

Trading Cycles for School Choice

Marek Pycia*

UCLA

M. Utku Ünver[†]

Boston College

October 18, 2011

Abstract

In this note we study the allocation and exchange of discrete resources in environments in which monetary transfers are not allowed. We allow each discrete resource to be represented by several copies, extend onto this environment the trading cycles mechanisms of Pycia and Ünver [2009], and show that the extended mechanisms are Pareto efficient and strategy-proof. In particular, we construct the counterpart of Pápai [2000] hierarchical exchange mechanisms for environments with copies.

Keywords: Mechanism design, group strategy-proofness, Pareto efficiency, matching, house allocation, house exchange, outside options.

JEL classification: C78, D78

*UCLA, Department of Economics, 8283 Bunche Hall, Los Angeles, CA 90095.

[†]Boston College, Department of Economics, 140 Commonwealth Ave., Chestnut Hill, MA 02467.

1 Introduction

The theory of mechanism design has informed the design of many markets and other institutions. Auction mechanisms have been developed to sell treasury bills, electricity, natural gas, radio spectra, timber, and foreclosed homes. Other mechanisms have been developed to allocate resources in environments in which transfers are not used or are prohibited: to allocate and exchange transplant organs – such as kidneys – through newly established regional programs such as the Alliance for Paired Donation (centered in Toledo, Ohio) and the New England Program for Kidney Exchange (centered in Newton, Massachusetts), and the new U.S. national program (managed by the United Network for Organ Sharing) [cf. Roth, Sönmez, and Ünver, 2004], to allocate school seats in New York City, Chicago, San Francisco, and Boston [cf. Abdulkadiroğlu and Sönmez, 2003], and to allocate dormitory rooms to students at US colleges [cf. Abdulkadiroğlu and Sönmez, 1999].

A primary concern in market design is the strategic behavior of market participants and its impact on resulting allocations. Strategic behavior becomes straightforward if the market mechanism is dominant-strategy incentive compatible: no agent can manipulate the system to his or her benefit. Non-manipulability is not the only benefit of using dominant-strategy incentive-compatible mechanisms. Dominant-strategy incentive-compatibility imposes minimal costs of searching for and processing strategic information, and do not discriminate among agents based on their access to information and ability to strategize [cf. Vickrey, 1961, Dasgupta et al., 1979, Pathak and Sönmez, 2008].

Pycia and Ünver [2009] constructed the full class of group strategy-proof and Pareto efficient mechanisms for environments without transfers in which agents have strict preferences over objects. They also show that some trading-cycles mechanisms improve upon all previously studied mechanisms, and allows one to implement efficient allocations that cannot be implemented by previously studied mechanisms.

In the present note we extend their construction to environments in which each object can be represented by several copies, and show that the resultant class is strategy-proof and efficient. In particular, ours is the first paper to extend Pápai [2000] hiererachical exchange mechanisms to environments with copies.

Let us highlight some common features of the above market design problems with no transfers. There is a finite group of agents, each of whom would like to consume a single indivisible object to which we will refer as a “house,” or “seat” using the terminology coined by Shapley and Scarf [1974] and Abdulkadiroğlu and Sönmez, 2003. Each house has a type (for instance each seat is in a particular school). Agents have strict preferences over types and are indifferent among objects of the same type. Some of the houses might be agents’

common endowment, while others belong to agents' private endowments. The outcome of the problem is a matching of agents and houses. Since we provide a unified treatment of both house allocation (from social endowment) and house exchange (among agents with private endowments), we refer to our environment as house allocation and exchange. By the revelation principle, we may restrict our attention to direct revelation mechanisms; that is, agents reveal their preferences over houses, and the mechanism matches each agent with a house (or the agent's outside option).

Both Pycia and Ünver [2009] and Pápai [2000] classes of mechanisms build on Gale's Top Trading Cycles (see Shapley and Scarf, 1974). Gale's mechanisms have been extended to environments with copies by Abdulkadiroğlu and Sönmez, 2003.¹

2 Model

2.1 Environment

Let I be a set of **agents** and H be a set of **schools (or houses,)**. We use letters i, j, k to refer to agents and h, g, e to refer to school. Each school h is represented by q_h copies (**seats**), and each agent demands one seat.

