Incentive Compatible Allocation and Exchange of Discrete Resources*

Marek Pycia† M. Utku Ünver‡

UCLA *Boston College*

This Draft: November, 2011

Abstract

Allocation and exchange of many discrete resources – such as kidneys or school seats – is conducted via direct mechanisms without monetary transfers. A primary concern in designing such mechanisms is the coordinated strategic behavior of market participants and its impact on resulting allocations. To assess the impact of this implementation constraint, we construct the full class of group dominant-strategy incentive compatible and Pareto efficient mechanisms. We call these mechanisms "Trading Cycles." This class contains new mechanisms as well as such previously studied mechanisms as top trading cycles, serial dictatorships, and hierarchical exchange. In some problems, the new trading-cycles mechanisms perform better than all previously known mechanisms. Just as importantly, knowing that all group incentive-compatible and efficient mechanisms can be implemented as trading cycles allows us to determine easily which efficient outcomes can and cannot be achieved in a group incentive-compatible way.

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

JEL classification: C78, D78

---

*We thank seminar participants at Pittsburgh, Rochester, UCLA, 2008 Caltech Mini Matching Workshop, 2008 Montreal SCW Conference, 2008 Pittsburgh ES North American Summer Meeting, Koç, Northwestern, U Washington St Louis, 2009 Bonn Mechanism Design Conference, Columbia, Harvard-MIT, 2010 Stony Brook International Conference of Game Theory, Microsoft Research Lab in New England, 2010 NBER Market Design Workshop, 2010 Yonsei University Market Design Conference, and at Rice, as well as Andrew Atkeson, Sophie Bade, Haluk Ergin, Manolis Galenianos, Ed Green, Matthew Jackson, Onur Kesten, Fuhito Kojima, Sang-Mok Lee, Vikram Manjunath, Szilvia Pápai, Al Roth, Andrzej Skrzypacz, Tayfun Sönmez, William Thomson, Özgür Yilmaz, and William Zame for comments. Ünver gratefully acknowledges the research support of National Science Foundation through grants SES #0338619, SES #0616689, and Microsoft Research Lab in New England.

†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 coordinated strategic behavior of market participants and its impact on resulting allocations. For instance, auction procedures such as the well-known Vickrey-Clarke-Groves mechanisms – though individually dominant-strategy incentive compatible – are susceptible to collusive behavior by bidders (cf. Vickrey, 1961; Clarke, 1971; Groves, 1973; Bennett and Conn, 1976). In environments in which transfers are not used, cooperation among market participants is also well documented. Coordination among parents participating in school choice was documented by Pathak and Sönmez (2008). In the allocation and exchange of transplant organs, a doctor acting on behalf of several patients can coordinate their reports if it benefits his or her patients. There are known cases of doctors gaming medical systems for the benefit of their patients.[1] In kidney exchange, transplant centers occasionally try to first conduct kidney exchanges using their internal patient-donor pool, and list their patient and donors in outside exchange programs only if they fail to find a suitable match, thus hindering the efficiency of regional exchange systems (cf. Sönmez and Ünver, 2010; Ashlagi and Roth, 2011).

Such coordinated strategic behavior is avoided if the market mechanism is group dominant-strategy incentive compatible. Group incentive compatibility means that no group of agents can jointly manipulate the system so that all of them weakly benefit from this manipulation, while at least one in the group strictly benefits. For instance, group incentive compatibility of a transplant allocation and exchange mechanism guarantees that no doctor is able to manipulate the mechanism on behalf of his or her patients without harming at least one of them. Non-manipulability is not the only benefit of using group incentive-compatible

[1] For instance, in 2003 two Chicago hospitals settled a Federal lawsuit alleging that some patients had been fraudulently certified as sicker than they were to move them up on the liver transplant queue (Warmbir, 2003).

mechanisms. Group incentive-compatible mechanisms impose 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, Hammond, and Maskin, 1979; Pathak and Sönmez, 2008).

What then are group incentive-compatible mechanisms? Which efficient outcomes can they achieve? In environments in which agents have quasilinear utilities and monetary transfers can be used, these questions have been answered by Green and Laffont (1977, 1979), who showed that no group dominant-strategy incentive-compatible mechanism leads to efficient outcomes, while every individually dominant-strategy incentive-compatible and efficient mechanism is equivalent to a Vickrey-Clarke-Groves mechanism. In contrast, in the above-mentioned environments without transfers, these questions remained unanswered.

This paper addresses these questions by constructing a new class of group strategy-proof and efficient mechanisms – *trading cycles* – and showing that every group dominant-strategy incentive-compatible and efficient mechanism is equivalent to a trading cycle mechanism. This equivalence is the main result of the paper. Some special cases of our class such as top trading cycles, serial dictatorships, and hierarchical exchange have already been studied; many mechanisms in our class have never been studied before. As illustrated below, in some problems, employing the new trading-cycles mechanisms does improve on all previously studied mechanisms. The new trading-cycles mechanisms also allow one to implement efficient allocations that cannot be implemented by previously studied mechanisms.

