On-Campus Housing: Theory vs. Experiment*

Pablo Guillen†and Onur Kesten‡

August 2007

Abstract

Many universities in the US offer on-campus housing opportunities to incoming as well as already enrolled students. Most common student assignment mechanism used in the US has been shown to suffer serious efficiency losses. In this paper we first show that a particular placement mechanism which is in use at the MIT for about two decades is in fact equivalent to a natural adaptation of the well-known Gale-Shapley mechanism of two-sided matching theory to this framework. Motivated from the increasing popularity and success of the Gale-Shapley mechanism in a number of markets, we next experimentally compare the performances of the MIT mechanism with that of the leading theory mechanism Top Trading Cycles. Contrary to theory, the MIT mechanism performs better in terms of efficiency and participation rates, while we observe no significant difference between the two mechanisms in terms of truth-telling rates.

1 Introduction

A *house allocation problem* consists of a set of agents and a set of indivisible objects (e.g., houses) that needs to be distributed among agents. Typical examples are assignment of tasks to workers, offices to professors, houses to prospective tenants, etc. A commonly observed example of this problem in the universities in the US is the assignment of housing units (or, dormitory rooms) to students. Specific to this application, not all of the participants have equal right over each house prior to the central assignment procedure. There may, for example, be *existing tenants*, who may already be occupying a house, and who may still be seeking a better one. This variation of the problem is known as the *house allocation problem with existing tenants* (Abdulkadiroglu and Sönmez, 1998).

A house allocation problem with existing tenants consists of two pieces of information: (1) a *priority ordering* over all participants, determined by the assignment policies of the particular university based on seniority, GPA, result of a lottery draw etc.; (2) a *list of preferences* of each participant over housing types, typically a rank-ordered-list of housing types which a participant

*Our special thanks are due to Yan Chen and Tayfun Sönmez for their kind assistance to this project. We are grateful to Harvard Business School for its hospitality and financial support. We also thank David Reiley for useful comments, Siu Kho for software assistance, and seminar and conference participants at Harvard, Pittsburgh, Rochester, Boston College, SCW 2006 (Istanbul), RED 2006 (Bodrum), and Matching Workshop 2007 (Barcelona).

†The University of Sydney

‡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.

decides upon comparing different housing types available. An assignment *mechanism* is a systematic procedure that chooses an assignment of participants to available housing units based on the aforementioned two pieces of information.

The performance of a mechanism is basically evaluated on four merits: (1) *individual rationality* (i.e., an existing tenant should be encouraged to participate by giving the guarantee of a house that is no worse than her current house); (2) *efficiency* (i.e., resources should be optimally allocated according to the likings of participants);(3) *fairness* (i.e., the assignment should respect the priority order), and (4) *incentive compatibility* (i.e., participants should be induced to act straightforwardly and reveal true preferences). Abdulkadiroglu and Sönmez (1998) examine some of the real-life mechanisms used in universities in the US, and show that most mechanisms currently in use lack either efficiency or individual rationality. They show that, quite surprisingly, the most common mechanism in the US, the *random serial dictatorship with squatting rights (RSDwSR)* lacks both of these properties, mainly because it discourages existing tenants from participating in the assignment procedure, and consequently leads to potential losses from trade. Abdulkadiroglu and Sönmez (1998) propose an alternative mechanism called the *top trading cycles (TTC)* mechanism[1]. TTC fully achieves the first three properties. It achieves fairness in a weak[2] sense. Chen and Sönmez (2002) experimentally compare the performances of TTC and RSDwSR, and find TTC to be significantly more efficient than the popular real-life mechanism RSDwSR.

The *school choice problem* is an important extension of the present problem. Differently than a house allocation problem, in a school choice problem each school has a multiple capacity of students it can adnmit, and typically a distinct priority ordering (which, for each school, is determined acccording to specific policies of school districts). After being advocated as a promising school choice mechanism by the pioneers of school choice, the well-known Gale-Shapley mechanism of two-sided matching theory has gained increasing popularity also among school districts in the US and replaced two deficient mechanisms in the 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 school choice as well is in two-sided matching markets (such as the National Resident Match; Roth and Peranson, 1997), we offer a natural and intuitive adaptation of the Gale-Shapley mechanism to the present context. In terms of our desireta, this adaptation achieves individual rationality, incentive compatibility, and fairness.