Each agent i has a **strict preference relation** over H , denoted by \succ_i .² Let \mathbf{P}_i be the set of strict preference relations for agent i , and let \mathbf{P}_J denote the Cartesian product $\times_{i \in J} \mathbf{P}_i$ for any $J \subseteq I$. Any profile from $\succ = (\succ_i)_{i \in I}$ from $\mathbf{P} \equiv \mathbf{P}_I$ is called a **preference profile**. For all $\succ \in \mathbf{P}$ and all $J \subseteq I$, let $\succ_J = (\succ_i)_{i \in J} \in \mathbf{P}_J$ be the restriction of \succ to J .

To simplify the exposition, we make two initial assumptions. Both of these assumptions are fully relaxed in subsequent sections. First, we initially restrict attention to allocation problems. A **house allocation problem** is the triple $\langle I, H, \succ \rangle$ [cf. Hylland and Zeckhauser, 1979]. Throughout the paper, we fix I and H , and thus, a problem is identified with its preference profile. In Section 5, we generalize the setting and the results to house allocation and exchange by allowing agents to have initial rights over houses. The results on allocation and exchange turn out to be straightforward corollaries of the results on (pure) allocation. Second, we initially follow the tradition adopted by many papers in the literature [cf. Svensson, 1999, Bogomolnaia and Moulin, 2001] and assume that $|H| \geq |I|$ so that each agent is allocated a house. This assumption is satisfied in settings in which each agent is always allocated a house (there are no outside options), as well as in settings in which agents' outside options are tradeable, effectively being indistinguishable from houses. In Section 6, we

¹Pycia and Ünver [2009] provide an overview of other related literature.

²By \succeq_i we denote the induced weak preference relation; that is, for any $g, h \in H$, $g \succeq_i h \iff g = h$ or $g \succ_i h$.

allow for non-tradeable outside options and show that analogues of our results remain true irrespective of whether $|H| \geq |I|$ or $|H| < |I|$.

An outcome of a house allocation problem is a matching. To define a matching, let us start with a more general concept that we will use frequently. A **submatching** is an allocation of a subset of houses to a subset of agents, such that no two different agents get the same house. Formally, a submatching is a function $\sigma : J \rightarrow H$ where $J \subseteq I$ and $|\sigma^{-1}(h)| \leq q_h$. Using the standard function notation, we denote by $\sigma(i)$ the assignment of agent $i \in J$ under σ , and by $\sigma^{-1}(h)$ the set of agents that got house $h \in \sigma(J)$ under σ . Let \mathcal{S} be the set of submatchings. For each $\sigma \in \mathcal{S}$, let I_σ denote the set of agents matched by σ . We write \overline{I}_σ for $I - I_\sigma$, and \overline{H}_σ for the set of houses h such that $q_h - \sigma^{-1}(h) > 0$. For all $h \in H$, let $\mathcal{S}_{-h} \subset \mathcal{S}$ be the set of submatchings $\sigma \in \mathcal{S}$ such that $h \in \overline{H}_\sigma$, i.e., the set of submatchings at which house h has unallocated copies. In virtue of the set-theoretic interpretation of functions, submatchings are sets of agent-house pairs, and are ordered by inclusion. A **matching** is a maximal submatching; that is, $\mu \in \mathcal{S}$ is a matching if $I_\mu = I$. Let $\mathcal{M} \subset \mathcal{S}$ be the set of matchings. We also write $\overline{\mathcal{M}}$ for $\mathcal{S} - \mathcal{M}$.

A **(direct) mechanism** is a mapping $\varphi : \mathbf{P} \rightarrow \mathcal{M}$ that assigns a matching for each preference profile (or, equivalently, allocation problem).

