

February 2011

USC LUSK CENTER FOR REAL ESTATE

ABSTRACT

Tiered Housing Allocation: an Experimental Analysis*

We study in the laboratory, a variant of the house allocation with existing tenants problem where agents are partitioned into tiers with different privileges. Members of higher tiers receive their allocation before those in lower tiers and can also take the endowment of a member of a lower tier if they wish to. In this tiered environment, we evaluate the performance of the modified versions of three well-known mechanisms - the Top Trading Cycle (TTC), the Gale-Shapley (GS) and the Random Serial Dictatorship (RSD). For all three mechanisms, we find low rates of participation (around 40%), high rates of truth-telling conditional on participation (around 90%) and efficiency levels that are high (above 90%) but below full efficiency. Also, of the three novelties introduced in our experiment tiered structure, multiple matches and known priority queue- only the last one has an impact on choices, with subjects being significantly more likely to participate the higher their position in the queue. Finally, the majority of subjects who do not play according to the theory still follow discernible patterns of participation and preference revelation.

JEL Classification: C78, D61 and D78

Keywords: house allocation, laboratory experiment and matching

Juan D Carrillo
Department of Economics
University of Southern California
3620 S. Vermont Ave.
Los Angeles, CA 90089-0253
USA
Email: juandc@usc.edu

Saurabh Singhal
Department of Economics
University of Southern California
3620 S. Vermont Ave.
Los Angeles, CA 90089
USA
Email: singhals@usc.edu

* We would like to thank Manuel Castro, Arya Gaduh, Rahul Nilakantan and Brijesh Pinto for helpful comments and Chris Crabbe for developing the software. We also thank the LUSK center for Real Estate and the Microsoft Corporation for financial support. All remaining errors are our own.

Submitted 07 February 2011

1 Introduction

This paper studies a version of the *house allocation with existing tenants* problem. In this problem, a set of indivisible goods must be allocated to a number of agents. Money exchanges are not feasible and some agents may be endowed with some of the goods. Examples of such situations include the assignment of offices to faculty members, on-campus housing to students, parking spaces to employees, etc. Following the existing literature, in the rest of the paper we refer to the indivisible goods as ‘houses’.

We investigate the class of one-sided matching problems where agents are *partitioned into tiers* with different privileges. Indeed, it is often the case that individuals belong to classes or groups with privileges that are identical within tiers and different across them. For example, in university departments, offices are allocated to full professors, associate professors, assistant professors, lecturers, and graduate students. Undergraduate housing is allocated to students in their senior, junior, sophomore and freshman year. Similar (and sometimes more rigid) tiered structures are present in other situations such as firms, fraternities or the armed forces. In such environments, individuals within a stratum are on a level playing field, but have rights which supersede the rights of agents in the strata below. In particular, a member of a higher tier is privileged in her ability to get her assignment before the individuals in lower tiers do, and take the endowment of a lower-tiered individual if she wants to.

A house allocation mechanism is a systematic procedure that assigns houses to prospective tenants, allotting at most one house to each agent. At the outset, some agents are endowed with a house (existing tenants) and some others are not (newcomers). Similarly, some houses are occupied by an agent while others are vacant. An allocation mechanism can be evaluated on four desirable properties: (a) *Pareto efficiency* (the houses should be optimally allocated given the preferences of the agents); (b) *fairness* (the assignment should respect the priority order); (c) *individual rationality* (an agent should be no worse-off by participating in the mechanism); and (d) *strategy-proofness* (agents should not benefit from misrepresenting their preferences).

Three leading allocation mechanisms have been proposed in the literature, each with different advantages. The most commonly used mechanism in real life applications, *the Random Serial Dictatorship with squatting rights* (RSD), satisfies properties (a,b,d) but not (c), thereby discouraging participation. The mechanism results in losses which, depending on the specificities of the payoff structure, may be substantial. To get around this problem Abdulkadiroğlu and Sönmez (1999) propose the *Top Trading Cycles* (TTC), a mechanism that satisfies properties (a,c,d) but not (b). Under this procedure, an agent

may end up with a worse allocation than someone below in the priority queue, a feature that may create tensions between agents due to issues related to envy or fairness. Finally, the well-known *Gale-Shapley* mechanism (GS) of two-sided matching theory (Gale and Shapley, 1962) has a natural counterpart in the housing allocation problem. Guillen and Kesten (2010) show that a mechanism used at one of MIT's undergraduate dormitories (the NH4) is theoretically equivalent to GS in the context of the housing allocation with existing tenants problem. This rule satisfies properties (b,c,d) but not (a). Naturally, Pareto inefficiencies create ex-post incentives for swaps.

With the imposition of a tiered structure, an individual in a higher tier can always expropriate the house of an individual in a lower one. As a result of this new hierarchical structure, two of the four properties above mentioned need to be adjusted. First, fairness must apply only to agents in the same tier. We call it *tiered fairness*. Second, since property rights across tiers are compromised, individual rationality can no longer be globally guaranteed. Following Title (1998), we consider the weaker notion of *tiered individual rationality*, which incorporates the fact that an agent may be forced to switch to another house if her endowment is preferred by another agent who belongs to a higher tier. The first step of our analysis consists in extending the three mechanism discussed above to a multi-tiered structure, which we denote tRSD, tTTC and tGS respectively. Under the proper modification of fairness and individual rationality, it is easy to show that the multi-tiered mechanisms keep the same theoretical properties as their single tier counterparts: tRSD satisfies all but tiered individual rationality, tTTC satisfies all but tiered fairness and tGS satisfies all but Pareto efficiency (Proposition 1). Therefore, tRSD is at a disadvantage as it implies that some agents should rationally opt out.¹ As for tTTC and tGS, we should observe differences in the final allocations between the two but not in the rates of participation or truthful revelation. Whether these predictions match the empirical behavior of agents in a controlled laboratory environment is the main subject of our research.

There exist two experimental tests on housing markets with existing tenants in a single tier environment. The seminal paper by Chen and Sönmez (2002), from now on [CS], compares RSD and TTC in a single shot experiment with 12 agents (9 tenants and 3 newcomers) and 12 houses under asymmetric information.² In that paper, TTC dominates RSD both in terms of efficiency (88% v. 74% of full efficiency) and participation rates (79% v. 47%) whereas no significant differences are found in truth-telling rates conditional on

¹Following the literature, Pareto efficiency is defined from the point of view of the participating agents only. This means that tRSD may result in Pareto inefficiencies over the set of all agents.

²As a robustness check, in a follow up paper Chen and Sönmez (2004) compare RSD and TTC under complete information (where subjects know the payoffs of everyone). The qualitative results are the same.

participation (71% v. 74%). More recently, Guillen and Kesten (2010), from now on [GK], compares TTC and GS (or, more precisely, the equivalent NH4 version) in an identical setting. Using an ordinal efficiency test, they find that GS is more likely to Pareto dominate TTC than the other way around. GS also yields more participation than TTC (78% v. 48%) and similar truth-telling rates (80% v. 69%).³

Our paper introduces four changes in the experimental design of [CS]. First and as already discussed, we consider a hierarchical structure with 12 agents divided into 4 tiers, with 2 tenants and 1 newcomer per tier. A stratified population seems empirically relevant for a large class of applications. Second, we compare all three mechanisms (tRSD, tTTC, tGS) within the same framework. Indeed, the results in [CS] and [GK] suggest that behavior in laboratory experiments on one-sided matching can vary substantially even when they follow identical protocols. It is therefore preferable to make comparisons between mechanisms in a framework that is as unified as possible. Third, we communicate to the agents not only their tier but also their positions in the priority queue *before* eliciting their participation decision and ranking of alternatives. This should have no effect on either tTTC or tGS, where participation and truthful revelation are dominant strategies. It should affect participation (but not ranking) under tRSD since agents have extra information about the likelihood of losing their endowed allocation. Fourth, we conduct 6 rounds of the game with an identical payoff matrix and set of players but with random re-assignment of preferences over houses and positions in the queue at the end of each round. In practice, these mechanisms are typically played only one or a few times. However, it is important to understand whether learning occurs if we want to provide some policy prescriptions. Repetitions also allow us to perform tests of individual behavior.

In general, the performance in our laboratory experiment is similar under the three mechanisms. In all cases, efficiency is high (above 90%) but below full efficiency, participation rates are low (around 40%) and truthful revelation rates conditional on participation are also high (around 90%). However, we also find some differences among them. Contrary to [CS], tRSD outperforms tTTC (and also tGS) in terms of cardinal efficiency and truthful revelation rates. Contrary to [GK], differences between tTTC and tGS are not statistically significant. The reason for the (surprising) dominance of tRSD, a theoretically inferior mechanism, lies on the combination of two factors: participation rates under tTTC and tGS are as low as under tRSD, and truthful revelation under tRSD is higher than under tTTC and tGS.