First we show that this adaptation of the Gale-Shapley mechanism is in fact equivalent to a mechanism that has already been in use at MIT for about two decades (Theorem 1). To the best of our knowledge, this is the second reported equivalence of the well-known Gale-Shapley mechanism to a real-life mechanism. The first such report is due to Roth (1984) who showed that the mechanism used by the National Resident Matching Program in the US between 1954 and 1986 to assign medical interns to hospital positions is actually an equivalent of the Gale-Shapley mechanism.

Our equivalence result in turn implies that the mechanism used in MIT is indeed Pareto superior to any mechanism that respects the given priority ordering of participants (Corollary 1).

[1] Top trading cycles based mechanisms have been extensively studied in recent inidivisible goods allocation problems. Two such problems that attracted much attention are the *school choice problem* and the *kidney exchange problem.*

[2] Precisely speaking, in the following sense: Given a priority ordering of agents, under TTC it is possible that an agent may get a house worse for him than the house an agent with lower priority gets. In fact, when this is the case, the lower priority agent is an existing tenant.

Finally, in order to compare the performances of the MIT mechanism with the leading theory mechanism TTC when participants are boundedly rational, we conduct a laboratory experiment to test the two mechanisms. We find that the MIT mechanism performs better in terms of participation rates and it is at least as good as TTC in terms of efficiency. We do not observe any significant difference between the two mechanisms in terms of truth-telling rates.

2 The Model

Prior to the centralized assignment procedure each existing tenant chooses whether to participate or not. Then a *house allocation problem (with existing tenants)* (Abdulkadiroglu and Sönmez, 1998) is given by

- ◦ a finite set of existing tenants I_{E+}who have chosen to participate,
- ◦ a finite set of existing tenants I_{E-}who have chosen not to participate,
- ◦ a finite set of new applicants I_N,
- ◦ a finite set of occupied houses $H_O = \{h_i\}_{i \in I_{E+} \cup I_{E-}}$,
- ◦ a finite set of vacant houses H_V,
- ◦ an ordering f over all agents but the non-participating existing tenants, and
- ◦ a list of strict preference relations $P = (P_i)_{I_{E+} \cup I_N}$ and,

Often times we will suppress the first six components assumming that they are exogenously given and fixed. Let $I = I_{E+} \cup I_{E-} \cup I_N$ denote the set of all agents, and $H = H_O \cup H_V \cup \{h_0\}$ denote the set of all houses plus the null house. Every existing tenant $i \in I_E$ is endowed with (i.e., currently lives in) the occupied house $h_i \in H_O$. For expositional simplicity, we assume that the null house is the last option for each agent.

An *allocation* μ is a list of assignments such that (1) every agent is assigned one house; (2) no house other than the null house is assigned to more than one agent; and (3). Let $\mu(i)$ denote the assignment of agent i under μ.

A *mechanism* φ is a systematic procedure that chooses an allocation for each problem. Let $\varphi(P)$ denote the allocation chosen by φ for the problem P.

2.1 Properties of mechanisms

We next define four desirable properties of mechanisms:

Individual rationality: No participating existing tenant ever gets a house that is worse than her endowment. (i.e., for every $i \in I_{E+}$ and every problem P, $\varphi_i(P)\ R_i\ h_i$.)

Pareto efficiency: The outcome cannot be Pareto improved, i.e., there is no allocation at which all agents are at least as well off and at least one agent is strictly better off. (i.e., for every problem P, there is no μ such that $\mu(i)\ R_i\ \varphi_i(P)$ for all $i \in I$ and $\mu(i)\ P_i\ \varphi_i(P)$ for some $j \in I$.)

Fairness: Among all the agents but the non-participating existing tenants, if an agent ever prefers another agent's assignment, then either of the following two should be true: (1) the other agent has higher priority (according to the priority ordering); (2) the other agent is an existing tenant who is assigned her own house. (i.e., for every problem P and every $i, j \in I_{E+} \cup I_N$, if $\varphi_j(P)\ P_i$

$\varphi_i(P)$, then either (1) $f(j) < f(i)$ or, (2) $\varphi_j(P) = h_j$.)

Incentive compatibility (strategy-proofness): Among all the agents but the non-participating existing tenants it is a dominant strategy for each agent to truthfully report her preferences. (i.e., for every problem P, every $i \in I_{E+} \cup I_N$, and every P_i', $\varphi_j(P)\ R_i\ \varphi_i(P_i', P_{-i})$.)

