' must have been rejected by all houses that she prefers to h' . Hence, any such house must be tentatively holding an agent who has higher priority than i' according to F . But the tentatively held agent has higher priority for that house according to F only if she has higher priority than i' according to f , or if she is the owner of that house. Thus, for all $j \in I'$, if $\mu'(j) P_i \mu(i')$, then either $f(j) < f(i')$, or $j \in I_{E+}$ with $\mu(j) = h_j$.

Q.E.D.

Lemma 2: *Given $I' \subset I$, let μ and μ' be two partial allocations defined on I' . Suppose the set of existing tenants who are assigned their own houses are the same at μ and μ' . Let J denote this set. If μ and μ' are both non-wasteful and fair for $I' \setminus J$, then $\mu = \mu'$.*

Proof. Let $H' = H \setminus \cup_{j \in J} \{h_j\}$. Consider the highest priority agent in $I' \setminus J$. Since both μ and μ' are fair and non-wasteful, they should both assign her her top choice in H' . Next, consider the second highest priority agent in $I' \setminus J$. Since both μ and μ' are fair and non-wasteful, they should also assign her her top choice among the remaining houses in H' . Continuing similarly, we conclude that for all $i \in I' \setminus J$, $\mu(i) = \mu'(i)$. Then $\mu = \mu'$. **Q.E.D.**

We give a direct proof of Theorem 1. We show that the set of existing tenants who are assigned their own houses are the same under the two mechanisms at the problem (f, P) . Then, since both algorithms' outcomes are fair by Lemma 1 and non-wasteful,³⁸ by Lemma 2 they have to choose the same allocation. Let α and γ respectively be the NH4 allocation at (f, P) and the GS allocation at (F, P) .

Let $S = \{i_1, i_2, \dots, i_T\} \subset I_{E+}$ be the set of agents who squatted in the NH4 algorithm. W.l.o.g., let the subscripts denote the order with which they squatted. We show that for each $i \in S$, $\gamma(i) = h_i$. We argue by induction.

Consider agent i_1 . We observe how the tentative and partial allocation is obtained until just before the first squatting conflict under the NH4 algorithm. First, agent $f(1)$ is tentatively assigned

³⁸By Proposition 1, GS is non-wasteful. The fact that NH4 is also non-wasteful can be straightforwardly verified.

her top choice house in H ; next, agent $f(2)$ is tentatively assigned her top choice among the remaining houses;...; and finally, agent $f(f^{-1}(i_1) - 1)$ is tentatively assigned her top choice among the remaining houses. Under the DA algorithm, first, agent $f(1)$ applies to her top choice house to which she is tentatively assigned;...; in general, agent $f(k)$ with $2 \leq k \leq f^{-1}(i_1) - 1$, eventually applies to her top choice house among those that no one has applied to before and is tentatively assigned to it. Let α^1 denote the partial allocation that forms under the NH4 algorithm just before agent i_1 squats. Note that α^1 is defined on $I'_1 = \{f(1), f(2), \dots, f(f^{-1}(i_1) - 1)\}$. Let γ^1 denote the partial allocation that forms under the DA algorithm just before agent i_1 starts to apply. Clearly, γ^1 is also defined on I'_1 . Moreover, upon contrasting the two procedures, it is straightforward to see that $\alpha^1 = \gamma^1$.³⁹ Under the NH4 algorithm, when it is the turn of (existing tenant) i_1 , by assumption, all available houses are worse for her than h_{i_1} . Under the DA algorithm, when it is the turn of agent i_1 to start to apply, since $\alpha^1 = \gamma^1$, she applies to all the houses she prefers to h_{i_1} (if any), and in turn gets rejected from each such house since it is now tentatively assigned to some agent who, by the construction of F , has higher priority for it than i_1 . Then agent i_1 applies to h_{i_1} . Since she has the highest priority for h_{i_1} , she is tentatively assigned to h_{i_1} . From this point on, any agent who applies to h_{i_1} is rejected, i.e., i_1 is permanently assigned to h_{i_1} . Hence, $\gamma(i_1) = h_{i_1}$.

Suppose that for each $i_2, \dots, i_{k-1} \in S$, we have $\gamma(i_2) = h_{i_2}, \dots, \gamma(i_{k-1}) = h_{i_{k-1}}$. [Induction hypothesis] We show that $\gamma(i_k) = h_{i_k}$. Let α^k be the partial allocation that forms under the NH4 algorithm just before agent i_k squats. Note that α^k is defined on $I'_k = \{f(1), f(2), \dots, f(f^{-1}(i_k) - 1)\}$. Let γ^k be the partial allocation that forms under the DA algorithm just before agent i_k starts to apply. Clearly, γ^k is also defined on I'_k . By Lemma 1, both α^k and γ^k are fair for I' . Since they are both also non-wasteful, by Lemma 2 and the induction hypothesis, $\alpha^k = \gamma^k$. Under the DA algorithm, when it is the turn of agent i_k to start to apply, since $\alpha^k = \gamma^k$, she applies to all the houses she prefers to h_{i_k} (if any), and in turn gets rejected from each such house since it is now tentatively assigned to some agent who has higher priority for it than i_k . Then agent i_k applies to

³⁹An alternative way to show this is to invoke Lemma 2 with $J = \emptyset$. Here we adopt the longer and direct proof to provide further insight into the equivalence result.

h_{i_k} . Since she has the highest priority for h_{i_k} , she is tentatively assigned to h_{i_k} . From this point on, any agent who applies to h_{i_k} is rejected, i.e., i_k is permanently assigned to h_{i_k} . Hence, $\gamma(i_k) = h_{i_k}$. This completes the proof. **Q.E.D.**