Interestingly, announcing the positions in the priority queue before eliciting choices has

³The paper also compares GS and NH4 and finds no difference in the laboratory between these two (theoretically equivalent) mechanisms.

a significant impact on the subjects' decision to opt out. Agents with an unfavorable draw in the priority queue realize that they are unlikely to improve their endowment and choose to play safe, that is, to keep their initial allocation. This is obviously suboptimal under tTTC and tGS, although it could be rationalized with an (ad-hoc) fixed cost of opting in, evaluating the alternatives or thinking through the game. Overall, among all the changes relative to the previous literature, knowing the priority queue has the greatest impact on behavior. Therefore, it is the best candidate to explain the differences between our results and those in [CS] and [GK]. As for the other features, we did not find any systematic tier effect on the behavior of subjects after controlling for potential gains of participation and opting out payoffs. Perhaps more surprisingly, we did not find any change in behavior over the course of the six matches under any mechanism. Still, having multiple observations per individual generated some interesting insights. In particular, many subjects never opt in and very few always do, suggesting that almost no subject realizes that participating is a dominant strategy under tTTC and tGS despite the transparency of the instructions. The pattern is the opposite when it comes to preference revelation: more than 80% of subjects *always* report truthfully and less than 10% misrepresent their preferences more than once, implying that the majority of subjects do realize the benefits of truthful revelation. Finally, we obtain an interesting result from the interaction of multiple matches and known priority queues. Indeed, the participation behavior of half of the subjects can be rationalized by a rule of the type "I participate only if my position in the queue is x or above", where x differs between subjects.

More generally and as noted above, there are substantial differences in the behavior of subjects across experiments playing under similar conditions.⁴ This suggests an important policy implication that deserves further consideration: changes in the protocol, experimental conditions or subject population may affect choices as much as changes in the mechanisms themselves.

Other related literature.

The theoretical literature on housing allocation includes the 'Hierarchical Exchange' mechanisms introduced by Pápai (2000) which characterizes the class of Pareto-efficient, reallocation-proof and group strategy-proof mechanisms. Pycia and Ünver (2007) introduce a new class of mechanisms, the 'Trading Cycles with Brokers and Owners', and show that all group incentive compatible and efficient mechanisms belong to that class. Under the tiered environment, Title (1998) develops the 'Tiered Exchange' mechanisms

⁴For example, participation rates under TTC are 79% in [CS] and 48% in [GK]. Participation rates under GS are 78% in [GK] and 40% in our experiment. Efficiency under RSD is 74% in [CS] and 93% in our experiment.

that characterizes the class of group strategy-proof, tiered individually rational, tiered envy-free and Pareto efficient mechanisms. Finally, the house allocation problem with existing tenants is related to other matching problems like the school choice problem (see e.g. Abdulkadiroğlu and Sönmez (2003)) and the kidney exchange problem (see e.g. Roth, Sönmez and Ünver (2004)).⁵

There exists also a rapidly expanding experimental literature on matching problems. Studies on one-sided matching problems include Olson and Porter (1994), Chen and Sönmez (2006), Pais and Pintér (2008), Calsamiglia, Haeringer and Klijn (2009) and Featherstone and Niederle (2009). The only study with a multi-tiered environment is the field experiment by Baccara et al. (2009) which analyzes the effect of network externalities in the housing allocation problem. A set of vacant offices are offered to faculty members partitioned into three groups. The priority queue of agents within each tier is known to everyone at the outset and the allocation is implemented using the RSD mechanism.⁶ The paper finds that externalities, such as institutional, co-authorship and friendship networks, strongly affect the subjects' choices. Finally, there is an experimental literature on two-sided matching (Nalbantian and Schotter (1995), Kagel and Roth (2000), Ünver (2005) and Haruvy, Roth and Ünver (2006)).⁷

2 The model

2.1 Basic definitions

Since we use a multi-tier extension of a model that is standard in the literature, we present it briefly (we refer the reader to Abdulkadiroğlu and Sönmez (1999) for a formal and comprehensive exposition of the single-tier version). A tiered housing allocation problem with existing tenants consists of a finite set of agents, exogenously partitioned into a finite number L of ordered sets or 'tiers', indexed by l . In each tier, agents are classified as either existing tenants (already occupying a house) or newcomers (not occupying a house). There is a finite set of houses, some of which are occupied by existing tenants while the others are vacant. Without loss of generality, we assume that the total number of agents is equal to the total number of houses. Also, there is a list of preference relations for each agent over all houses. Agents require at most one house. They have strict preferences over all houses and, without loss of generality, strictly prefer occupying a house rather than not.

⁵For a comprehensive survey of the literature see Sönmez and Ünver (2011).

⁶There are no existing tenants and occupied houses so the concept of tiered individual rationality does not apply.

⁷See Roth and Sotomayor (1990) for an early review of the literature on two-sided matching theory.

In a tiered environment, we model the privileges associated with belonging to a higher tier in two ways. First, members of higher tiers choose houses before those in lower tiers. The allocation procedure for tier l only begins after all the agents in tier $l - 1$ have received their allocation. Second, a member of a higher tier can expropriate the house occupied by an agent in a lower tier. Hence, the existing tenant in tier l who currently occupies a house is entitled to keep it only if it is not requested by some agent in tiers 1 to $l - 1$.

The outcome of a tiered allocation problem is a ‘matching’, that is, an assignment of houses to agents such that each agent is assigned at most one house and no house is assigned to more than one agent. The typical properties that a matching mechanism should promote are efficiency, fairness, participation and truthful revelation. However, due to the existence of hierarchical privileges, the standard definitions of fairness and participation need to be adapted to our stratified structure. We describe how these basic properties are defined in a tiered mechanism.

- (a) A matching mechanism \mathbf{M} is *Pareto efficient* if there is no other matching which assigns each agent a weakly preferred house and at least one agent a strictly preferred house.
- (b) A matching mechanism \mathbf{M} is *tiered fair* if, when an agent prefers another agent’s assignment, then (i) the other agent belongs to a higher tier, or (ii) the other agent is in the same tier and is ranked higher in the priority queue, or (iii) the other agent is in the same tier and is assigned her own house.⁸
- (c) A matching mechanism \mathbf{M} is *tiered individually rational* if no agent gets a house that is worse than her endowment at the beginning of the allocation process for her tier.⁹
- (d) A matching mechanism \mathbf{M} is *strategy proof* or incentive compatible in dominant strategies if no agent can benefit by unilaterally misrepresenting her preferences independently of the preferences and announcements of the other agents.

Notice that Pareto efficiency ensures that agents do not want to ex-post exchange their allocations. Tiered fairness in the present context corresponds to the well-known

⁸The analogue of this property in the single tier environment is sometimes called *justified envy-freeness*. Title (1998) introduces *tiered envy-freeness* to describe a matching where agents prefer their allocation to that of subjects in strictly lower tiers. By definition, all the mechanisms discussed in this paper are *tiered envy-free*, as members of higher tiers receive their allocation before those in lower tiers.

⁹This corresponds to her initial endowment only if she does not lose it to someone in a higher tier. If she does lose it, then it corresponds to a randomly selected house that was either vacant or previously occupied by someone in a higher tier and vacated.

(within tier) pairwise stability condition in the college admission problem. Tiered individual rationality guarantees participation conditional on the endowment inherited *after* the choice of agents in higher tiers. Finally, strategy proofness is defined in dominant strategies, which ensures that truthful revelation should not be affected by the (rational or irrational) participation and announcement decision of other players in the game.

2.2 The tiered allocation mechanisms

In this section we describe the tiered versions of the Top Trading Cycles (tTTC), Gale-Shapley (tGS) and Random Serial Dictatorship with Squatting Rights (tRSD) mechanisms. All mechanisms $\mathbf{M} \in \{\text{tTTC}, \text{tGS}, \text{tRSD}\}$ have a common structure. The allocation of houses is done sequentially, by tiers. We start with tier 1 while participants in tiers 2 to L wait. When the allocation for tier 1 is completed, we proceed to tier 2 while participants in tiers 3 to L wait, and so on.

Starting with the top tier, the allocation for each tier l takes place in the following way. First, all the houses that have been assigned to members of tiers 1 to $l-1$ are not available anymore. If an existing tenant's house has not been allocated to someone at a higher tier, then the agent continues to occupy her house. If, however, that agent's house has been taken by a member of a higher tier, she is compensated with a randomly selected house that was either vacant or occupied by someone in a higher tier and vacated. Second, an ordering of agents in tier l , also called a priority queue, is randomly chosen from a given distribution of orderings. The ordering is communicated to all the members of the tier, who also observe the house occupied by each agent in the tier. Third, each existing tenant in the tier decides whether to participate in the allocation mechanism or opt out. Those who opt out are assigned their houses and removed from the process. Fourth, the participating agents, i.e., the existing tenants who opt in and the newcomers, report their preferences over all the remaining available houses. Fifth, using the priority queue and the submitted preferences of the agents, the allocation in tier l takes place according to the mechanism \mathbf{M} being used. The agents and their allocations are removed from the process, and the same steps are repeated in tier $l+1$.¹⁰

The procedure always follows these general principles. The next sections describe the modifications included in each of the three mechanisms \mathbf{M} in order to account for the tiered structure of our problem.

¹⁰Notice that there is an important difference in the timing relative to the literature, where the ordering is chosen *after* the participation decision and preference revelation of agents. The reasons for this choice are twofold. First, to study the effect of this new feature relative to the existing experimental literature. Second, to ensure that the equilibrium allocations in tTTC and tGS are unique and different from each other.