2.2 Strategy-Proofness and Pareto Efficiency

A mechanism is group strategy-proof if there is no group of agents that can misstate their preferences in a way such that each one in the group gets a weakly better house, and at least one agent in the group gets a strictly better house. Formally, a mechanism φ is **group strategy-proof** if for all $\succ \in \mathbf{P}$, there exists no $J \subseteq I$ and $\succ'_J \in \mathbf{P}_J$ such that

$$\varphi[\succ'_J, \succ_{-J}](i) \succeq_i \varphi[\succ](i) \text{ for all } i \in J,$$

and

$$\varphi[\succ'_J, \succ_{-J}](j) \succ_j \varphi[\succ](j) \text{ for at least one } j \in J.$$

A mechanism φ is **(individually) strategy-proof** if for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ'_i \in \mathbf{P}_i$ such that

$$\varphi[\succ'_i, \succ_{-i}](i) \succ_i \varphi[\succ](i).$$

A matching is Pareto efficient if no other matching would make everybody weakly better off, and at least one agent strictly better off. That is, a matching $\mu \in \mathcal{M}$ is Pareto efficient if there exists no matching $\nu \in \mathcal{M}$ such that for all $i \in I$, $\nu(i) \succeq_i \mu(i)$, and for some $i \in I$, $\nu(i) \succ_i \mu(i)$. A mechanism is **Pareto efficient** if it finds a Pareto-efficient matching for

every problem.

3 Trading-Cycles Mechanism

We turn now to our extension of Pycia and Ünver [2009] trading cycles (TC). TC is a recursive algorithm that matches agents and houses in exchange cycles over a sequence of rounds. TC is more flexible, however, as it allows two types of intra-round control rights over houses that agents bring to the exchange cycles: ownership and brokerage.

In our description of the TTC class, each TTC mechanism was determined by a consistent ownership structure. Similarly, each TC mechanism is determined by a consistent structure of control rights.

Definition 1. A **structure of control rights** is a collection of mappings

$$\{(c_\sigma, b_\sigma) : \overline{H}_\sigma \rightarrow \overline{I}_\sigma \times \{\text{ownership, brokerage}\}\}_{\sigma \in \overline{\mathcal{M}}}.$$

The functions c_σ of the control rights structure tell us which unmatched agent controls any particular unmatched house at submatching σ . Agent i **controls** house $h \in \overline{H}_\sigma$ at submatching σ when $c_\sigma(h) = i$. The type of control is determined by functions b_σ . We say that the agent $c_\sigma(h)$ **owns** h at σ if $b_\sigma(h) = \text{ownership}$, and that the agent $c_\sigma(h)$ **brokers** h at σ if $b_\sigma(h) = \text{brokerage}$. In the former case we call the agent an **owner** and the controlled house an **owned house**. In the latter case we use the terms **broker** and **brokered house**. Notice that each controlled (owned or brokered) house is unmatched at σ , and any unmatched house is controlled by some uniquely determined unmatched agent.

The consistency requirement on TC control rights structures consists of three constraints on brokerage at any given submatching (the *within-round* requirements) and three constraints on how the control rights are related across different submatchings (the *across-rounds* requirements).

Within-round Requirements. Consider any $\sigma \in \overline{\mathcal{M}}$.

(R1) There is at most one brokered house at σ .

(R2) If i is the only unmatched agent at σ then i owns all unmatched houses at σ .

(R3) If agent i brokers a house at σ , then i does not own any houses at σ . Furthermore, if h is brokered at σ then $q_h - |\sigma^{-1}(h)| = 1$.

The conditions allow for different houses to be brokered at different submatchings, even though there is at most one brokered house at any given submatching.

Requirements R1-R2 are what we need for the TC algorithm to be well defined (R3 is necessary for Pareto efficiency and individual strategy-proofness; see Appendix ??). With these requirements in place, we are ready to describe the TC algorithm, postponing the introduction of the remaining consistency requirements until the next section.