Just as importantly, knowing that all group dominant-strategy incentive-compatible and efficient mechanisms may be represented in our class allows us to easily determine what can and cannot be achieved in a group incentive-compatible way. The equivalence result radically simplifies any analysis of such questions because it allows us to restrict our attention to trading cycles without loss of generality. In this sense, the class of trading-cycles mechanisms is a natural benchmark of which efficient outcomes can be achieved in a group incentive-compatible way in the no-transfers environments we study, just as the Vickrey-Clarke-Groves class provides such a benchmark for individual incentive compatibility in environments with transfers.

Before describing the trading-cycles mechanisms, 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, using the terminology coined by Shapley and Scarf (1974). Agents have strict preferences over the houses.[2] 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

[2]We will discuss the strictness assumption later in the introduction.

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).

As a first step in the description of the trading-cycles let us discuss its special case – Gale's top-trading-cycles mechanism (reported by Shapley and Scarf, 1974) – in a special environment in which there are as many houses as agents and each agent is endowed with one house (and owns or controls the house). Gale's top trading cycles resembles decentralized trading and matches agents and houses in a sequence of rounds. In each round, each house points to its owner and each agent points to his most preferred unmatched house. Since there is a finite number of agents, there exists at least one pointing cycle in which an agent, say agent 1, points to a house, say house A; the agent who controls house A points to house B, etc.; and finally the last agent in the cycle points to the house controlled by agent 1. The pointing cycles might be short (agent 1 points to house A, which points back to agent 1) or might involve many agents. The procedure then matches each agent in each pointing cycle with the house he points to. The pointing cycles thus become cycles of trading. Rounds are repeated until no agents and houses are left unmatched. The top-trading-cycles algorithm can be used in more general environments, for instance when all houses are socially endowed. In such environments, we need to specify for each round and each house which agent plays the role of the owner of this house (see Abdulkadiroğlu and Sönmez, 1999; Pápai, 2000).

Our trading-cycles algorithm builds on the top-trading-cycles idea. However, while top trading cycles matches agents only along "top" cycles – cycles in which each agent points to and is matched with his most preferred previously unmatched house – trading cycles allows exchanges along any cycles that preserve group strategy-proofness and efficiency. Our results show that not all such cycles are top cycles: in addition to top cycles, group strategy-proof and efficient trading can be conducted along certain cycles in which one of the agents (whom we call "broker" below) points to and is matched with his second-most preferred unmatched house.

We are now ready to describe the trading-cycles algorithm. The algorithm matches houses and agents in a sequence of rounds. At each round some agents and houses are matched and the matches are final. At the beginning of the round, each previously unmatched house is controlled by a unique unmatched agent. We distinguish two forms of control over a house, which we call *ownership* and *brokerage*. The trading-cycles algorithm takes the allocation of control rights in each round as given and such that there is at most one broker and one brokered house. Each house points to the agent that controls it, and each agent except the broker (if there is one) points to his most preferred unmatched house. If there is a broker, he points to his most preferred unmatched house other than the brokered house. Since there is a finite number of agents, there exists at least one pointing cycle. The algorithm clears the pointing cycle by matching each agent with the house he points to.[3]

The allocation of control rights in each round is fully determined by how agents and houses were matched prior to that round. The above-described algorithm takes as given the mapping from such partial matchings to control rights. Each such mapping that satisfies certain consistency conditions determines a mechanism in our class. For expositional purposes, we first formulate our main result for settings in which all houses are social endowments, and hence there are no additional exogenous constraints on the allocation of control rights (cf. Hylland and Zeckhauser, 1979). For instance, at some universities, the dormitory rooms are treated as social endowments. At other universities, however, some students, such as sophomores, have the right to stay in the room they lived in the preceding year. In kidney exchange, patients (interpreted as agents) come with a paired-donor (interpreted as a house) and have to be matched with at least their paired-donor. Such exogenous control rights are straightforwardly accommodated by our mechanism class. When some houses are private endowments of agents it is natural to require that the participation in the mechanism is individually rational so that each agent likes the mechanism's outcome at least as much as the best house from his endowment. The class of group strategy-proof, efficient, and individually rational direct mechanisms equals the class of individually rational trading-cycles mechanisms. A trading-cycles mechanism is individually rational if and only if it may be represented by a consistent control rights structure in which each agent is given ownership rights over all houses from his endowment.[4]

Recognizing the role of brokers in house allocation and exchange is crucial to obtaining the entire class of group strategy-proof and Pareto-efficient mechanisms.[5] Previously only trading-cycles mechanisms with ownership control rights were studied (see also Shapley and

[3] We study environments both with and without outside options. The results are the same in both environments, but the above algorithm needs to be slightly generalized in the case of outside options by allowing agents to point to houses or their outside options. We also need to postpone matching a broker with his outside option until a round in which an agent who owns a house lists the brokered house as his most preferred one.

[4] In particular, this implies that when each agent has a private endowment then top-trading-cycles mechanisms are the unique mechanisms that are group strategy-proof, efficient, and individually rational. A variant of this corollary in the special setting in which there are as many houses as agents and each agent is endowed with exactly one house was earlier proven by Ma (1994) (Pápai (2007) extended Ma's result to multi-unit environments in which all objects are initially owned).

[5] It is natural to ask whether we can run an analogue of trading cycles with more than one broker in a given round. The answer is negative; such a mechanism would not be strategy-proof and efficient. As we explain in the paper, at-most-one-broker-per-round is an inherent feature of group strategy-proofness and efficiency, and not merely a convenient simplification.