2.2.1 Tiered Top Trading Cycle (tTTC)

Building on Gale’s top trading cycle idea (Shapley and Scarf, 1974), Abdulkadiroğlu and Sönmez (1999) proposed the *top trading cycle* (TTC) mechanism.¹¹ Consider the agents in tier l who decide to opt in. Using the preferences submitted and the exogenously determined priority queue, the allocation procedure in our tiered-modified version of the mechanism works as follows. Starting from the top of the queue, assign each agent in turn her top choice from among the available houses until someone requests the house of an existing tenant in the same tier. 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 the existing tenant to the top of the queue and proceed. Similarly, insert any existing tenant in tier l who has not yet been assigned a house at the top of the queue once her house is demanded. If at any point, a loop forms, it can only be formed by existing tenants in the same tier where each tenant demands the house of the tenant next in the loop.¹² In such cases remove all agents in the loop by assigning them the houses they demand and proceed. Once all the agents in tier l have received their allocation, move on to tier $l + 1$. When the number of tiers is 1, this procedure reduces to the TTC mechanism proposed by Abdulkadiroğlu and Sönmez (1999, p.251).

2.2.2 Tiered Gale-Shapley (tGS)

[GK] have adapted the GS mechanism to the house allocation with existing tenants problem. We simply extend that mechanism to our multi-tiered environment. For tier l and using the exogenously determined priority queue of agents, construct a priority ordering for each available house as follows. If the house is vacant or occupied by someone in a lower tier, then the corresponding queue for this house is the same as the priority queue of the agents. If the house is currently occupied by someone in the same tier, then assign the highest priority for this house to the corresponding existing agent, and assign the remaining priorities without changing the relative ordering of the remaining agents. Then, using this priority ordering for each house and the submitted preferences of the agents, apply the following version of the *deferred acceptance algorithm* due originally to Gale and Shapley (1962):

Step 1. Each agent applies to her top choice house. For each house, look at its pool of

¹¹This version is called “you request my house - I get your turn” in Abdulkadiroğlu and Sönmez (1999) and shown to be theoretically equivalent to Gale’s idea of top trading cycles in the context of housing allocation with existing tenants. This version of TTC has been used in the experiments of [CS] and [GK].

¹²A 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 , ..., agent i_k demands the house of agent i_1 .

applicants and tentatively assign the house to the agent with the highest priority according to the priority ordering for that house and reject the rest of the applicants.

In general,

Step k. Each rejected agent applies to her 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, assign the house to the agent with the highest priority according to the priority ordering for that house and reject the rest.

The process for tier l ends when no agent in that tier is rejected. At that point, all tentative assignments are finalized, and we move on to tier $l + 1$.

2.2.3 Tiered Random Serial Dictatorship with Squatting Rights (tRSD)

The RSD algorithm is the simplest and most commonly employed one. In each tier l , the tRSD determines the allocation by simply working its way down the priority queue, assigning each participating agent, in turn, her top choice from among the remaining available houses. The unassigned houses are then passed on to the next tier.

Under this mechanism, existing tenants may not wish to participate since they are not guaranteed a house that is at least as good as the one they are assigned at the beginning of the allocation process for their tier, resulting in loss of potential gains from trade. However, since members of a higher tier get to choose before those in lower tiers, it is tiered fair. Note that once an existing agent decides to participate, truthful preference revelation is a dominant strategy.

2.3 Properties

Allocation mechanisms are evaluated (both theoretically and empirically) according to the properties they satisfy. It is well known that for the case of a single tier, no (one-sided or two-sided) matching mechanism satisfies simultaneously the four properties described in section 2.1.¹³ For the same reason, no mechanism can satisfy the four properties in our tiered environment either. The proposition below describes which properties hold under each of the tiered mechanisms.

Proposition 1 *For any ordering of agents in each tier, the mechanisms satisfy the following properties:*

- *tTTC is Pareto efficient, tiered individually rational and strategy-proof;*
- *tGS is tiered fair, tiered individually rational and strategy-proof;*

¹³See e.g. [GK] and the references therein.

- *tRSD is Pareto efficient, tiered fair and strategy-proof.*

Proof. The notions of fairness and individual rationality have been adequately modified to be applied to each tier sequentially. The allocation algorithm of each mechanism has also been modified only to be applied to each tier sequentially. Since the properties hold for any given tier, the sequential application by tiers implies that they must necessarily hold for the tier modified mechanism. \square

The idea is simple. In a stratified population, the notions of fairness and participation must be modified in order to be applied to each tier sequentially, that is, in order to ensure identical privileges within tiers and hierarchical privileges between tiers. The key issue is then to find the proper definition of fairness and participation. Once this is achieved, it is straightforward to extend the logic of the single-tier mechanisms to our multi-tiered environments. Not surprisingly, the very same properties that have been shown for the single-tier setting also hold in our multi-tier framework. In particular, tTTC lacks tiered fairness as an agent may get a house that is worse than the one assigned to an existing tenant lower in the priority queue. Similarly, tGS lacks Pareto efficiency as its outcome may be Pareto dominated by a matching that lacks tiered fairness.¹⁴ Finally, tRSD lacks tiered individual rationality since an agent who participates has to give up her endowment. If her position in the priority queue is low, she is likely to end up with a house which is worse than if she opted out. Naturally, whether such theoretical properties hold in practice is an empirical question. [CS] and [GK] document interesting deviations from the theoretical predictions in the single-tier case. The experimental analysis conducted in the next sections report other deviations in our multi-tier environment with known priorities and repetition.

3 Experimental Design

Our experiment is designed to compare the tTTC, tGS and tRSD mechanisms in terms of participation of existing tenants, truthful preference revelation by participating agents and overall efficiency. To this purpose we conduct three treatments which differ exclusively in \mathbf{M} , the allocation mechanism employed.

For each treatment we ran three sessions with twelve subjects per session for a total of 108 subjects. The subjects were undergraduate students at the University of California, Los Angeles who were recruited by email solicitation. All sessions were conducted at the California Social Science Experimental Laboratory (CASSEL). The interaction between

¹⁴See Abdulkadiroğlu and Sönmez (1999) for a simple example demonstrating that the NH4 outcome may not be Pareto efficient.

subjects was fully computerized using an extension of the open source software package Multistage Games.¹⁵ No subject participated in more than one session.

In each session, participants played a total of 6 matches. In the first match, the twelve participants were randomly assigned a role, labeled 1 to 12, and divided into 4 tiers of 3 participants each. Subjects in roles 1-2-3 were in tier 1, subjects in roles 4-5-6 were in tier 2, and so on. Subjects in roles 1-2-4-5-7-8-10-11 were existing tenants while subjects in roles 3-6-9-12 were newcomers. Hence, each tier consisted of two existing tenants and one newcomer. There were twelve different houses to be allocated, labeled A to L.

Table 1 shows the monetary payoff to each participant (in dollars) as a function of the house she holds at the end of the allocation process. The square bracket, $[\cdot]$, indicates that the participant is occupying that particular house at the beginning of the match. Note that due to the existence of hierarchical privileges, existing tenants in tiers 2 to 4 might not be occupying these particular houses when the allocation mechanism reaches their tier. However, they will be occupying some house. The payoffs differ from previous experiments and the priority queues are also pre-specified and announced to the participants. Payoffs and queues are chosen with several objectives in mind. First, there are four possible Pareto efficient house allocations within each tier, giving a total of $4^4 = 256$ aggregate Pareto efficient house allocations, all with an identical aggregate payoff of \$235. Second, the initial payoffs are such that no existing tenant (except role 10) is occupying her most preferred house (again, remember that endowments may change for agents in tiers 2 to 4). Third, payoffs range from \$4 to \$25 providing significant variation. Fourth, with these payoffs and queues the equilibrium allocations under tTTC and tGS are unique and different from each other.

All treatments are implemented as games of incomplete information. At the beginning of each match, each subject knows her own tier, her position in the priority queue, and her payoff of holding each of the twelve houses. Subjects also know the number of tiers, the number of existing tenants and newcomers in each tier and that the payoff tables of other participants may differ. Finally, when the allocation process reaches their tier, they know which houses are not available anymore, which house is occupied by someone in the same tier with a higher position in the priority queue (if any), which house is occupied by someone in the same tier with a lower position in the priority queue (if any), which houses are occupied by members of lower tiers, and which houses are vacant.

¹⁵This contrasts with [CS] and [GK] who ran their experiment fully and partly by hand, respectively. We do not think that this minor modification had an impact on the subjects' behavior. On the other hand, having the game fully computerized allowed us to play multiple matches in a relatively short period of time (6 matches in 1 hour).

Table 1: Payoff matrix

Tier	Role	Houses											
		A	B	C	D	E	F	G	H	I	J	K	L
1	1	[19]	22	11	5	13	10	7	4	6	17	15	8
	2	25	[22]	5	11	10	13	4	7	8	17	6	15
	3	25	22	13	7	8	15	10	11	5	17	4	6
2	4	5	25	[17]	19	4	6	15	13	10	22	8	11
	5	22	17	25	[19]	6	8	15	5	11	4	10	13
	6	22	17	25	19	6	11	15	4	8	5	13	10
3	7	8	6	11	22	[17]	19	5	10	13	25	15	4
	8	17	4	6	25	22	[19]	11	13	10	8	15	5
	9	17	4	25	10	22	19	5	8	13	6	15	11
4	10	11	5	8	17	10	4	[25]	22	15	13	6	19
	11	13	11	4	10	5	25	8	[17]	15	6	22	19
	12	10	17	5	13	25	4	6	22	15	8	11	19

In each session, the allocation of houses is done sequentially, by tiers. We start with tier 1 while participants in lower tiers wait. The existing tenants first have the option of keeping their current house (by opting ‘out’) or participating (by opting ‘in’). If they opt out, they keep their own house and the process ends for them. The existing tenants opting ‘in’ and the newcomers are simultaneously asked to submit their preferences over the remaining houses.¹⁶ The participants are assigned houses according to the allocation mechanism \mathbf{M} ($\in \{tTTC, tGS, tRSD\}$) employed in that session, and these houses become unavailable for the lower tiers. Once the allocation for tier 1 is over, we move on to tier 2. If an existing tenant has lost her house to a member of a higher tier, then she is compensated with a randomly selected house which was either vacant at the beginning of the match or previously occupied by a member of a higher tier and vacated in the process. The allocation process in tier 2 then follows the same steps as in tier 1. The process continues with tiers 3 and 4, at which point the match ends.