The TC algorithm. The algorithm consists of a finite sequence of rounds $r = 1, 2, \dots$. In each round some agents are matched with houses. By σ^{r-1} we denote the submatching of agents and houses matched before round r . Before the first round the submatching is empty, that is, $\sigma^0 = \emptyset$. If $\sigma^{r-1} \in \mathcal{M}$, that is, when every agent is matched with a house, the algorithm terminates and gives matching σ^{r-1} as its outcome. If $\sigma^{r-1} \in \overline{\mathcal{M}}$, then the algorithm proceeds with the following three steps of *round* r :

Step 1. Pointing. Each house $h \in \overline{H_{\sigma^{r-1}}}$ points to the agent who controls it at σ^{r-1} . If there exists a broker at σ^{r-1} , then he points to his most preferred house among the ones owned at σ^{r-1} . Every other agent $i \in \overline{I_{\sigma^{r-1}}}$ points to his most preferred house in $\overline{H_{\sigma^{r-1}}}$.

Step 2. Trading cycles. There exists $n \in \{1, 2, \dots\}$ and an exchange cycle

$$h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots h^n \rightarrow i^n \rightarrow h^1$$

in which agents $i^\ell \in \overline{I_{\sigma^{r-1}}}$ point to houses $h^{\ell+1} \in \overline{H_{\sigma^{r-1}}}$ and houses h^ℓ points to agents i^ℓ (here $\ell = 1, \dots, n$ and superscripts are added modulo n);

Step 3. Matching. Each agent in each trading cycle is matched with the house he is pointing to; σ^r is defined as the union of σ^{r-1} and the set of newly matched agent-house pairs.

The algorithm terminates when all agents are matched or when no unmatched house copies remain.

Looking back at the example of the previous section, we see that it was TC and that agent i_1 brokered house h_1 while other agents owned houses. We may now also see that requirements R1 and R2 are needed to ensure that in Step 1 there always is an owned house

for the broker to point to. The difference between TTC and TC is encapsulated in Step 1; the other steps are standard and were already present in Gale’s TTC idea [Shapley and Scarf, 1974]. The existence of the trading cycle follows from there being a finite number of nodes (agents and houses), each pointing at another. The matching of Step 3 is well defined, as (i) each agent points to exactly one house, and (ii) each matched house is allocated to exactly one agent (no two different agents pointing to the same house h can belong to trading cycles because there is a unique pointing path that starts with house h). Finally, since we match at least one agent-house pair in every round, and since there are finitely many agents and houses, the algorithm stops after finitely many rounds.

Our algorithm builds on Gale’s top-trading-cycles idea, but allows more general trading cycles than top cycles. In TC, brokers do not necessarily point to their top-choice houses. In contrast, all previous developments of Gale’s idea, such as the top trading cycles with newcomers [Abdulkadiroğlu and Sönmez, 1999], hierarchical exchange [Pápai, 2000], top trading cycles for school choice [Abdulkadiroğlu and Sönmez, 2003], and top trading cycles and chains [Roth, Sönmez, and Ünver, 2004], allowed only top trading cycles and had all agents point to their top choice among unmatched houses. All these previous developments may be viewed as using a subclass of TC in which all control rights are ownership rights and there are no brokers.³

The terminology of owners and brokers is motivated by a trading analogy. In each round of the algorithm, an owner can either be matched with the house he controls or with another house obtained from an exchange. A broker cannot be matched with the house he controls; the broker can only be matched with a house obtained from an exchange with other agents. At any submatching (but not globally throughout the algorithm), we can think of the broker of house h as representing a latent agent who owns h but prefers any other house over it. The analogy is, of course, imperfect and, ultimately, our choice of terminology is arbitrary.

4 Results

Introduced in the previous section, the TC algorithm with a control rights structure satisfying R1-R3 provides a Pareto-efficient mechanism that maps profiles from \mathbf{P} to matchings in \mathcal{M} . The recursive argument for the efficiency of the non-TTC mechanism from Section ?? applies.

Proposition 1. *The TC algorithm produces a Pareto-efficient matching for all control rights structures that satisfy R1-R3.*

³In particular, TC can easily handle private endowments, as explained in Section 5.

We are about to see that the TC-induced mapping is group strategy-proof if the control rights structure also satisfies the following across-round consistency requirements.

Across-round Requirements. Consider any submatchings σ, σ' such that $|\sigma'| = |\sigma| + 1$ and $\sigma \subset \sigma' \in \overline{\mathcal{M}}$, and any agent $i \in \overline{I_{\sigma'}}$ and any house $h \in \overline{H_{\sigma}}$:

(R4) If i owns house h at σ then i owns h at σ' .