We start with a negative result. It turns out any three of the above four properties are compatible except for the first three:

Proposition 1: *No mechanism is individually rational, Pareto efficient, and fair.*

3 Two mechanisms

3.1 Top Trading Cycles

Abdulkadiroglu and Sönmez (1998) propose the *top trading cycles (TTC)* mechanism which is based on Gale's top trading cycles idea. Mechanisms based on this idea have been proposed and extensively studied in the recent literature mainly for two other important problems: the *school choice problem* (Abdulkadiroglu and Sönmez, 2003) and the *kidney exchange problem* (Roth, Sönmez, and Ünver, 2005).

TTC works as follows: Consider a given house allocation problem with a given priority ordering f of agents. Assign the first agent (according to f) his top choice, the second agent his top choice among the remaining houses, and son 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 him to the top and proceed. Similarly, insert any existing tenant who is not already served at the top of the line once his or her house is demanded. If at any point, a loop forms, it is formed by exclusively existing tenants and each of them demands the house of the tenant next in the loop. (A *loop* is an ordered list of agents $(i_1, i_2, \ldots, i_k)$ where agent i_1 demands the house of agent i_2, agent i_2 demands the house of agent $i_3, \ldots$, 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. Chen and Sönmez (2002) report that TTC is significantly more efficient than the popular real-life mechanism RSD-wSR.

3.2 MIT-NH4 Mechanism

The following mechanism is in use at residence NH4 of MIT. Consider a given house allocation problem with a given priority ordering f of agents:

1. The first agent (according to f) is *tentatively* assigned his top choice among all houses, the next agent is *tentatively* assigned his 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 his 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 his 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 again with the conflicting agent. Every squatting conflict that occurs afterwards is resolved in a similar way.

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

It is not much difficult to show that the MIT-NH4 mechanism achieves all the desirable properties except Pareto efficiency.

4 A modified Gale-Shapley mechanism & an equivalence

The Gale-Shapley mechanism has long dominated two-sided matching theory due to its attractive stability[3] and incentive features. It has also been adopted by a number of real-world matching markets (see Roth and Rothblum, 1998 for an extensive list of these 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 problems. Shortly after its proposal for school choice by Abdulkadiroglu and Sönmez (1998), the Gale-Shapley mechanism has attracted the attention of education authorities in NYC and Boston, and replaced two controversial school choice mechanisms in these places. Even though school choice and the present problem are mathematically similar,[4] no counterpart of the popular Gale-Shapley mechanism has so far been considered for house allocation.

We first transform the present problem into a school choice problem, and next propose a direct adaptation of the Gale-Shapley mechanism. In a school choice problem, for each school there is a (possibly different) priority ordering determined based on certain criteria of school districts. Using the given priority ordering f of agents, first construct a priority ordering for a given house as follows:

(1) if it is a vacant house, then the corresponding ordering for this house is also f,

(2) if it is an occupied house, then assign the highest priority for this house to the corresponding existing tenant, and do not change the relative ordering of the remaining agents.

Given the constructed priority ordering for each house, the outcome of the *modified Gale-Shapley mechanism* is computed by applying the following *deferred acceptance* algorithm (Gale and Shapley, 1962):

Step 1: Each agent applies to his top choice house. For each house, consider its applicants. The agent with the highest priority according to the priority ordering for that house is tentatively placed. The rest are rejected.

[3] In two-sided matching a matching is stable if no two partcipants from the two sides of the market would refuse their current matches and rather form a blocking coalition with each other. See Kelso and Crawford for a large domain of preferences ensuring existence of stable matches and Roth and Sotomayor (1991) for a comprehensive survey on two-sided matching.

[4] The only differences between the two are: (1) in school choice for each school there is a separate (often different) priority ordering of students, and (2) in school choice individual rationality is irrelevant since there is no counterpart of existing tenants .

In general;

Step k: Each rejected agent applies to his next top choice house. For each 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. The rest are rejected.

The process is over when no agent is rejected any more.[5]

Much to our surprise the above natural modification of the Gale-Shapley deferred acceptance procedure in fact yileds the same outcome with the MIT-NH4 mechanism.