At the end of a match, subjects are randomly reassigned a role (1 to 12). Roles 1-2-3 are always associated with tier 1, 4-5-6 with tier 2, and so on. The payoffs associated with each role are always the same as described in Table 1. Subjects in roles 1-2-4-5-7-8-10-11 are always existing tenants and subjects in roles 3-6-9-12 are always newcomers. The only

¹⁶When submitting their rankings, the participants know the total number of existing tenants and newcomers who have decided to opt in, but not their specific roles.

difference is that the position in the priority queue for a participant in a given role changes from match to match.¹⁷ Once the new roles and priorities are reassigned, subjects play the same game under the same rules and using the same allocation mechanism. In each session subjects made decisions over a total of 6 matches.

At the beginning of each session, instructions were read by the experimenter standing on a stage in the front of the experiment room. The experimenter fully explained the rules with special emphasis on the details of the allocation procedure and answered all questions. Next, the subjects went through a practice match in order to familiarize themselves with the computer interface and procedures. Subjects had to complete an interactive computerized comprehension quiz before they could proceed to the paid matches. Subjects were then asked to make their decisions over 6 matches. After each match, subjects were randomly reassigned a role. As explained above, the payoffs associated with each role were identical in all matches but the positions in the priority queue for an agent in a given role changed. At the end of the session, one of the matches was randomly selected and subjects were paid privately, in cash, their earnings in that match. Each session lasted for an average of 1 hour and the average earnings in each session were \$18 plus a show up fee of \$5. Table 2 summarizes the details of each session. Appendix B provides the instructions used in sessions 1-3 for the tTTC mechanism.¹⁸

Table 2: Session details

mechanism	session	date	# of subjects	# of matches
tTTC	1	7/20/10	12	6
	2	7/20/10	12	6
	3	7/20/10	12	6
tGS	1	7/21/10	12	6
	2	7/21/10	12	6
	3	7/21/10	12	6
tRSD	1	7/22/10	12	6
	2	7/22/10	12	6
	3	7/22/10	12	6

¹⁷We changed the ordering of the queue between matches to make sure that even if an individual happens to draw the same role as in a previous match, she still faces a different decision problem. However, the six orderings we selected for the six matches all satisfy the properties described above.

¹⁸Instructions for the other sessions are similar and available from the authors upon request.

4 Results

We evaluate the performance of the three mechanisms on three fronts: efficiency, participation rates and truthful preference revelation.

4.1 Efficiency

First we compare the relative efficiency of the three mechanisms using a cardinal measure. The empirical efficiency is calculated as the ratio of the sum of actual earnings to the unique Pareto efficient earnings (\$235). Table 3 reports the empirical efficiency for each match in each session under each mechanisms. It also reports the theoretical efficiency which, by definition, is 1.0 for tTTC and strictly smaller than 1.0 for tGS. For tRSD, two theoretical benchmarks are computed: the efficiency if all existing tenants were forced to participate and the efficiency if all tenants knew the preferences of the other agents (and therefore could evaluate the optimality of opting in).

Empirical and theoretical efficiency can be compared using a simple t-test. Empirical efficiency is significantly lower than theoretical under tTTC ($p < 0.001$) while there is no statistical difference under tGS ($p = 0.221$). Empirical efficiency under tRSD is significantly lower than theoretical conditional on all agents participating ($p < 0.001$) but significantly higher than theoretical under full information ($p < 0.001$). Finally, we can compare the empirical efficiency between the mechanisms. Wilcoxon-Mann-Whitney tests shows that while empirical efficiencies are not statistically different between tTTC and tGS ($p = 0.71$), the empirical efficiency of tRSD is significantly higher than that of tTTC ($p = 0.005$) and tGS ($p = 0.046$). This first result is summarized as follows.

Result 1 (Cardinal Efficiency) *Efficiency is high in all three mechanisms: 0.90 in tTTC, 0.91 in tGS and 0.93 in tRSD. Only for tTTC the efficiency is below the theoretical prediction. The empirical efficiency of tRSD is higher than the empirical efficiency of both tTTC and tGS.*

The result highlighting the efficiency dominance of tRSD is surprising in light of the previous experimental literature. Notice in particular that no match in tRSD has an efficiency level below 0.89, whereas efficiency can reach levels as low as 0.85 under tTTC and 0.84 under tGS.

[GK] suggest an *ordinal efficiency test* (OET) to evaluate the relative efficiency of different allocation mechanisms. For example, to compare the allocations under tTTC and tGS the test works as follows. We pick each outcome under tTTC and Pareto compare it to each outcome under tGS. We then count the number of times tTTC dominates tGS and

Table 3: Empirical efficiency

Session	Match	Efficiency		
		tTTC*	tGS**	tRSD***
1	1	0.89	0.94	0.90
	2	0.90	0.94	0.92
	3	0.92	0.84	0.94
	4	0.90	0.90	0.98
	5	0.90	0.94	0.95
	6	0.95	0.92	0.90
2	1	0.91	0.90	0.92
	2	0.92	0.97	0.92
	3	0.90	0.90	0.92
	4	0.92	0.90	0.89
	5	0.92	0.84	0.92
	6	0.90	0.97	0.92
3	1	0.90	0.90	0.95
	2	0.87	0.90	0.93
	3	0.90	0.90	0.90
	4	0.90	0.90	0.96
	5	0.95	0.90	0.94
	6	0.85	0.90	0.94
Overall		0.90	0.91	0.93

* Theoretical = 1.00; ** Theoretical = 0.90

*** Theoretical if all participate = 1.00, Theoretical under full information = 0.90

the number of times tGS dominates tTTC, and use a sign rank Wilcoxon test for equality of the matched pairs of dominations (comparisons between tTTC and tRSD and between tGS and tRSD are done in a similar manner). Since we have 18 matches under each mechanism, we get a total of $18 \times 18 = 324$ paired comparisons. We are unable to reject the null hypothesis of equality in the number of times that one mechanism dominates another. Out of 324 comparisons, tTTC dominated tGS 37 times and tGS dominates tTTC 14 times ($p = 0.448$), tGS never dominates tRSD and tRSD dominates tGS 1 time ($p = 0.317$) and, finally, tTTC never dominates tRSD and tRSD never dominates tTTC.

Result 2 (Ordinal Efficiency) *Using the OET, we find that no mechanism is more likely to Pareto dominate the others.*

[GK] argue that since we are concerned with eliciting ordinal preferences from agents,

the OET is the most appropriate test. While we concur with the argument, it is also true that (non-equilibrium) deviations from truthful revelation are likely to be affected by the cardinality of payoffs. Hence, empirical performance should take into account such effect. More importantly, in [GK] the ordering is announced once the agents have made their decision. The efficiency comparison can then be performed with respect to the 10,000 possible priority orderings. In contrast, we announce the ordering before eliciting preferences, so it only makes sense to consider that priority ordering. This results in a dramatic decrease in the number of observations which reduces the statistical power of the OET test.¹⁹

4.2 Participation Rates

Next we compare the mechanisms on the basis of participation rates. Table 4 shows the participation rates under the different mechanisms broken down by session, match, tier and position in the priority queue. In general, the participation rate of existing tenants is low: 36.8% for tTTC, 39.6% for tGS and 41.0% for tRSD. A one-tailed t-test of proportions shows that the overall participation rates are not statistically different between tTTC and tGS ($p = 0.627$), between tTTC and tRSD ($p = 0.765$) or between tGS and tRSD ($p = 0.595$). Furthermore, differences in participation rates between mechanisms are not statistically significant for any given tier, any match, any position in the queue or for existing tenants who have or have not lost their original house.

Result 3 (Overall Participation) *Participation rates are low in all three mechanisms: 36.8% for tTTC, 39.6% for tGS and 41.0% for tRSD. Also, differences across mechanisms are not statistically significant.*

Next we look at the determinants of participation rates within mechanisms. First, standard t-tests of proportions performed on the data in Table 4 reveals that tenants in tiers 3 and 4 are more likely to participate than tenants in tiers 1 and 2 under all three mechanisms ($p = 0.025$ for tTTC, $p = 0.061$ for tGS, $p = 0.062$ for tRSD). Participation is particularly low in tier 1. The test, however, does not control for other factors, such as the payoff of opting out (which depends on her endowment at the beginning of the allocation process for her tier) or the maximum possible gain. Second and again without controlling for other factors, we find that under all mechanisms an existing tenant is more likely to participate when she is higher in the priority queue. On aggregate, participation rates in priority positions 1, 2 and 3 are 61%, 42% and 24% respectively. Also, t-tests of proportions show that participation rates are significantly higher in position 1 than in 2

¹⁹As explained above, we have 6 matches in each session for a total of 18 matches per mechanism.