Scarf, 1974; Abdulkadiroğlu and Sönmez, 1999; Pápai, 2000). The introduction of brokers is also useful in some mechanism design problems.

As an example of a mechanism design problem in which brokerage rights are useful, consider a manager who assigns n tasks $t_1, ..., t_n$ to n employees $w_1, ..., w_n$ with strict preferences over the tasks. One of the employees, w_1, is the manager's nephew. The manager wants the allocation to be Pareto efficient with regard to the employees' preferences. Within this constraint, she would like to avoid assigning task t_1 to her nephew w_1. She wants to use a group strategy-proof direct mechanism, because she does not know employees' preferences. The only way to do this using the previously known mechanisms is through Pápai (2000)'s class, endowing employees $w_2, ..., w_n$ with the tasks, let them find the Pareto-efficient allocation via the top-trading-cycles algorithm, and then allocating the remaining task to the nephew. Every procedure like that privileges all employees $w_2, ..., w_n$ over the nephew w_1: the nephew is allocated the task no one else wanted. Using a trading-cycles mechanism, the manager can improve the outcome of her nephew while still achieving her objective. To do so, she makes the nephew the broker of t_1, allocates the remaining tasks among $w_2, ..., w_n$ (for instance, she may make w_i the owner of t_i, $i = 2, ..., n$), and runs the trading cycles. The allocation of employee w_1 in this trading-cycles mechanism is better than in any top-trading-cycles procedure satisfying the manager's constraints: the allocation is weakly better regardless of agents' preference profile, and it is strictly better for some preference profiles. In particular, trading cycles can be used to implement Pareto-efficient outcomes that cannot be implemented by any previously known mechanism.

Our results may be also used to understand what cannot be achieved in a group strategy-proof way. For instance, if the manager in the above example would like to avoid assigning two tasks t_1 and t_2 to her nephew w_1, then she must run a trading-cycles mechanism in which the nephew initially does not have control rights over any task, and can acquire such control rights only when one of the tasks t_1, t_2 is allocated. As before, many brokered trading-cycles mechanism would make the nephew better off than all top-trading-cycles mechanisms. However, the manager cannot allow the nephew to participate in the initial round of trading. This form of discrimination against the nephew is unavoidable and is caused by the need to learn employees' preferences in a group strategy-proof way.

To illustrate the applicability of our results, in Appendix B we show how some of the deepest insights in the no-transfer market design theory are immediate corollaries of our theorems: (i) Pápai (2000)'s result that all reallocation-proof, efficient, and group strategy-proof mechanisms are implementable as hierarchical exchange and (ii) Svensson (1999)'s result that a mechanism is neutral and group strategy-proof if and only if it is a serial dictatorship (we provide more details on these two results in the overview of the literature below). Appendix B also shows how our theorems can be used to generate new insights into subclasses of group strategy-proof and efficient mechanisms.

The no-transfers market design has been studied extensively. We already compared our result to one of the two most influential mechanism: top trading cycles. Another mechanism that has been very influential in studies of object allocation is serial dictatorships (Satterthwaite and Sonnenschein, 1981; Svensson, 1994, 1999; Ergin, 2000). It is group strategy-proof and efficient, and hence is a special case of trading cycles. When houses are social endowments, we can run a serial dictatorship by ordering agents in a priority queue and matching each agent, in order, with his most preferred house among houses that were not matched with higher-priority agents. The trading cycles can replicate any serial dictatorship by endowing the first agent in the priority queue with ownership rights over all houses in the first round, endowing the second agent in the priority queue with ownership rights over all unmatched houses in the second round, etc.[6]

The largest subclass of group-strategy proof and efficient allocation and exchange mechanisms in the literature prior to our study was constructed by Pápai (2000) in a very insightful paper. She focuses on the allocation problem and constructs a class of mechanisms referred to as top trading cycles or hierarchical exchange, which use the same algorithm as Gale's top-trading-cycles mechanism with the exception that the mechanism takes as an input a structure of control rights (ownership only) over houses that – for each round of the mechanism and each unmatched house – determines the agent to whom the house points to. Her class characterizes group strategy-proofness and Pareto efficiency together with an additional property that she refers to as *reallocation-proofness*. A mechanism is reallocation-proof in the sense of Pápai (2000) if there is no profile of preferences with a pair of agents and a pair of preference manipulations such that (i) if both of them misrepresent their preferences, both of them weakly gain and one of them strictly gains by swapping their assignments, and (ii) if only one of them misrepresents his preferences, he cannot change his assignment. Pápai also notes that the stronger reallocation-proofness-type property obtained by dropping condition (ii) conflicts with group strategy-proofness and Pareto efficiency or allowing the swap of assignments among more than two agents. We do not use reallocation-proofness in our results.