Theorem 1*: The MIT-NH4 mechanism and the modified Gale-Shapley mechanism are equivalent.*

Theorem 1, to the best of our knowledge, is the second time the Gale-Shapley deferred acceptance procedure is shown to coincice with a real-life mechanism. Roth (1984) showed that the National Resident Matching algorithm used in the US between 1954 and 1986 is a hospital-proposing deferred acceptance procedure. Interestingly, the MIT-NH4 mechanism is an agent-proposing (intern-proposing in the context of Roth, 1984) procedure.

The equivalence in Theorem 1 allows MIT-NH4 to claim all the attractive properties of Gale-Shapley mechanism. By Balinski and Sönmez (1998) the following corollary is now immediate.

Corollary 1*: The MIT-NH4 mechanism (as well as the modified Gale-Shapley mechanism) Pareto dominates any other fair mechanism.*

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

5 Experimental Design

Our design compares the performance 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. We tried to keep our design as close as possible to the one in Chen and Sönmez (2002).

We run five replications, or five independent groups, for each treatment. 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 [Fischbacher (2007)]. Each group consists of 12 participants. Participants #1 to #8 are existing tenants. Participants #9 to #12 are newcomers. There are also 12 houses of 8 different types to be allocated. House types go from A to H. Participants #1 to #12 are existing tenants, each living in a house type A to H. There are four additional vacant houses of types A, B, C and D. Table 1 shows the payment 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

[5] Note that since the capacity for the null house is unlimited, any agent who applies to it, is assigned this house.

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 to each payment on top of the payments 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 monetary salient difference of $14 between the top and the last choice.

Both treatments, NH4 and TTC, are implemented as one shot games of incomplete information. Each participant knew its own payoff table but not the others' payoff table. Participants did know the number of existing tenants and newcomers and that payoff tables may differ. In both treatments existing tenants are given an option to keep their houses and then not participate in the allocation mechanism.

The experiment was conducted as follows. Once each participant was allocated 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 one randomly chosen participant drew them one by one in order to generate the initial priority order. 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

	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

6 Experimental Results

To evaluate the aggregate performance of NH4 vs. TTC, we compare the efficiency generated by each mechanism. We look at three different efficiency measures - observed efficiency, expected efficiency and the recombinant estimation of mean efficiency. Observed efficiency is calculated by taking the ratio of the sum of actual earnings of all subjects in a session and the Pareto optimal earnings of the group. The unique Pareto optimal group earning is 231. Column 3 in Table 2 shows the observed efficiency for each group in the two treatments.

TABLE 2. EFFICIENCY (Standard Errors in Brackets)

Mechanisms	Group	Observed efficiency	Expected efficiency	Recombinant estimation of mean e
NH4	NH4-1	1	1 (0)	
	NH4-2	.8701299	.8738139 (0.025)	$\widehat{\mu} = 0.893$ (0.0923)
	NH4-3	.9047619	.8874434 (0.028)	$\sigma^2 = 0.0045$
	NH4-4	.8528138	.8445917 (0.028)	$\varphi = 0.0036$
	NH4-5	.8138528	.8050089 (0.033)	
TTC	TTC-1	.8095238	.8166418 (0.025)	$\widehat{\mu} = 0.813$ (0.0545)
	TTC-2	.8398268	.8742235 (0.017)	$\sigma^2 = 0.0018$
	TTC-3	.7705628	.7692916 (0.022)	$\varphi = 0.0012$
	TTC-4	.8354979	.8287112 (0.021)	
	TTC-5	.7532467	.7902112 (0.020)	

Result 1 (Observed Efficiency): A permutation test shows that the observed efficiency of NH4 is significantly higher than that of TTC: $p = 0.0167$ (one-tailed) for the original treatment.

Observed efficiency only takes into account the particular priority order randomly determined in the experimental lab. In order to obtain a measure as independent as possible of a particular priority order we calculate the expected efficiency for each of the five groups in our two treatments. Expected efficiency is computed by randomly generating 1 million priority orders for each group. Hence each priority order results in one allocation. For each allocation the ratio of the sum of total earnings is calculated. Finally, the expected efficiency for each group is the average calculated over the 1 one million ratios. Column 4 in Table 2 summarizes expected efficiency for each group.

Result 2 (Expected Efficiency): A permutation test shows that the expected efficiency of NH4 is significantly higher than that of TTC: $p = 0.0498$ (one-tailed) for the original treatment.

The truly one-shot nature of the experimental design allows for the use of the recombinant test techniques described in Mullin and Reiley (2006). The basis of this method is to recombine the strategies of different players in order to obtain the result if the grouping had been different. We randomly generated two million groups for each treatment and one priority order for each group. Then we estimated the mean, variance and covariance (see column 5 in Table 2) of the data in order to compute the recombinant z-value.