Table 4: Existing tenants' participation decision

		tTTC	tGS	tRSD
Session	1	21/48	25/48	21/48
	2	21/48	15/48	18/48
	3	11/48	17/48	20/48
Tier	1	8/36	9/36	13/36
	2	12/36	14/36	11/36
	3	15/36	13/36	17/36
	4	18/36	21/36	18/36
Match	1	6/24	9/24	7/24
	2	11/24	8/24	10/24
	3	9/24	7/24	8/24
	4	8/24	9/24	11/24
	5	9/24	10/24	11/24
	6	10/24	14/24	12/24
Priority	1	29/45	24/45	29/45
	2	9/27	11/27	14/27
	3	15/72	22/72	16/72
Original House	Lost	16/19	13/18	14/20
	Not Lost	37/125	44/126	45/124
Overall		0.37	0.40	0.41

or 3 for tTTC ($p=0.01$ and $p < 0.001$), in position 1 than in 3 for tGS ($p = 0.014$) and in positions 1 and 2 than in 3 for tRSD ($p < 0.001$ and $p = 0.004$). The result is an accordance to theory for tRSD, where participation is a dominant strategy for the first agent in the queue and becomes less interesting as the likelihood of losing the endowment to someone higher in the queue increases. Under tTTC and tGS participation is (weakly) dominant independent of the position. In general, it reflects the idea that participants are reluctant to opt in if they rationally (for tRSD) or irrationally (for tTTC or tGS) believe that participation may result in a loss. Third, we analyze the evolution of participation over the course of the experiment. T-tests of proportions reveal no significant differences in participation rates between match 1 and match 6, between matches 1-2 and matches 5-6 or between matches 1-3 and matches 4-6, under any mechanism. One could a priori think that agents would learn through repetition that participation is always optimal under tTTC and tGS. Surprisingly, this does not seem to be the case in our experiment. A possible reason is that we provide very limited feedback to our subjects (only the final

vector of allocations and the subject's own payoff). In any case, that result deserves further exploration.

We now turn to examine other factors that may be affecting the decision of agents to participate. Table 5 shows the number of agents who occupy a house (# tenants) and their participation rate (IN) as a function of two variables: (i) the payoff of opting out (own house), that is, the value of the house they occupy when deciding whether to participate and (ii) the potential gain of opting in (max. gain IN) defined as the maximum possible net gain from participation.

Table 5: Participation as a function of maximum gain in and own house payoff

	tTTC		tGS		tRSD	
	# tenants	IN	# tenants	IN	# tenants	IN
Own house						
25	2	.00	5	.20	1	1.00
22	18	.22	18	.22	20	.20
19	64	.42	58	.38	64	.47
17	51	.31	51	.41	51	.35
15	6	.50	9	.67	7	.86
10	1	1.00	-	-	-	-
8	-	-	-	-	1	.00
6	1	1.00	1	1.00	-	-
5	-	-	1	1.00	-	-
4	1	1.00	1	1.00	-	-
Max. gain IN						
≥ 9	3	1.00	3	1.00	1	.00
7-8	6	.50	9	.56	9	.67
5-6	21	.43	21	.43	16	.44
3-4	64	.39	57	.35	65	.42
0-2	50	.26	54	.37	55	.35

Outside option and potential gains have the expected effect: participation by existing tenants increases as the payoff of keeping their own house decreases and as the maximum attainable gain from participation increases. This suggests that even if subjects in tTTC and tGS do not always participate contrary to their best interest, the comparative statics

on participation retain some degree of rationality. Once again, behavior is similar across mechanisms.

In order to investigate in more detail the determinants of participation, we perform two logit regressions where the dependent variable is a discrete choice variable that takes the value 1 if the existing tenant opts in and 0 otherwise. The first specification uses the maximum possible payoff gain as an explanatory variable (columns 2-4) whereas the second specification uses the payoff of keeping her own house (columns 5-7). The common set of independent variables include dummies for tiers, position in the priority queue, matches and whether or not the existing tenant has lost her original house. Table 6 reports the results.

Table 6: Logit models of participation decisions

	tTTC	tGS	tRSD	tTTC	tGS	tRSD
Tiers 3&4	0.720 (0.459)	0.599 (0.496)	0.747 (0.462)	0.346 (0.395)	0.203 (0.479)	0.568 (0.428)
Position 2	-0.798 (0.538)	-0.037 (0.499)	-0.062 (0.576)	-0.759 (0.535)	-0.105 (0.546)	-0.048 (0.587)
Position 3	-1.461*** (0.526)	-0.667 (0.492)	-1.668*** (0.566)	-1.518*** (0.548)	-0.916* (0.538)	-1.692*** (0.572)
Lost house	1.148 (0.875)	0.764 (0.763)	0.357 (0.567)	1.500* (0.826)	0.737 (0.764)	0.511 (0.523)
Max. possible gain	0.214* (0.875)	0.161* (0.763)	0.100 (0.102)			
Own house payoff				-0.083 (0.097)	-0.200** (0.090)	-0.058 (0.099)
Constant	-1.758* (0.960)	-1.054* (0.592)	-0.985 (0.717)	0.751 (2.021)	3.508* (1.883)	0.505 (2.071)
N	144	144	144	144	144	144
Pseudo R^2	0.199	0.116	0.161	0.188	0.132	0.159

All regressions included dummies for the matches. Clustered standard errors in parenthesis.

*, **, ***: Significant at 90%, 95% and 99% confidence level, respectively.

The first specification confirms to a large extent the previous results on mean comparisons. After controlling for maximum possible gain, we find that participation is unaffected by tier and negatively affected by being last in the priority queue (under tTTC and tRSD). Maximum gain has a positive effect on participation under tTTC and tGS. Finally, no

dummy variable for match is statistically significant (results not reported for brevity). The second specification is also informative. The position in the priority queue has still a significant effect on participation (specially for tRSD and tTTC) and tier has still no impact. Loss of the original house makes an existing tenant more likely to participate under tTTC, while payoff in own house has a significant negative impact on participation under tGS. Once again, the match dummy has no effect. The determinants of participation are summarized in the following result.

Result 4 (Determinants of Participation) *Participation increases when (i) agents are higher in the priority queue, (ii) the maximum possible gain is high and (iii) the payoff of opting out is low. Participation does not change over time or across tiers. These determinants are similar across mechanisms.*

An advantage of having subjects play multiple matches is that we can study behavior at the individual level. In particular, it is instructive to determine whether participation is bimodal (with some individuals always participating and others never) or if every individual sometimes participates and sometimes not. Table 7 presents the frequency of participation as a function of the number of times a subject was assigned the role of existing tenant (tenant). There are very few subjects who always participate and many who never do (1, 5 and 4 vs. 13, 6 and 7 under tTTC, tGS and tRSD respectively). The lower average participation rate under tTTC reported in Result 3 may be driven by the important number of subjects who always opt out. Such behavior may be simply due to a conservative strategy by individuals with problems to understand the rules governing the allocation mechanism.

Table 7: Frequency of participation

	Frequency - tTTC							Frequency - tGS						Frequency - tRSD							
	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6
tenant																					
1	-	-	-	-	-	-	-	-	-	-	-	-	-	-	1	0	-	-	-	-	-
2	3	2	0	-	-	-	-	0	0	1	-	-	-	-	1	1	1	-	-	-	-
3	3	1	2	1	-	-	-	2	5	2	2	-	-	-	2	3	1	1	-	-	-
4	3	1	2	5	0	-	-	3	4	3	2	2	-	-	0	6	1	3	1	-	-
5	3	1	3	0	2	0	-	0	4	2	1	0	0	-	3	2	1	4	1	1	-
6	1	1	0	1	1	0	0	1	1	0	0	1	0	0	0	2	0	0	0	0	0

Further analysis of the individual participation decisions reveals some interesting patterns. As developed above, priority in the queue is an important explanatory variable. In Table 8, we classify subjects into those who: (i) always participate; (ii) participate only when first or second in the queue; (iii) participate only when first in the queue; (iv) never participate.²⁰ This classification explains the behavior of 44% to 64% of the subjects depending on the mechanism. Note that these types are all characterized by some level of rationality, as they realize that the likelihood of improving their allocation by opting in is increasing in their position in the queue. At the same time, it is hard to imagine why would anyone choose not to participate when they are not endowed with their preferred house and they are first in the queue.

Table 8: Classification of individual participation behavior

		tTTC	tGS	tRSD
(i)	Always IN	1	6*	4
(ii)	IN if 1st or 2nd	7*	3	6*
(iii)	IN only if 1st	2	1	2
(iv)	Always OUT	13	6	7
	Total	23 (64%)	16 (44%)	19 (53%)

* Includes an individual who did not participate when occupying her most preferred house

Result 5 (Individual Participation) *Only 9% of subjects always participate in the mechanism whereas 24% never do. We can classify 54% of subjects according to a monotone participation rule based exclusively on their ranking in the priority queue.*

4.3 Truthful Preference Revelation

Truthful preference revelation is a dominant strategy for agents participating in any of the three mechanisms under consideration. Table 9 presents for each mechanism the proportion of truthful preference revelation by session, tier, match, position in the queue, and house endowment. Even though participants submit rankings over all available houses, we only consider the *relevant rankings*. Indeed, if n agents choose to participate in a given

²⁰Some subjects never play in a certain position. So, for example, an agent who chooses IN when first, OUT when second and never plays in third position would be classified as a type (iii). Note that agents who chooses IN when first, OUT when third and who never play in second position could be classified as either (ii) or (iii). We chose to put them in (ii).

tier, they can only end up in a house ranked 1 to n , so we restrict attention to preference revelation over that set.²¹ The overall proportion of truthful revelation is high: 86.4% under tTTC, 87.6% under tGS and 94.7% under tRSD. T-tests of equality of proportions show that truthful revelation is higher under tRSD than under tTTC ($p = 0.023$) or tGS ($p = 0.045$) whereas no significant differences are found between tTTC and tGS ($p = 0.777$). When we compare truth-telling rates across mechanisms for subsets of the data, the most significant difference is that newcomers are more likely to tell the truth under tRSD than tTTC ($p = 0.71$) while participating existing tenants are more likely to tell the truth under tRSD than tGS ($p = 0.041$).