All the above work share with our paper the assumption that agents have strict preferences. This is the standard modeling assumption in analysis of matching and house allocation and exchange because – as Ehlers (2002) shows – "one cannot go much beyond strict preferences if one insists on efficiency and group strategy-proofness." The full preference domain

[6]The serial dictatorships also belong to Pápai's top-trading-cycles class. The connections between serial dictatorships and top trading cycles have been also explored by Abdulkadiroğlu and Sönmez (1999).

gives rise to an impossibility result, i.e., when agents can be indifferent among houses, there exists no mechanism that is group strategy-proof and Pareto efficient.[7] For this reason, participants are frequently allowed to submit only strict preference orderings to real-life direct mechanisms in various markets, such as dormitory room allocation, school choice, and matching of interns and hospitals.[8]

The study of strategy-proof mechanisms has a long tradition. Gibbard (1973) and Satterthwaite (1975) have shown that all strategy-proof and unanimous voting rules are dictatorial. Satterthwaite and Sonnenschein (1981) extended this result to public goods economies with production, Zhou (1991) extended it to pure public goods economies, and Hatfield (2009) to group strategy-proof quota allocations. In social choice models, Dasgupta, Hammond, and Maskin (1979) have proved that every Pareto-efficient and strategy-proof social choice rule is dictatorial. In exchange economies, Barberà and Jackson (1995) showed that strategy-proof mechanisms are Pareto inefficient. Characterizations of Pareto-efficient and strategy-proof mechanisms that are non-dictatorial have been obtained by Green and Laffont (1977) for design problems with monetary transfers and quasi-linear utilities (cf. Vickrey, 1961; Clarke, 1971; Groves, 1973; Roberts, 1979); by Barberà, Jackson, and Neme (1997) for sharing a perfectly divisible good among agents with single-peaked preferences over their shares(cf. Sprumont, 1991); and by Barberà, Gül, and Stacchetti (1993) for voting problems with single-peaked preferences (cf. Moulin, 1980).[9]

2 Model

2.1 Environment

Let I be a set of **agents** and H be a set of **houses.** We use letters i, j, k to refer to agents and h, g, e to refer to houses. Each agent i has a **strict preference relation** over H, denoted

[7] Ehlers also characterizes group strategy-proof and Pareto-efficient mechanisms in the maximal subset of full preference domain such that such a mechanism exists. Under strict preferences, his class of mechanisms is a subclass of ours, and substantially different from the general class. See Bogomolnaia, Deb, and Ehlers (2005) for another characterization with indifferences.

[8] In school choice there is a tension between efficiency and priority-based fairness; see for example Balinski and Sönmez (1999); Ergin (2002); Abdulkadiroğlu and Sönmez (2003); Ehlers and Klaus (2006); Kesten (2006); Erdil and Ergin (2008); Kojima and Manea (2010); Abdulkadiroğlu and Che (2010); Abdulkadiroğlu, Che, and Yasuda (2011).

[9] Sönmez (1999) studies generalized matching problems in which each agent is endowed with a good. The class of such problems non-trivially intersects with the class of house allocation and exchange problems studied in this paper. He shows that there exists a Pareto-efficient, strategy-proof, and individually rational mechanism if and only if the core is nonempty and agents are indifferent between all core allocations. He also shows that any such mechanism is group strategy-proof (cf. Shapley and Scarf, 1974; Roth and Postlewaite, 1977; Roth, 1982; Ma, 1994). A related problem of preference aggregation has also been intensively studied, beginning with Arrow (1950).

by $\succ_i$.[10] 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 $\succ_J = (\succ_i)_{i \in J} \in P_J$ 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 house 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 6, 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) 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 7, we 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 one-to-one function $\sigma : J \to H$; where for $J \subseteq I$, 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 agent 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 and $H_\sigma \subseteq H$ denote the set of houses matched by σ. For all $h \in H$, let $\mathcal{S}_{-h} \subset \mathcal{S}$ be the set of submatchings $\sigma \in \mathcal{S}$ such that $h \in H - H_\sigma$, i.e., the set of submatchings at which house h is unmatched. 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 will write $\overline{I_\sigma}$ for $I - I_\sigma$, and $\overline{H_\sigma}$ for $H - H_\sigma$ for short. We will also write $\overline{\mathcal{M}}$ for $\mathcal{S} - \mathcal{M}$.

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

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

2.2 Group 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

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

In our domain group strategy-proofness has a non-cooperative interpretation, and is equivalent to the conjunction of two non-cooperative properties: individual strategy-proofness and non-bossiness. The strategy-proofness of a mechanism means that the truthful revelation of preferences is a weakly dominant strategy: 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

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

Non-bossiness (Satterthwaite and Sonnenschein, 1981) means that no agent can misreport his preferences in such a way that his allocation is not changed but the allocation of some other agent is changed: a mechanism φ is **non-bossy** if for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ'_i \in \mathbf{P}_i$ such that

$$\varphi[\succ'_i, \succ_{-i}](i) = \varphi[\succ](i) \quad \text{and} \quad \varphi[\succ'_i, \succ_{-i}] \neq \varphi[\succ].$$

The following lemma due to Pápai (2000) states the non-cooperative interpretation of group strategy-proofness:

Lemma 1. *Pápai (2000) A house-allocation mechanism is group strategy-proof if and only if it is individually strategy-proof and non-bossy.*

Another useful formulation of group strategy-proofness builds on Maskin (1999). A mechanism φ is **Maskin monotonic** if $\varphi[\succ'] = \varphi[\succ]$ whenever $\succ' \in \mathbf{P}$ is a φ-monotonic transformation of $\succ \in \mathbf{P}$. A preference profile $\succ' \in \mathbf{P}$ is a φ**-monotonic transformation** of $\succ \in \mathbf{P}$ if

$$\{h \in H : h \succeq_i \varphi[\succ](i)\} \supseteq \{h \in H : h \succeq'_i \varphi[\succ](i)\} \text{ for all } i \in I.$$

Thus, for each agent, the set of houses better than the base-profile allocation weakly shrinks when we go from the base profile to its monotonic transformation. The following lemma was proven by Takamiya (2001) for a subset of the problems we study. His proof extends word for word to our domain of problems.