(R5) Assume that at least two agents from $\overline{I_{\sigma'}}$ own houses at σ . If i brokers house h at σ then i brokers h at σ' .

(R6) Assume that at σ agent i controls h and agent $i' \in \overline{I_{\sigma}}$ controls $h' \in \overline{H_{\sigma}}$. Then, i' owns h at $\sigma \cup \{(i, h')\}$, and if, in addition, i brokers h at σ but not at σ' and $i' \in \overline{I_{\sigma'}}$, then i' owns h at σ' .

Requirements R4 and R5 postulate that control rights persist: agents hold on to control rights as we move from smaller to larger submatchings, or through the rounds of the algorithm. Requirements R4 and R6 imply that whenever an agent i' takes over the control of the brokered house, then the broker i is in line for ownership rights of i' after i' is matched (i becomes the heir to i'). We refer to this implication as a *broker-to-heir transition*. For discussion of the assumptions see Pycia and Ünver [2009].

To sidestep the complications of R5 and R6 in the first reading, the reader is invited to keep in mind a smaller class of control rights structures in which both of these requirements are replaced by the following strong form of brokerage persistence: “If $|\sigma'| < |I| - 1$ and agent i brokers house h at σ then i brokers h at σ' .” We think that by restricting attention to this smaller class of control rights structures, we are not missing much of the flexibility of the TC class of mechanism. We must stress, however, that the complication is there for a reason: see Pycia and Ünver [2009].

We are now ready to define our mechanism class and state our main results.

Definition 2. A control rights structure is **consistent** if it satisfies requirements R1-R6. The class of **TC mechanisms** (trading cycles) consists of mappings from agents’ preference profiles \mathbf{P} to matchings \mathcal{M} obtained by running the TC algorithm with consistent control rights structures.

The TTC and the non-TTC mechanism of Pycia and Ünver [2009] Section 3 are examples of TC. We will denote by $\psi^{c,b}$ the TC mechanism obtained from a consistent control rights structure $\{(c_{\sigma}, b_{\sigma})\}_{\sigma \in \overline{\mathcal{M}}}$.

Theorem 1. *Every TC mechanism is individually strategy-proof and Pareto efficient.*

The proof of this result is the same as in Pycia and Ünver [2009].

5 House Allocation and Exchange

In this section, we generalize the model by allowing agents to have private endowments. The characterizations in the resulting allocation and exchange domains are straightforward corollaries of our main results. We also relate the results to allocation and exchange market design environments.

Let $\mathcal{H} = \{H_i\}_{i \in \{0\} \cup I}$ be a collection of $|I| + 1$ pairwise-disjoint subsets of H (some of which might be empty) such that $\cup_{i \in \{0\} \cup I} H_i = H$. We interpret houses from H_0 as the social endowment of the agents, and houses from H_i , $i \in I$, as the private endowment of agent i . A **house allocation and exchange problem** is a list $\langle H, I, \mathcal{H}, \succ \rangle$. Since we allow some of the agents to have empty endowment, the allocation model of Section 2 is contained as a special case with $\mathcal{H} = \{H, \emptyset, \dots, \emptyset\}$. We may fix H, I and \mathcal{H} , and identify the house allocation and exchange problem just by its preference profile \succ . Matchings and mechanisms are defined as in the allocation model of Section 2.

Pareto efficiency and group strategy-proofness are defined in the same way as in Section 2. In particular, the equivalence between group strategy-proofness and the conjunction of individual strategy-proofness and non-bossiness continue to hold true. In addition to efficiency and strategy-proofness, satisfactory mechanisms in this problem domain should be individually rational. A mechanism is **individually rational** if it always selects an individually rational matching. A matching is individually rational, if it assigns each agent a house that is at least as good as the house he would choose from among his endowment. Formally, a matching μ is individually rational if

$$\mu(i) \succeq_i h \quad \forall i \in I, \forall h \in H_i.$$

For agents with empty endowments, $H_i = \emptyset$, this condition is tautologically true.