Table 9: Proportion of truthful preference revelation (relevant ranking)

		tTTC	tGS	tRSD
Session	1	38/45	40/49	45/45
	2	41/45	33/39	39/42
	3	29/35	40/41	40/44
Tier	1	21/26	24/27	29/31
	2	29/30	25/32	27/29
	3	27/33	29/31	34/35
	4	31/36	35/39	34/36
Match	1	14/18	19/21	19/19
	2	20/23	15/20	21/22
	3	18/21	18/19	20/20
	4	17/20	19/21	22/23
	5	20/21	19/22	19/23
	6	19/22	23/26	23/24
Priority	1	50/56	45/51	55/56
	2	46/54	52/56	54/59
	3	12/15	16/22	15/16
Endowment	newcomer	60/72	64/72	67/72
	tenant	48/53	49/57	57/59
Overall		0.86	0.88	0.95

Result 6 (Overall Truth-telling) *Truth-telling is very high in all three mechanisms: 86.4% for tTTC, 87.6% for tGS and 94.7% for tRSD. Truth-telling is higher under tRSD*

²¹The case where all rankings are considered is discussed in Appendix A. Naturally, as we enlarge the set of houses ranked, the likelihood of truthful revelation decreases. However, we show that the same qualitative properties hold under the full ranking specification.

than under *tTTC* or *tGS*.

Next, we look at the determinants of truthful preference revelation within mechanisms. Given the very high rates of truth-telling, differences in behavior when we partition the data are likely not to be significant or to be significant but small in magnitude. For example, we find that under *tGS* truth-telling rates are higher in positions 1 and 2 than in position 3 ($p = 0.101$ and $p = 0.017$ respectively). The position in the priority queue has no effect under *tTTC* or *tRSD* although, in the latter case, it is mainly because agents almost always reveal truthfully. Also, in general we do not find statistically significant differences under any mechanism when we look at truth-telling rates by matches or when we compare the behavior of newcomers and existing tenants. As in section 4.2, we also perform logit regressions where the dependent variable is a discrete choice variable that takes value 1 if the agent reveals his ranking truthfully and 0 otherwise. The regressions are not informative due to the little variation in the truth-telling rates (results omitted for brevity).

Result 7 (Determinants of Truth-telling) *Truth-telling rates are not significantly affected by the position in the queue, tier or match.*

Again, multiple matches allow us to perform some analysis at the individual level. Table 10 tabulates the frequency of truth-telling by subjects as a function of the number of times the subject submitted preferences over the available houses. Less than 10% misreport their preferences more than once (5, 3 and 1 subject, in *tTTC*, *tGS* and *tRSD* respectively) whereas 80% always report truthfully. It suggests that the observed aggregate differences in truth-telling across mechanisms (which are significant but small in magnitude) are driven by the behavior of only a few subjects and therefore they should not be over-emphasized. On the other hand, these differences may also reflect the greater difficulty to understand some mechanisms than others.

Our final analysis at the individual level is reported in Table 11 and consists in classifying heuristically the 40 instances of preference misrepresentation into categories. We find five categories, labeled (i) to (v), with 5 to 10 observations each. In (i), subjects rank a vacant house above their first true preference. In (ii), subjects switch their top two choices. In (iii), subjects list their own house first and then randomly assign the rest. In (iv), subjects follow some pattern in the rankings but there is not enough data to put them in a separate category (for example, a subject second in the priority queue submits the first two rankings truthfully and the third one randomly). In (v), we put the observations with no discernible pattern.²² One could argue that (i), (iii) and possibly (iv) follow some

²²These categories are similar to those found in [CS].

Table 10: Frequency of truth-telling (relevant ranking)

	Frequency - tTTC							Frequency - tGS							Frequency - tRSD							
	0	1	2	3	4	5	6	0	1	2	3	4	5	6	0	1	2	3	4	5	6	
submitted																						
1	0	4	-	-	-	-	-	0	1	-	-	-	-	-	1	4	-	-	-	-	-	-
2	1	2	1	-	-	-	-	0	0	7	-	-	-	-	0	0	2	-	-	-	-	-
3	0	2	1	5	-	-	-	0	0	2	6	-	-	-	0	0	0	9	-	-	-	-
4	0	0	0	1	6	-	-	0	1	0	1	8	-	-	0	1	0	1	7	-	-	
5	0	0	1	0	1	9	-	0	1	0	0	0	3	-	0	0	0	0	2	5	-	
6	0	0	0	1	0	0	0	1	0	0	0	0	0	4	0	0	0	0	0	0	4	

reasonable logic, and (v) corresponds to individuals who did not pay attention or did not understand the game. The most puzzling behavior is (ii) since we find no plausible reason for this type of misrepresentation.

Table 11: Classification of Misrepresentations

		tTTC	tGS	tRSD	Total
(i)	Vacant house	5	2	2	9
(ii)	Switch-top-two	4	3	1	8
(iii)	Random after own	2	4	2	8
(iv)	Unclassified	2	3	0	5
(v)	Random	4	4	2	10
Total		17	16	7	40

Result 8 (Individual Truth-telling) *9% of subjects misreport their preferences more than once and 80% always report them truthfully. Less than 4% of the observations follow an undiscernible pattern.*

5 Conclusion

In this paper we have studied the *tiered housing allocation problem with existing tenants*. In our setting, as is often the case in society, individuals are partitioned into tiers with

hierarchical privileges. We have proposed modified versions of three well-known mechanisms –Top Trading Cycles, Gale-Shapley and Random Serial Dictatorship– and analyzed their theoretical properties.

We have then evaluated the mechanisms in a laboratory experiment on three fronts and showed that performance of the three mechanisms is similar, with high efficiency, low participation rates and high rates of truthful revelation conditional on participation. We have reported some performance differences with the previous literature. Also, of the three novelties introduced in our experiment –tiered structure, multiple matches and known priority queues– only the last one has a significant impact on choices. According to our analysis at the individual level, only 8% of subjects realize that participating is a dominant strategy under tTTC and tGS whereas 80% realize the optimality of truthful preference revelation conditional on participation under all mechanisms. More generally, the majority of subjects follow discernible individual strategies. For instance, participation only as a function of the position in the priority queue accounts for the behavior of half of our subjects. Finally, the substantial behavioral differences between the existing one-sided matching laboratory experiments suggests that some minor and apparently innocuous features of the protocol and experimental conditions may have as much impact in the outcome as the type of allocation mechanism employed.

References

- [1] Abdulkadiroğlu, A. and T. Sönmez (1999) “House Allocation with Existing Tenants.” *Journal of Economic Theory*, 88, 233-260.
- [2] Abdulkadiroğlu, A. and T. Sönmez (2003) “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93, 729-747.
- [3] Baccara, M., Imrohoroglu, A., Wilson, A., and L. Yariv (2009) “A Field Study on Matching with Network Externalities.” Working Paper.
- [4] Calsamiglia, C., Haeringer, G. and F. Klijn (2009) “Constrained School Choice: An Experimental Study.” Working paper.
- [5] Chen, Y. and T. Sönmez (2002) “Improving Efficiency of On-Campus Housing: An Experimental Study.” *American Economic Review*, 92, 1669-1686.
- [6] Chen, Y. and T. Sönmez (2004) “An Experimental Study of House Allocation Mechanisms.” *Economics Letters*, 83, 137-140.
- [7] Chen, Y. and T. Sönmez (2006) “School Choice: An Experimental Study.” *Journal of Economic Theory*, 127, 202-231.
- [8] Featherstone, C. and M. Niederle (2009) “School Choice Mechanisms under Incomplete Information: An Experimental Investigation.” Working Paper.
- [9] Gale, D. and L. Shapley (1962) “College Admissions and the Stability of Marriage.” *American Mathematical Monthly*, 69, 9-15.
- [10] Guillen, P. and O. Kesten (2010) “Matching Markets with Mixed Ownership: The Case for A Real-life Assignment Mechanism.” Working paper.
- [11] Haruvy, E., Roth, A. and U. Ünver (2006) “The dynamics of law clerk matching: An experimental and computational investigation of proposals for reform of the market.” *Journal of Economic Dynamics and Control*, 30(3), 457-486.
- [12] Kagel, J. and A. Roth (2000) “The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by A Natural Experiments.” *Quarterly Journal of Economics*, 115(1), 201-235.
- [13] Nalbantian, H.R. and A. Schotter (1995) “Matching and Efficiency in the Baseball Free-Agent System: An Experimental Examination.” *Journal of Labor Economics*, 13(1), 1-31.