Lemma 2. *A house-allocation mechanism is Maskin monotonic if and only if it is group strategy-proof.*

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.

Pareto efficiency is a very weak requirement when imposed on group strategy-proof mechanisms. Every group strategy-proof mechanism that maps $\mathbf{P}$ onto the entire set of matchings $\mathcal{M}$ is Pareto efficient. This surjectivity property, which refer to as *full range*, is implied, for instance, by *unanimity* of Gibbard (1973) and Satterthwaite (1975). A house allocation mechanism is unanimous if the mechanism allocates all agents their most-preferred houses whenever no two agents rank the same house as their most-preferred choice (that is, the overall matching most preferred by all agents obtains whenever the agents agree on the most preferred matching).

3 Beyond Top Trading Cycles

3.1 Top Trading Cycles

To set the stage for our trading-cycles (TC) mechanism, let us look at the well-known top-trading-cycles (TTC) algorithm adapted by Pápai (2000) to house allocation problems.[11] The class of mechanisms presented in this section is identical to Pápai's "hierarchical exchange" class. Our presentation, however, is novel and aims to simultaneously simplify the earlier constructions of Pápai's class, and to introduce some of the terminology we will later use to introduce our class of all group strategy-proof and efficient mechanisms (TC).

TTC is a recursive algorithm that matches houses to agents in a sequence of rounds. In each round, some agents and houses are matched. The matches will not be changed in subsequent rounds of the algorithm.

At the beginning of each round, each unmatched house is "owned" by an unmatched agent. The algorithm creates a directed graph in which each unmatched house points to the

[11] The algorithm was originally proposed by David Gale for the special case of house exchange (cf. Shapley and Scarf, 1974).

agent who owns it, and each unmatched agent points to his most preferred house among the unmatched houses. In the resultant directed graph there exists at least one exchange cycle in which agent 1's most preferred house is owned by agent 2, agent 2's most preferred house is owned by agent 3, ..., and finally, for some $k = 1, 2, ...$, agent k's most preferred house is owned by agent 1. Moreover, no two exchange cycles intersect. The algorithm matches all agents in exchange cycles with their most preferred houses.

The algorithm terminates when all agents are matched. As at least one agent-house pair is matched in every round, the algorithm terminates after finitely many rounds.

As we see, the outcome of the TTC algorithm is determined by two types of inputs: agents' preferences and agents' rights of ownership over houses. The preferences are, of course, submitted by the agents. The ownership rights are defined exogenously as part of the mechanism.[12] We formalize this aspect of the mechanism via the following concept.

Definition 1. A **structure of ownership rights** is a collection of mappings $\left\{c_\sigma : \overline{H_\sigma} \to \overline{I_\sigma}\right\}_{\sigma \in \overline{\mathcal{M}}}$. The structure of ownership rights $\{c_\sigma\}_{\sigma \in \overline{\mathcal{M}}}$ is **consistent** if

$$c_\sigma^{-1}(i) \subseteq c_{\sigma'}^{-1}(i) \text{ if } \sigma \subseteq \sigma' \in \overline{\mathcal{M}} \text{ and } i \in \overline{I_{\sigma'}}.$$

The structure of ownership rights tells us at each submatching which unmatched agent owns any particular unmatched house. Agent i owns house h at submatching σ when $c_\sigma(h) = i$. Consistency means that whenever an agent owns a house at a submatching (σ) then he also owns it at any larger submatching (σ') as long as he is unmatched.

Each consistent structure of ownership rights $\{c_\sigma\}_{\sigma \in \overline{\mathcal{M}}}$ determines a *hierarchical exchange mechanism* of Pápai (2000). This class of mechanisms consists of mappings from agents' preferences $\mathbf{P}$ to matchings $\mathcal{M}$ obtained by running the TTC algorithm with consistent structures of ownership rights. Because of this, we will also refer to hierarchical exchange as **TTC mechanisms**. Pápai showed that all TTC mechanisms are group strategy-proof and Pareto efficient.

Example 1. As an example, consider the TTC mechanism to allocate four houses $h_1, ..., h_4$ to three agents $i_1, ..., i_3$ given by the structure of ownership rights that allocates ownership

[12] Recall that we are studying an allocation problem in which objects are a collective endowment. In Section 6 we will enlarge the analysis to include exchange problems among agents with private endowments. In exchange problems, some of the mechanism's ownership rights are determined by individual rationality constraints.

of houses according to the following table:

h_1	h_2	h_3	h_4
i_1	i_2	i_3	i_1
i_3	i_1	i_2	i_3
i_2	i_3	i_1	i_2

That is, house h_1 is initially owned by i_1; at submatchings i_3 and it are not matched and i_1 is matched, it is owned by i_3; at submatchings i_1 and i_3 are matched and it is not matched, it is owned by i_2. Notice that the owner is uniquely determined and the ownership structure is consistent.