The following result for house allocation and exchange is now an immediate corollary of the results of previous section.

Theorem 2. *In house allocation and exchange problems, TC mechanisms are Pareto efficient, and strategy-proof.*

Furthermore, it is straightforward to identify individually rational TC mechanisms. Referring to control rights at the empty submatching as the initial control rights, let us formulate the criterion for individual rationality as follows.

Proposition 2. *In house allocation and exchange problems, a TC mechanism is individually rational if and only if it may be represented by a consistent control rights structure in which each agent is given the initial ownership rights of all houses from his endowment.*

The proof is the same as in Pycia and Ünver [2009].

The subclass of TC mechanisms without brokers is the extension of Pápai [2000] hierarchical exchange to the setting with object copies.

6 Outside Options

In this final section, we drop the assumption that $|H| \geq |I|$ and allow agents to prefer their (non-tradeable) outside options to some of the houses. Thus, some agents may be matched with their outside options, and we need to slightly modify some of the definitions. As before, I is the set of agents and H is the set of houses. Each agent i has a strict preference relation \succ_i over H and his outside option, denoted y_i . We denote the set of outside options by Y . The houses preferred to the outside option are called **acceptable** (to the agent); the remaining houses are called **unacceptable** to this agent. As before, we denote by \mathbf{P}_i the set of agent i 's preference profiles, and $\mathbf{P}_J = \times_{i \in J} \mathbf{P}_i$ for any $J \subseteq I$.

Let us initially restrict our attention to house allocation problems. This restriction can be easily relaxed as in Section 5, and we do so at the end of the section. As before, a house allocation problem is the triple $\langle I, H, \succ \rangle$. We impose no assumption on the cardinalities of I and H ; in particular, we allow both $|H| \geq |I|$ and $|H| < |I|$.

We generalize the concept of submatching as follows: For $J \subseteq I$, a submatching is a one-to-one function $\sigma : J \rightarrow H \cup Y$ such that each agent is matched with a house or his outside option.

A terminological warning is in order. A natural interpretation of the outside option is remaining unmatched. We will not refer to the outside option in this way, however, in order to avoid confusion with our submatching terminology. As in the main body of the paper, whenever we say that an agent is unmatched at σ , we refer to agents from $\overline{I}_\sigma = I - I_\sigma$. An agent is considered matched even if he is matched to his outside option.

As before, \mathcal{S} is the set of submatchings, I_σ denotes the set of agents matched by σ , $\overline{H}_\sigma \subseteq H$ denotes the set of houses with unallocated copies at σ , and we use the standard

function notation so that $\sigma(i)$ is the assignment of agent $i \in I_\sigma$, $\sigma^{-1}(h)$ is the set of agents that got house $h \in \sigma(I_\sigma)$, and $\sigma^{-1}(Y)$ is the set of agents matched to their outside options. A matching is a maximal submatching, that is, $\mu \in \mathcal{S}$ is a matching if $I_\mu = I$. As before, $\mathcal{M} \subset \mathcal{S}$ is the set of matchings. A (direct) mechanism is a mapping $\varphi : \mathbf{P} \rightarrow \mathcal{M}$ that assigns a matching for each preference profile (or, equivalently, allocation problem). Mechanisms, efficiency, and group strategy-proofness are defined as before.

The control rights structures (c, b) and their consistency R1-R6 are defined as before (notice though that the meaning of some terms such as submatching has changed, as explained above). In particular, (i) only houses are owned or brokered, the outside options are not; and (ii) control rights are defined for all submatchings, including submatchings in which some agents are matched with their outside options. Notice that if a control rights structure is consistent on the domain with outside options, and $|H| \geq |I|$, then the restriction of the control rights structure to submatchings in which all agents are matched with houses is a consistent control rights structure in the sense of Sections 4-5.

We will adjust the definition of the TC algorithm by adding two clauses.

Clause (a). We add the following provision to Step 1 (pointing) of round r :

- If an agent prefers his outside option to all unmatched houses, the agent points to the outside option. If there is a broker for whom the brokered house is the only acceptable house, such a broker also points to his outside option. The outside option of each agent points to the agent.