- [14] Olson, M. and D. Porter (1994) "An experimental examination into the design of decentralized methods to solve the assignment problem with and without money." *Economic Theory*, 4(1), 11-40.
- [15] Pais, J. and A. Pintér (2008) "School choice and information: An experimental study on matching mechanisms." *Games and Economic Behavior*, 64, 303-328.
- [16] Pápai, S. (2000) "Strategyproof Assignment by Hierarchical Exchange." *Econometrica*, 68(6), 1403-1433.
- [17] Pycia, M. and U. Ünver (2007) "A Theory of House Allocation and Exchange Mechanisms." Working paper.
- [18] Roth, A.E., Sönmez, T. and U. Ünver (2004) "Kidney Exchange." *Quarterly Journal of Economics*, 119(2), 457-488.
- [19] Roth, A.E. and M. Sotomayor (1990) *Two-Sided Matching: A Study on Game - Theoretic Modelling*. Cambridge University Press.
- [20] Shapley, L. and H. Scarf (1974) "On Cores and Indivisibility." *Journal of Mathematical Economics*, 1, 23-28.
- [21] Sönmez, T. and U. Ünver (2011) "Matching, Allocation and Exchange of Discrete Resources", in J. Benhabib, A. Bisin, and M. Jackson (eds.) *Handbook of Social Economics* Vol. 1A, Elsevier.
- [22] Title, L.R. (1998) "Platonic Mechanism Design." Working Paper.
- [23] Ünver, U. (2005) "On the survival of some unstable two-sided matching mechanisms." *International Journal of Game Theory*, 33(2), 239-254.

Appendix A: preference revelation (full ranking)

In this section we check the robustness of our preference revelation analysis by considering all the rankings submitted by subjects. Naturally, the overall proportion of truthful preference revelation decreases. It becomes 81% under tTTC, 78% under tGS and 84% under tRSD, which for tTTC and tRSD are still above those found in the previous literature. Table 12 shows the analogue of Table 9 when all rankings are considered. We find that most of the drop in truth-telling happens in tiers 1 and 2. This is not surprising: in higher tiers subjects have to submit rankings over a larger set of available houses, so they are more liable to make “mistakes” after the first few (relevant) rankings.

Table 12: Proportion of truthful preference revelation (full ranking)

		tTTC	tGS	tRSD
Session	1	36/45	38/49	40/45
	2	37/45	29/39	30/42
	3	28/35	34/41	38/44
Tier	1	16/26	17/27	22/31
	2	28/30	21/32	18/29
	3	26/33	28/31	34/35
	4	31/36	35/39	34/36
Match	1	14/18	17/21	17/19
	2	19/23	14/20	17/22
	3	17/21	16/19	17/20
	4	16/20	18/21	21/23
	5	18/21	17/22	17/23
	6	17/22	19/26	19/24
Priority	1	46/56	40/51	44/56
	2	44/54	46/56	51/59
	3	11/15	15/22	13/16
Endowment	newcomer	57/72	55/72	58/72
	tenant	44/53	46/57	50/59
Overall		0.81	0.78	0.84

Overall, there are 385 observations where individuals submit preferences. In 345 cases, subjects reveal truthfully the relevant rankings (Table 9) and in 310 cases they reveal truthfully the entire ranking (Table 12). Therefore, there are 35 observations where truthful rankings are submitted over the relevant ranking but not over the entire one. In Table 13 we take a closer look at these observations. As we can see, in almost half of these observations subjects make a mistake either in the last two rankings or in the ranking immediately after the relevant ones.

Table 13: Misrepresentations after relevant ranking

	tTTC	tGS	tRSD	Total
Mistake in last two ranks only	0	4	2	6
Mistake immediately after relevant ranks	3	3	5	11
Other mistakes	4	5	9	18
Total	7	12	16	35

Appendix B: sample of instructions (tTTC mechanism)

This is an experiment in the economics of decision making and you will be paid for your participation in cash at the end of the experiment. The entire experiment will take place through computer terminals, and all interaction between participants will take place through the computers. You will remain anonymous to me and to all the other participants during the entire experiment; the only person who will know your identity is the Lab Manager who is responsible for paying you at the end. It is important that you not talk or in any way try to communicate with other participants during the experiment. Remember that you are not being deceived and you will not be deceived: everything I tell you is true.

In this experiment we are going to simulate a house allocation process. The procedure and payment rules will be described in detail below. We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. You must take a quiz after the instruction period, so it is important that you listen carefully. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

Different participants may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. At the end of the session, you will be paid the sum of what you have earned in the experiment and a show-up fee of \$5. Everyone will be paid in private and you are under no obligation to tell others how much you earned. The experiment consists of 6 matches. The procedure in each match is exactly the same and is as follows:

The Procedure is as follows:

- There are 12 participants divided into 4 tiers. Your participation ID and TIER is mentioned on your screen. [SLIDE #2]
- In each tier there are 3 participants. 2 of them are EXISTING tenants, that is, they currently occupy a house. 1 of them is a NEWCOMER who does not have a house yet. In all there are 8 existing tenants and 4 newcomers. Your ROLE of existing tenant or newcomer is also mentioned on your screen. [SLIDE #2]
- There are 12 houses labeled A-L to allocate. [SLIDE #2] Each house must be allocated to one and only one participant.

- Your payoff for the match, denominated in dollars, depends on the house you hold at the end of the match and it is given in the payoff table like this [SLIDE#2]. For example, if you hold house K at the end of the match, then your payoff is \$17.
- Should you be the current tenant of a house, then this fact is also indicated on your computer screen [SLIDE #2]. Note that different participants might have different payoff tables and these payoffs are privately known.
- In the experiment, the allocation of houses is done sequentially, by tiers. We start with tier 1 while participants in lower tiers wait. When the allocation is done for tier 1 we proceed to tier 2 while participants in lower tiers wait and so on.
- The match ends when the allocation process for all the tiers is over and we move on to the next match. For the next match the computer randomly reassigns the tiers and the roles of existing tenants and newcomers. The new assignments do not depend in any way on the past decisions of any participant including you and are done completely randomly by the computer.
- The second match then follows the same rules as the first match. This continues for 6 matches after which the experiment ends.
- At the end of the experiment the computer randomly selects with equal probability one of the 6 matches and your payoff in the experiment is equal to your payoff in the match selected by the computer.

Classification of Houses:

The houses are classified into the following categories indicated by their STATUS [SLIDE #2]. There are 2 main categories: Not Available and Available.

- NOT AVAILABLE houses are houses which have already been assigned and are no longer available for allocation. These are indicated by the color BLUE and labeled as NA.
The AVAILABLE houses can be further classified into:
- If the house is occupied by you, then it is indicated by the color LIGHT GREEN and labeled (O).
- If it is occupied by someone else in your tier then it is labeled as (OS) and colored RED if the tenant is ranked higher than you in the priority queue (we will explain in a minute what this is) and DARK GREEN if the tenant is ranked lower than you in the priority queue.
- If the house is occupied by someone in a tier lower than yours then it is indicated by the color ORANGE and labeled (L).
- If the house is vacant, that is, not occupied by anyone, it is indicated by the color WHITE and labeled (V).

House Allocation is as follows:

Tier 1:

The house allocation process starts with tier 1 while the other tiers wait. Within tier1 the house allocation takes place in the following way.

- All participants in the tier are lined in a pre-determined priority queue. Your position in the queue is indicated on your screen. [SLIDE #2] Note that your position in the queue does not depend on any of your or anyone else's past decisions.
- Existing tenants first, simultaneously, choose between taking part in the allocation process by choosing IN and not taking part by choosing OUT. [SLIDE #2]
 - If you are an existing tenant and choose OUT, you will keep your current house and the allocation process is over for you.
 - If you are an existing tenant and choose IN, you will then need to rank the available houses. [SLIDE #3]
- If you are a newcomer you cannot choose OUT. You need to rank the available houses.
- Note that the participating existing tenants and the newcomers simultaneously submit their rankings and they need to rank all the available houses. No two houses should be given the same rank.
- Once the participating existing tenants and the newcomers have submitted their ranking of the available houses, the house allocation takes place in the following way:
- We proceed from the top of the priority queue. Based on her chosen ranking of the houses, for the first participant in the queue, we look at the status of her top ranked house from among the remaining houses.
- If the Status of the house is L or V, i.e., if the house is occupied by someone in a lower tier or not occupied by anyone, then the participant at the top of the queue is assigned to it. Note that during the allocation process houses occupied by lower tiers are treated in the same manner as those that are vacant.
- Note that an existing tenant vacates her current house, once she is assigned another house.
- If the requested house is not V or L, it means the requested house is the current house of an existing tenant in the tier. In this case, the existing tenant is moved to the top of the priority queue, directly in front of the requester. This way the existing tenant is always guaranteed a house which is at least as good as the house she is living in, based on her chosen ranking of the houses. The process continues afterwards with the modified queue.
- If a cycle of requests are formed (e.g., I want John's house, John wants your house and you want my house), all members of the cycle are given what they want, and their new houses are removed from the system.
- The process continues until all participants in tier1 are assigned a house.

Tier 2:

- All participants are lined in a pre-determined priority queue. Your position in the queue is indicated on your screen. Note that your position in the queue does not depend on any of your or anyone else's past decisions.

- The houses assigned to members of tier 1 are not available (NA). In tier 2, if an existing tenant's house has not been allocated to someone in tier 1, then the agent continues to occupy her house. If however, the agent's house has been taken by a member of a higher tier, she is compensated with a randomly selected house which was either previously occupied by a member of tier 1 or not occupied by anyone (V).
- Existing tenants first, simultaneously, choose between taking part in the allocation process by choosing IN and not taking part by choosing OUT.
 - If you are an existing tenant and choose OUT, you will keep your current house and the allocation process is over for you.
 - If you are an existing tenant and choose IN, you will then need to rank the available houses.
- If you are a newcomer you cannot choose OUT. You need to rank the available houses.