To see how the TTC algorithms run, let us apply this mechanism to the preference profile in which all agents i have the same preferences $\succ_i$:

$$\text{agent } i \text{ preferences: } h_1 \succ_i h_2 \succ_i h_3 \succ_i h_4.$$

In the first round, all agents point to house h_1, houses h_1 and h_4 point to agent i_1, house h_2 points to i_2, and house h_3 points to i_3. In this round, there is one exchange cycle, in which i_1 is matched with h_1.

In the second round, agents i_2,i_3 and houses h_2, h_3, h_4 are unmatched. House h_2 is still owned by i_2, while houses h_3, h_4 are still owned by i_3. In the resultant directed graph, there is again one exchange cycle in which i_2 points to h_2 and h_2 points to i_2, and they are matched.

In the third round, agent i_3 owns all unmatched houses, is matched with h_3, and the algorithm terminates.

The second round of this example illustrates two phenomena. First, we cannot allocate the ownership unconditionally, as this would leave unresolved the ownership of house h_4 after its initial owner, agent i_1, is matched with house h_1. Second, it illustrates the need for the consistency condition. If the ownership structure was not consistent, and, say, h_2 was owned by i_3 at $\sigma = \{(i_1, h_1)\}$ (that is, after i_1 left with h_1), then agent i_2 would have an incentive to misreport his preferences and claim that he prefers h_2 over all other houses.

Although under the above preference structure, all exchange cycles involve only one agent and one house, this is not generally true. Consider, for instance, the following preference profile in which i_2's preference between h_2 and h_3 is reversed:

$$\text{agent } i_1 \text{ preferences: } h_1 \succ_{i_1} h_2 \succ_{i_1} h_3 \succ_{i_1} h_4,$$
$$\text{agent } i_2 \text{ preferences: } h_1 \succ_{i_2} h_3 \succ_{i_2} h_2 \succ_{i_3} h_4,$$
$$\text{agent } i_3 \text{ preferences: } h_1 \succ_{i_3} h_2 \succ_{i_3} h_3 \succ_{i_3} h_4.$$

When this profile is reported, the first round is the same as above, but the exchange cycle in the second round has agent i_2 pointing to h_3, h_3 pointing to i_3, i_3 pointing to h_2, and h_2 pointing to i_2.

To appreciate the generality of the Pápai's class, notice that the serial dictatorship of Satterthwaite and Sonnenschein (1981) and Svensson (1994) is a special case of the TTC mechanisms in which at each submatching there is an agent who owns all unmatched houses.

3.2 Beyond Top Trading Cycles: An Example

How might a group strategy-proof and efficient non-TTC mechanism look like? To give an example, we will modify the TTC mechanism of Example 1.

Example 2. Consider three agents $i_1, ..., i_3$ and three houses $h_1, ..., h_3$ and an ownership structure that allocates ownership of houses according to the following table (obtained by dropping house h_4 in the ownership structure of the example of Subsection 3.2):

h_1	h_2	h_3
i_1	i_2	i_3
i_3	i_1	i_2
i_2	i_3	i_1

The owner is uniquely determined and the ownership structure is consistent. Given this structure, let us run TTC with one modification: agent i_1 is not allowed to point to house h_1 as long as there are other unmatched agents. In rounds with other unmatched agents (and hence other unmatched houses), agent i_1 will point to his most preferred house among unmatched houses other than h_1.[13]

For instance, if each agent i has the preference $h_1 \succ_i h_2 \succ_i h_3$ then in the first round agents i_2 and i_3 will point to h_1, but agent i_1 will point to his second-choice house, h_2. We

[13] Pápai (2000) gives an example of a non-TTC mapping from $\mathbf{P}$ to $\mathcal{M}$. Her construction is different from ours though the resultant mappings are identical. As we will show in the next section, the advantage of our construction lies in its generalizability to cover the whole class of group strategy-proof and efficient mechanisms.

will then have an exchange cycle in which i_1 is matched with h_2 and i_2 is matched with h_1. In the second round, the algorithm matches agent i_3 and house h_3, and terminates.

This mechanism is group strategy-proof and Pareto efficient. An easy recursion may convince us that at each round the submatching formed is Pareto efficient for matched agents. Indeed, if an agent matched in the first round does not get his top choice then he gets his second choice, and getting his first choice would harm another agent matched in that round. In general, agents matched in the n'th round get their first or second choice among houses available in the n'th round, and giving one of these agents a better house would harm some other agent matched at the same or earlier round. The intuition behind its group strategy-proofness is more complex, and we defer its discussion until our formal results.

The mechanism of Example 2 turns out to be different from all TTC mechanisms. To see this, first observe that the mechanism matches house h_1 with agent i_2 under the illustrative preference profile analyzed above, whereas it would match h_1 with another agent, i_3, if agent i_1 submitted preferences $h_1 \succ_{i_1} h_3 \succ_{i_1} h_2$ (and other agents $i \neq i_1$ continued to have preferences $h_1 \succ_i h_2 \succ_i h_3$). However, any TTC mechanism would match h_1 with the same agent in these two preference profiles. Indeed, TTC ownership structure uniquely determines which agent owns h_1 at the empty submatching, and this agent would be matched with h_1 in the first round of the algorithm under any preference profile in which all agents rank h_1 as their first choice.