References

1. Abdulkadiroğlu, A. and T. Sönmez, House allocation with existing tenants, *Journal of Economic Theory* 88 (1999), 233-260.
2. Abdulkadiroğlu, A., P. Pathak, A. Roth, and T. Sönmez, The Boston public school match, *American Economic Review, Papers and Proceedings*, 95 (2005), 368-371.
3. Abdulkadiroğlu, A., P. Pathak and A. Roth, The New York City high school match, *American Economic Review, Papers and Proceedings*, 95 (2006), 364-367.
4. Abdulkadiroğlu, A. and T. Sönmez, School choice: A mechanism design approach, *American Economic Review* 93 (2003), 729-747.
5. Balinski, M. and T. Sönmez, A tale of two mechanisms: Student placement, *Journal of Economic Theory* 84 (1999), 73-94.
6. Calsamiglia, C, G. Haeringer, and P. Klijn, Constrained school choice: An experimental study, *American Economic Review*, forthcoming.
7. Chen, Y. and T. Sönmez, Improving efficiency of on-campus housing: An experimental study, *American Economic Review* 92 (2002), 1669-1686.
8. Chen, Y. and T. Sönmez, School choice: An experimental study, *Journal of Economic Theory* 127 (2006), 202-231.
9. Dubins, L.E. and D.A. Freedman, Machiavelli and the Gale-Shapley algorithm, *American Mathematical Monthly* 88 (1981), 485-494.

10. Erdil, A. and H. Ergin, What's the matter with tie-breaking? Improving efficiency in school choice, *American Economic Review* 98 (2008), 669-689
11. Ergin, H., Efficient resource allocation on the basis of priorities, *Econometrica* 70 (2002), 2489-2497.
12. Frechette, G., A. Roth, and U. Ünver, Unraveling yields inefficient matching: Evidence from post-season college football bowls, *RAND Journal of Economics*, forthcoming.
13. Gale, D. and L.S. Shapley, College admissions and the stability of marriage, *American Mathematical Monthly* 69 (1962), 9-15.
14. Harrison, G. and 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.
15. Kagel, J. and A. Roth, The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment, *Quarterly Journal of Economics* 115 (1) (2000), 201-235.
16. Kesten, O., On two competing rules for priority based allocation problems, *Journal of Economic Theory* 127 (2006), 155-171.
17. Kesten, O., School choice with consent, *Quarterly Journal of Economics* 125 (2010), 1297-1348.
18. McVitie D. and B. Wilson (1970), Stable marriage assignment for unequal sets, *BIT* 10, 295-309.
19. Mullin C., and D. Reiley, Recombinant estimation for normal-form games, with applications to auctions and bargaining, *Games and Economic Behavior* 54 (2006), 159-182.
20. Nalbantian, H. and A. Schotter, Matching and efficiency in the baseball free-agent system: An experimental examination, *Journal of Labor Economics* 13 (1995), 1-31.

21. Olson, M. and D. Porter, An experimental examination into design of decentralized methods to solve the assignment problem with and without money, *Economic Theory* 4 (1994), 11–40.
22. Pais, J. and Á. Pintér, School choice and information: An experimental study on matching mechanisms, *Games and Economic Behavior* 64 (2008), 303-328.
23. Rapaport, F. T., The case for a living emotionally related international kidney donor registry, *Transplantation Proceedings* 18 (1986), 5-9.
24. Roth, A., The evolution of labor market for medical interns and residents: A case study in game theory, *Journal of Political Economy* 92 (1984), 991-1016.
25. Roth, A., A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom, *American Economic Review* 81 (1991), 414-440.
26. Roth, A., The economics of matching: Stability and incentives, *Mathematics of Operations Research* 7 (1982), 617-628.
27. Roth, A. and E. Peranson, The redesign of the matching market for American physicians: Some engineering aspects of economic design, *American Economic Review* 89 (1999), 748-780.
28. Roth, A. and U.G. Rothblum, Truncation strategies in matching markets –In search of advice for participants, *Econometrica* 67 (1999), 21-43.
29. Roth, A. and M. Sotomayor, Two-sided matching, New York: Cambridge University Press (1990).
30. Roth, A., T. Sönmez, and U. Ünver, Kidney exchange, *Quarterly Journal of Economics* 119, no. 2 (2004), 457-488.
31. Roth, A., T. Sönmez, and U. Ünver, A kidney exchange clearinghouse in New England, *American Economic Review Papers & Proceedings*, 95-2 (2005a), 376-380.

32. Roth, A., T. Sönmez, and U. Ünver, Pairwise kidney exchange, *Journal of Economic Theory*, 125 (2005b), 151-188.
33. Roth, A., T. Sönmez, and U. Ünver, Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences, *American Economic Review*, 97-3 (2007), 828-851.
34. Roth, A. and X. Xing, Jumping the gun: Imperfections and institutions related to the timing of market transactions, *American Economic Review* 84 (1994), 992-1044.
35. Shapley, L. and H. Scarf, On cores and indivisibility, *Journal of Mathematical Economics* 1 (1974), 23-28.
36. Sönmez, T. and U. Ünver, House allocation with existing tenants: An equivalence, *Games and Economic Behavior* (2005) 52 , 153-185.
37. Sönmez, T. and U. Ünver, House allocation with existing tenants: A characterization, *Games and Economic Behavior*, (2010) 69, 425-445.
38. Sönmez, T. and U. Ünver, Course bidding at business schools, *International Economic Review*, forthcoming.
39. Ünver, U., Backward unraveling over time: The evolution of strategic behavior in the entry-level British medical labor markets, *Journal of Economic Dynamics and Control* 25 (2001), 1039-1080.
40. Zizzo, D., Experimenter demand effects in economic experiments, *Experimental Economics* (2010) 13, 75-98.