The rest of the steps are exactly as those for tier 1. When the allocation for TIER 2 is over, we move on to TIER 3. The steps for TIER 3 are exactly the same as those for TIER 2. When the allocation for TIER 3 is over, we move on to TIER 4, where once again, the steps are exactly the same.

EXAMPLE:

We will now go through a simple example to illustrate how the allocation method works. [Slide #4] Suppose that there are six participants in two tiers. Participants 1, 2 and 3 belong to tier 1 while 4, 5 and 6 belong to tier 2. In tier 1, participants 1 and 2 are existing tenants occupying houses W and X respectively, while 3 is a newcomer. In tier 2, participants 4 and 5 are existing tenants occupying houses Y and Z respectively, while 6 is a newcomer. In addition, houses U and V are vacant.

We start with tier 1. Suppose the pre-determined priority queue is: 3-1-2. Suppose player 1 chooses OUT and player 2 chooses IN. Player 1 is allocated her current house W and the allocation process is over for her. Then suppose that players 2 who chose IN and player 3 who is automatically in because she is a newcomer, given their payoffs, enter the following ranking of houses: [SLIDE #5]

	Participant 2	Participant 3
Rank 1	V	Z
Rank 2	Y	U
Rank 3	U	Y
Rank 4	X	V
Rank 5	Z	X

Then the allocation for Tier 1 takes place in the following manner:

	Priority Queue	Remaining Houses	Top choice among remaining houses and its status	Actions taken at the the end of the step
Step 1	3 – 2	X, Y, Z, U, V	Z, status-L	3 gets Z
Step 2	2	X, Y, U, V	V, status-V	2 gets V

Step1: We start with participant 3. His top choice is house Z which has status L. Participant 3 is assigned house Z.

Step 2: Only participant 2 remains. His top choice among the remaining houses is V which is vacant. So participant 2 is assigned house V.

[SLIDE #6] Now we move on to tier 2. Suppose the priority queue is 6-5-4. Since player 5 has lost her house to tier 1 she is compensated by a house randomly chosen from the set of houses that were previously occupied by tier 1 or vacant. Suppose she is compensated by house X. Now suppose the players 4, 5 choose IN and then players 4, 5 and 6 submit the following ranking of houses.

	Participant 4	Participant 5	Participant 6
Rank 1	X	Y	X
Rank 2	Y	U	Y
Rank 3	U	X	U

Then the allocation for Tier 2 takes place in the following manner:

	Priority Queue	Remaining Houses	Top choice among remaining houses and its status	Actions taken at the end of the step
Step 1	6 – 5 – 4	X, Y, U	X , occupied by 5	6-5-4 becomes 5-6-4
Step 2	5 – 6 – 4	X, Y, U	Y , occupied by 4	5-6-4 becomes 4-5-6
Step 3	4 – 5 – 6	X, Y, U	X , occupied by 5	4 gets X , 5 gets Y

Step 1: The priority queue is 6-5-4. Participant 6 has ranked house X as his top choice which is currently occupied by participant 5. Participant 5 is moved to the top of the queue.

Step 2: The modified priority queue is 5-6-4. Participant 5 has ranked house Y as his top choice which is currently occupied by participant 4. Participant 4 is moved to the top of the queue.

Step 3: The modified priority queue is 4-5-6. Participant 4 has ranked house X as his top choice which is currently occupied by participant 5. Now a cycle is created where Participant 4 wants the house of participant 5 and participant 5 wants the house of participant 4. So participant 4 is given house X and participant 5 is given house Y.

Step 4: Now only participant 6 is left. He gets house U.

The following slide summarizes the rules of the experiment: [Read summary slides #7 and #8]

***** PRACTICE SESSION *****

We will now begin the Practice session and go through a practice match to familiarize you with the computer interface and the procedures. During the practice match, please do not hit any keys until you are asked to do so, and when you enter information, please do exactly as asked. Remember, you are not paid for this practice match. At the end of the practice match you will have to answer some review questions. Are there any questions before we begin?

[AUTHENTICATE CLIENTS]

You have just received your first match. Notice your Role and tier.

- The existing tenant occupying house A in Tier 1 will see a screen like this [SLIDE #9]. Notice that you are asked to choose between taking part in the allocation process by choosing IN and not taking part by choosing OUT. In this the Existing tenant is occupying house A which is colored LIGHT GREEN. House B is colored DARK GREEN as it is occupied by someone else in the tier but with a lower position in the priority queue. The houses occupied by participants in tiers 2, 3 and 4 are colored Orange while houses C,F,K and L are colored white as they are not occupied by anyone.
- Newcomer in Tier 1 will see a screen like this [SLIDE #10]. Notice that you are asked to wait for your turn while the existing tenants in the tier choose between taking part in the allocation process or not. Notice that houses A and B are colored RED as they are occupied by participants with higher position in the priority queue. The houses occupied by participants in lower tiers are colored Orange those that are not occupied by anyone are colored white.
- Those not in tier 1 will see a screen like this [SLIDE #11]. Notice that you are asked to wait for your turn. If you are an existing tenant this is indicated on your screen and the house you occupy is colored LIGHT GREEN.
- Existing Tenants in Tier 1 please click "OUT". Notice that you have been allocated the house that you were occupying. For example, the Existing tenant who was occupying house A will see a screen like this [SLIDE #12].
- Newcomer in Tier 1 will see a screen like this [SLIDE #13]. Please give rank 1 to house C, rank 2 to house D, rank 3 to house E... and so on till rank 10 to house L.
- The Newcomer is allocated house C.

Now we have moved to tier 2

- The existing tenant occupying house E in Tier 2 will see a screen like this [SLIDE #14]. Notice that now houses A, B and C have become Not Available as they have already been assigned to someone in tier 1. In this the Existing tenant is occupying house E which is colored LIGHT GREEN. House D is colored RED as it is occupied by someone else in the tier but with a higher position in the priority queue. The houses occupied by participants in tiers 3 and 4 are colored Orange while houses not occupied by anyone are colored white.
- Notice that all participants in tiers 3-4 are asked to wait while we finish the allocation of tier 2. [SLIDE #15] You will also see that houses A, B and C have become Not Available as they have already been assigned to someone in a higher tier. If you are an existing tenant this is indicated on your screen and the house you occupy is colored LIGHT GREEN.
- Existing Tenants in Tier 2 please click IN. [SLIDE #16] Now Existing Tenants and Newcomer in Tier 2 please give rank 1 to house D, rank 2 to house E, rank 3 to house F,... and so on till rank 9 to house L.
- The participant first in the priority queue is allocated house D, the participant second in the priority queue is allocate house E and the participant third in the priority queue is allocated house F.

Now we have moved to tier 3

- The existing tenant occupying house G in Tier 3 will see a screen like this [SLIDE #17]. Notice that now houses D, E and F have also become Not Available as they have already been assigned to someone in tier 2. As the Existing tenant is occupying house G it is colored LIGHT GREEN. House H is colored DARK GREEN as it is occupied by someone else in the tier 3 but with a lower position in the priority queue. The houses occupied by participants in tier 4 are colored Orange while houses not occupied by anyone are colored white.
- Notice that all participants in tiers 4 are asked to wait while we finish the allocation of tier 3. You will also see that houses D, E and F have become Not Available as they have already been assigned to someone in a higher tier. If you are an existing tenant this is indicated on your screen and the house you occupy is colored LIGHT GREEN.
- Existing Tenants in Tier 3 please click OUT. Notice that you have been allocated the house that you were occupying. For example, the Existing tenant was occupying house G will see a screen like this [SLIDE #18].
- Newcomer in Tier 3 will see a screen like this [SLIDE #19]. Newcomer in Tier 3 please give rank 1 to house I, rank 2 to house J, rank 3 to house K and rank 4 to house L.
- The Newcomer is allocated house I.

Now we have moved to tier 4

- All the participants in tier 4 will see that houses G,H and I have also become Not Available.
- The existing tenant who was earlier occupying house I has lost his house but has now been compensated by another house, chosen randomly from the set of vacant houses. For example, if he is compensated with house L then he will see a screen like this [SLIDE #20].
- Existing Tenants in Tier 4 please click IN. [SLIDE #21] Now Existing Tenants and Newcomer in Tier 4 please give rank 1 to house J, rank 2 to house K, rank 3 to house L.
- The participant first in the priority queue is allocated house J, the participant second in the priority queue is allocate house K and the participant third in the priority queue is allocated house L.

***** END OF PRACTICE SESSION *****

The practice match is over. Please complete the quiz. It has 6 questions. If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you.

[START QUIZ]

[WAIT for everyone to finish the Quiz]

Are there any questions before we begin with the paid session? We will now begin with the 6 paid matches. If there are any problems or questions from this point on, raise your hand and an experimenter will come and assist you.

[START MATCH 1]

[After MATCH 6 read:]

This was the last match of the experiment. Your payoff is displayed on your screen. Please record this payoff in your record sheet. [CLICK ON WRITE OUTPUT]

Your Total Payoff is this amount plus the show-up fee of \$5. We will pay each of you in private in the next room in the order of your Subject ID number. Remember you are under no obligation to reveal your earnings to the other participants.