For future use, notice that in the above example, agent i_1 does not have full ownership right over h_1. Unless he is the only agent left, he cannot form the trivial exchange cycle that would match him with h_1. He does have some control right over h_1, however: he can trade h_1 for houses owned by other agents. In our general trading-cycles algorithm, we will refer to such weak control rights as "brokerage."

4 Trading-Cycles Mechanism

We turn now to our new algorithm, trading cycles (TC), an example of which we saw in the previous section. Like TTC, the 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 2. A **structure of control rights** is a collection of mappings

$$\left\{(c_\sigma, b_\sigma) : \overline{H_\sigma} \to \overline{I_\sigma} \times \{\text{ownership,brokerage}\}\right\}_{\sigma\in\overline{\mathcal{M}}}.$$

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 $b_\sigma(h)$ =ownership, and that the agent $c_\sigma(h)$ **brokers** h at σ if $b_\sigma(h)$ =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.

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 A). 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 $r = 1, 2, ...$ 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, $\sigma^0 = \varnothing$. 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 $n \in \{1, 2, ...\}$ and an exchange cycle

$$h^1 \to i^1 \to h^2 \to ...h^n \to i^n \to h^1$$

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 $\ell = 1, ..., n$ 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 or all houses are matched.

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 described in Section 3.1, 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.[14]

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.

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 3.2 applies.

Proposition 1. *The outcome of the TC algorithm is Pareto-efficient for all control rights structures that satisfy R1-R3.*

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 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. R4 (*persistence of ownership*) is identical to the consistency assumption we imposed on TTC.

[14]In particular, TC can easily handle private endowments, as explained in Section 6.

The first example of Section 3.1 illustrated why we need such a persistence assumption for the resultant mechanism to be individually strategy-proof. A similar example might convince us that individual strategy-proofness relies also on requirement R5 (*persistence of brokerage*); see Appendix A.

Requirement R6 has two related parts. The first part (*consolation for lost control rights)* postulates that when an agent i is matched with a house controlled by i' then i' owns the houses previously controlled by i. R6's second part (*brokered-to-owned house transition*) postulates who obtains the control right over a house when a broker loses his brokerage right. By R5, the broker can only lose the brokerage right between σ and σ' when no more than one agent is a σ-owner and σ'-owner; this is the agent who obtains the control right at σ' over the house brokered at σ. A key implication of R6 is the transfer of ownership rights to ex-brokers: if i brokers h at σ but not at σ', and $i' \in \overline{I_{\sigma'}}$ owns $h' \in \overline{H_{\sigma'}}$ at σ, then R6's first part implies that i owns h' at $\sigma \cup \{(i', h)\}$, and R4 further implies that i owns h' at $\sigma' \cup \{(i', h)\}$. We refer to this consequence of R6 (and R4) as *broker-to-heir transition*. Requirement R6 is needed to guarantee both the non-bossiness and individual strategy-proofness of the mechanisms; see Appendix A.

We are now ready to define our mechanism class.

Definition 3. 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 mechanisms of Section 3.1 and the non-TTC mechanism of Section 3.2 are examples of TC. We will denote by $\psi^{c,b}$ the TC mechanism obtained from a consistent control rights structure $\{(c_\sigma, b_\sigma)\}_{\sigma \in \overline{\mathcal{M}}}$. In Section 6 we adapt this class of mechanisms to exchange problems, and in Section 7 we enlarge it to allow for agents' outside options (in particular, in Section 7 we allow agents to ranks some objects as unacceptable).

To sidestep the complication of condition 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, one is not missing much of the flexibility of the TC class of mechanism. We hasten to stress, however, that the complication is there for a reason: there are TC mechanisms that cannot be replicated by TC control rights structures satisfying the above strengthening of R5-R6; such TC mechanisms are group strategy-proof and efficient by the results of the next section. Let us also stress that, a priori, we could expect the class of group strategy-proof and efficient mechanisms to be much more complex than it turned out to be. In fact, Section 6 and Appendix B show that brokers and condition R6 are actually quite easy to work with.

Example 3. Can we replace R5-R6 by the following simpler (and stronger) property "if $|\sigma'| < |I| - 1$ and agent i brokers house h at σ and is unmatched at $\sigma' \supset \sigma$, then i brokers h at σ' " (an analogue of R4 for brokers)? The following example shows that we cannot. Consider an environment with four agents, i_1, i_2, i_3, i_4, four houses, h_1, h_2, h_3, h_4, and a TC mechanism $\psi^{c,b}$ whose control rights structure (c, b) is explained below and illustrated by the table in Figure 1.

$\mathbf{h_1}$	$\mathbf{h_2}$		$\mathbf{h_3}$	$\mathbf{h_4}$	
i_1,o	i_2,o		i_1,o	i_4,b	
i_3,o	$(i_1,h_1)(i_2,h_4)$ ↙	↘ *otherwise*	i_3,o	(i_1,h_1) ↙	↘ *i_2 and i_3 are matched, and i_1 is matched with h_2 or h_3*
i_2,o	i_4,o	i_1,o	i_2,o	i_2,o	i_4,o
i_4,o	i_3,o	i_3,o	i_4,o	i_4,o	
		i_4,o		i_3,o	

Figure 1: A control rights structure with broker-to-heir transition

Houses h_1, h_3 are owned by agent i_1 (denoted by "o" next to i_1 in the figure); he continues owning them as long as he is unmatched (R4 is satisfied). When i_1 is matched the unmatched of the two houses is owned by i_3 (if he is still unmatched). When both i_1 and i_3 are matched and h_1 or h_3 is unmatched, the house is owned by i_2. When all agents are matched and one of the houses h_1 or h_3 is unmatched, the house is owned by i_4.