Result 3 (Mean Efficiency): A recombinant test shows how the mean efficiency of NH4 does not significantly differ from than of TTC. The recombinant estimation of mean efficiency is 89.3% for NH4 and 81.3% for TTC. Estimated variance is 0.0045 for NH4 and 0.0018 for TTC. The estimated covariances are 0.0036 and 0.0012 for NH4 and TTC respectively. Thus, the recombinant z-test yields $z = 0.21$ and $p = 0.42$.

We also tested whether participation rates and truthful revelation differ from NH4 to TTC.

Table 3. Participation and Truthful Preference Revelation

Mechanisms	Group	Participation rate	Proportion of truth
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 participation rates in column 3 and proportions of truthful preference revelation for each group in column 4.

Result 4 (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 T-test of proportions show that the participation rate of existing tenants under NH4 is significantly higher than that of TTC: $z = 2.7713$ ($p = 0.0028$).

Result 5 (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.

T-tests of proportions shows that the proportion of truthful preference revelation under NH4 is not significantly different from that of TTC: $z = 1.2250$ ($p = 0.1103$).

Results 1 to 3 show efficiency is at least not lower in NH4 than in TTC. Since we do not find significant differences in truthful preference revelation, result 5, we can conclude that participation is the key to understand why NH4 is outperforming TTC even the theory does not support this result.

7 Conclusion

Chen and Sönmez (2002) proposed TTC as a serious candidate to replace RSDwSR in house allocation systems often used by universities to allocate graduate students to campus housing. They had good reasons to do so: both in theory and in laboratory experiments TTC outperforms RSDwSR. TTC is Pareto efficient, individually rational and incentive compatible, but not fair. We analyzed the MIT house allocation mechanism known as NH4. We find out that NH4 is individually rational, incentive compatible and fair, but not Pareto efficient. By our equivalence result, NH4 is, however, the most efficient of all mechanisms that are fair. We designed an experiment in which NH4 and TTC go head to head in terms of comparing efficiency. Despite the theoretical advantage of TTC, NH4 turns out to be no worse than TTC in terms of efficiency. Part of this can be explained by the higher participation rate we find in NH4. This, we believe, might be because boundedly rational individuals may find NH4 much easier to understand than TTC feeling more reluctant to participate.

Our finding is also consistent with that of Chen and Sönmez (2006) whose experiments showed that for school choice applications (again, contrary to theory) the Gale-Shapley mechanism performs better in terms of efficieny than TTC. A second reason to be optimistic about the efficiency performance of NH4 comes from a result due to Ergin (2002): the Gale-Shapley mechanism tends to be more efficient as priority orderings for each school tend to be more 'correlated.' A feauture of the modiefied Gale-Shapley mechanism that might contribute to this possibility is that all the priority orderings for the modified Gale-Shapley mechanism (the equivalent of NH4) are, by construction, generated from the same ordering.

In practical terms, NH4 outperforms TTC which in turn outperforms RSDwSR. Our result suggests that a widespread replacement of RSDwSR by NH4 could be beneficial.

8 Appendix

Proof of Proposition 1: Suppose $I_N = \{1, 2\}$, $I_{E+} = \{3\}$, $H_V = \{a, b\}$, and $H_O = \{h_3\}$. Let the priority ordering f be 1-2-3. Suppose the preferences of the agents 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 that 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 to house a. This means, by individual rationality agent 3 should be assigned to house h_3. Then agent 1 is assigned to house b. But this clearly violates fairness.

Q.E.D.

Proof of Theorem 1: It is easy to show that for a given house allocation problem an allocation is individually rational, fair, and non-wasteful[6] if and only if it is stable for the corresponding marriage problem where house preferences are constructed from the priority ordering in the way described previously. Modified Gale-Shapley mechanism is stable, and therefore individually rational, fair, and non-wasteful. It is well-known that the outcome of the Gale-Shapley mechanism is preferred by each agent to any other stable allocation. Hence, it Pareto dominates any other stable mechanism.

We give a direct proof of Proposition 1. We show that for any given house allocation problem the set of existing tenants who are allocated their own house are the same under the two mechanisms. Then since both algorithms' outcomes are fair and non-wasteful, they have to choose the same allocation.