- We modify the definition of σ^r in Step 3 (matching); σ^r is defined as the union of σ^{r-1} and the set of agent-house pairs and agent-outside option pairs matched in Step 3.

Clause (b). In Step 3, we do not match agents in the cycle containing the broker except if leaving this cycle unmatched implies that no cycle is matched in the current round.

Clause (a) accommodates outside options. Clause (b) is added to ensure that we do not match a broker with his outside option when he prefers the brokered house to the outside option and the brokered house is not allocated to any other agent. Notice that the broker is matched only if any pointing sequence that starts with an owner ends by pointing to the broker.

We will refer to the algorithm of Section 4 modified by clauses (a) and (b) as outside options TC, and when there is no risk of confusion, simply as TC. We will refer to the mechanism $\psi^{c,b}$ resulting from running the outside options TC on consistent control rights structures as **outside options TC**, or **TC**. Using the same name is justified because the mechanism described above can be used to allocate houses in the setting of Sections 2-6, and – when restricted to the case of $|H| \geq |I|$ and the subdomain of preferences in which

all agents prefer any house to their outside option in the setting – is identical with the TC mechanism of Section 5. Indeed, in the restricted setting clause (a) is never invoked, and presence or absence of clause (b) has no impact on the allocation. This follows from the group strategy-proofness of TC of Section 5. Given a profile of agents’ preferences, agents who are brokers along the run of the TC without clause (b) can replicate the run of TC with clause (b) as follows: The first agent who becomes a broker along the path of the algorithm reports all houses that are matched in cycles not involving the broker ahead of the house the broker will be allocated, while keeping his preference profile otherwise intact. If another agent becomes a broker after the first broker is matched or loses his brokerage right, we modify this agent’s preferences in the same way, and the same for other brokers. If the outcomes of the mechanism were dependent on whether the brokers simulate clause (b) or not, there would be a preference profile in which one of the brokers could either improve his outcome or boss other agents; contrary to group strategy-proofness. By Theorem 1 this is not possible. The fact that clause (b) does not impact allocation in the setting without outside options is analogous to the well-known fact that in TTC the order in which we match the cycles of agents does not matter.

Theorem 3. *In the environment with outside options, every TC mechanism is strategy-proof and Pareto efficient.*

We are now ready to extend the characterization to the general allocation and exchange setting with outside options. As in Section 5, the social endowment $H_0 \subset H$ and agents’ endowments $H_i \subset H$, $i \in I$, are disjoint and sum up to H . A house allocation and exchange problem is a list $\langle H, I, \mathcal{H}, \succ \rangle$ where $\mathcal{H} = \{H_i\}_{i \in \{0\} \cup I}$. The results of Section 5 translate to the setting with outside options.

Theorem 4. *In house allocation and exchange problems with outside options, TC mechanisms are Pareto efficient, and strategy-proof.*

As before, it is straightforward to identify individually rational TC mechanisms.

Proposition 3. *In house allocation and exchange problems with outside options, a TC mechanism is individually rational if and only if it may be represented by a consistent control rights structure in which each agent is given the initial ownership rights of all houses from his endowment.*

In the environment with outside options, we define the TTC mechanisms as TC with no brokers.

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88:233–260, 1999.
- Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93:729–747, 2003.
- Anna Bogomolnaia and Herve Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.
- Partha Dasgupta, Peter Hammond, and Eric Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies*, 46:185–216, 1979.
- Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87:293–314, 1979.
- Szilvia Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68:1403–1433, 2000.
- Parag A. Pathak and Tayfun Sönmez. Leveling the playing field: Sincere and sophisticated players in the boston mechanism. *American Economic Review*, 98:1636–1652, 2008.
- Marek Pycia and M. Utku Ünver. Incentive compatible allocation and exchange of discrete resources. UCLA and Boston College, unpublished mimeo, 2009.
- Alvin E. Roth, Tayfun Sönmez, and M. Utku Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119:457–488, 2004.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1:23–37, 1974.
- Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16:557–567, 1999.
- William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.