[6] An allocation is *non-wasteful* if no agent other than a non-participating existing tenant prefers an unassigned house to her current assignment.

McVitie and Wilson (1970) show that under the DA algorithm, the ordering according to which agents make proposals to their mates on the other side of the market has no effect on the outcome, and provide an equivalent version of the DA algorithm, where agents make their proposals according to any given ordering. Take any house allocation problem with a given ordering f of agents. To prove Proposition 1 we use the McVitie-Wilson version of the DA algorithm in which agents propose in turn according to the ordering f and where the priority order for each house is constructed as described previously in the text.

First consider the NH4 algorithm applied to the given problem. Consider the tentative (partial) assignment that is obtained at the end of the first squatting conflict. Suppose it is the turn of existing tenant i_1 with house h_{i_1}, and all available houses are worse for him than h_{i_1}. Since both algorithms are fair and non-wasteful, at this point the tentative assignment of each agent under the NH4 algorithm who has higher priority than agent i_1 is the same as her tentative assignment under the DA algorithm right before agent i_1 starts to make proposals (*).

For the NH4 algorithm suppose it is some agent $j \neq i_1$ who is currently assigned h_{i_1}. Then agent i_1 is permanently assigned house h_{i_1} and removed from the process, all tentative assignments starting with agent j and up to agent i_1 are erased, and the process starts over with agent j. Note that this is the same as removing agent i_1 (with his house h_{i_1}), without changing the relative ordering of the remaining agents under f, and starting the NH4 algorithm all over from the beginning.

By (*), under the DA algorithm when it is the turn of agent i_1 to move, he starts his proposals with his top choice house, and in turn gets rejected from every house that she prefers to h_{i_1} (because each such house is now tentatively assigned to some agent who has higher priority for it than i_1). Then she proposes to h_{i_1}. Since she has the highest priority for h_{i_1}, she is permanently assigned to h_{i_1}, and from this point on any agent whomever proposes to h_{i_1} is rejected. Since the ordering of agents' moves under the DA algorithm has no effect on the outcome, the outcome does not change if we start the algorithm over, and apply it to the reduced problem which is obtained by removing agent i_1 (with her house h_{i_1}) without changing the relative ordering of the remaining agents under f.

We next apply the NH4 algorithm to the reduced problem and identify the first squatting conflict. Using the above argument once again we show that the existing tenant that participates in this conflict is permanently assigned his house under both algorithms, and remove him from the problem. Applying this argument repeatedly we conclude that the set of existing tenants who are allocated their own house must be the same under the two algorithms.

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. and T. Sönmez, School choice: A mechanism design approach, *American Economic Review* 93 (2003), 729-747.

3. Balinski, M. and T. Sönmez, A tale of two mechanisms: Student placement, *Journal of Economic Theory* 84 (1999), 73-94.

4. Chen, Y. and T. Sönmez, School choice: An experimental study, *Journal of Economic Theory* 127 (2006), 202-231.

5. Dubins, L.E. and D.A. Freedman, Machiavelli and the Gale-Shapley algorithm, *American Mathematical Monthly* 88 (1981), 485-494.

6. Gale D. and L.S. Shapley, College admissions and the stability of marriage, *American Mathematical Monthly* 69 (1962), 9-15.

7. Kelso A. S. and V. P. Crawford, Job matching, coalition formation, and gross substitutes, *Econometrica* 50 (1982), 1483-1505.

8. McVitie D. and B. Wilson (1970), Stable marriage assignment for unequal sets, *BIT,* 10**,** 295-309.

9. 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.

10. 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.

11. 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.

12. Roth, A., The economics of matching: Stability and incentives, *Mathematics of Operations Research* 7 (1982), 617-628.

13. Roth, A. and E. Peranson, The effects of a change in the NRMP matching algorithm, *American Economic Review* 89 (1999), 748-780.

14. Roth, A and U.G. Rothblum, Truncation strategies in matching markets –In search of advice for participants, *Econometrica* 67 (1999), 21-43.

15. Roth, A. and M. Sotomayor, Two-sided matching, New York: Cambridge University Press (1990).

16. Roth, A., T. Sönmez, and U. Ünver. Kidney Exchange, *Quarterly Journal of Economics* 19, no. 2 (May 2004): 457-488.

17. Shapley, L. and H. Scarf, On cores and indivisibility, *Journal of Mathematical Economics* 1 (1974), 23-28.