House h_2 is owned by i_2. When i_2 is matched but h_2 is not then h_2 is inherited by one of the unmatched agents; who inherits h_2 depends on how the matched agents are matched (the submatching). If i_1 is matched with h_1 and i_2 is matched with h_4, then the next owners of h_2 are i_4 and i_3, in this order. In all other cases, the order of next owners of h_2 is i_1, i_3, and i_4.

House h_4 is initially brokered by agent i_4 (denoted by "b" next to i_4 in the figure). Agent i_4 continues to broker h_4 as long as he is unmatched with two exceptions: (i) if i_1 is matched with h_1 then i_4 looses the brokerage right, and h_4 becomes an owned house with the order of owners i_2, i_4, and i_3; and (ii) if i_4 is the only remaining agent, then he owns h_4. Notice that the first exception satisfies condition R6, while the second is dictated by R2. In general, we can check that the control right structure is consistent.

Let us now check that the TC mechanism defined by this control rights structure is different from all TC mechanisms with consistent control rights structures in which the simple analogue of R4 for brokers holds true: "if $|\sigma'| < |I| - 1$ and agent i brokers house h at σ and is unmatched at $\sigma' \supset \sigma$, then i brokers h at σ' ." By way of contradiction, let us assume that there is a TC mechanism ψ with a control rights structure satisfying the above strong form of brokerage persistence and produces the same allocation as $\psi^{c,b}$ for each profile of agents' preferences.

First, notice that at the empty submatching, i_4 is the broker of h_4 in ψ. This is so because h_4 is not owned by any agent at the empty submatching $\varnothing$ as $(\psi[\succ])^{-1}(h_4) = (\psi^{c,b}[\succ])^{-1}(h_4)$ varies with $\succ \in \mathbf{P}$ (that is, across profiles at which all agents rank h_4 first). Hence, there is an agent who has the brokerage right over h_4, and it must be i_4, as $\psi[\succ](i_4) = \psi^{c,b}[\succ](i_4) = g$ for all $\succ \in \mathbf{P}$ such that all agents rank h_4 first and any $g \in \{h_1, h_3, h_2\}$ second.

Second, consider the submatching $\sigma = \{(i_1, h_1)\}$ and a preference profile $\succ \in \mathbf{P}$ such that i_1 ranks h_1 first, others rank h_4, h_3, h_2, and h_1 in this order. In mechanism ψ, agent i_4 would continue to be the broker of h_4 at σ, and thus

$$\psi[\succ](i_4) = h_3.$$

However,

$$\psi^{c,b}[\succ](i_4) = h_2.$$

This contradiction shows that indeed the TC mechanism of the example cannot be represented by a control right structure in which brokerage satisfies the analogue of R4 for brokers (in particular it cannot be represented without brokers).

5 Main Results

Our main results tell us that the class of Trading Cycles mechanisms coincides with the class of Pareto-efficient and group strategy-proof direct mechanisms. In this section we state and prove it for the model of allocation in which all objects are acceptable. In the following two sections we relax both of these simplifying assumptions.

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

Theorem 2. *Every group strategy-proof and Pareto-efficient direct mechanism is TC.*

Let us start the discussion of the proof of Theorem 1 with the following observation about the TC algorithm.

Lemma 3. *If an agent i is unmatched at a round r of the algorithm under preference profiles $[\succ_i, \succ_{-i}]$ and $[\succ'_i, \succ_{-i}]$, then the control rights structure at round r is the same under $[\succ_i, \succ_{-i}]$ and $[\succ'_i, \succ_{-i}]$.*

The lemma obtains because its assumption implies that the same submatching was formed before round r whenever agent i submitted preference ranking $\succ_i$ or $\succ'_i$. Hence, the control rights structures must also be the same at round r. The lemma has an important implication: as long as an agent is unmatched, he cannot influence when he becomes an owner, a broker, or enters the broker-to-heir transition (see R6) by choosing which preferences to submit.

To see intuitively why trading cycles are strategy-proof, notice that the above lemma implies that no agent i can improve his match by being matched earlier. Owners cannot benefit by waiting since they get the best available house at the time they match under $\succ$. Checking that brokers cannot benefit by waiting is only slightly more subtle. We provide the details below.

As observed by Pápai (2000), to show not only that individual agents cannot benefit from manipulation, but also that groups of agents cannot, it is enough to show that the mechanism is non-bossy. Proving non-bossiness is harder and this part of the proof is relegated to Appendix C. To get a sense for the proof, consider a TC mechanism without brokers, and an agent i who gets the same object whether he submits preferences $\succ_i$ or $\succ'_i$. An inductive argument then shows that the algorithm will go though the same cycles under $\succ = (\succ_i, \succ_{-i})$ and $\succ' = (\succ'_i, \succ_{-i})$ even if the rounds at which these cycles are formed may differ. If brokers were strongly persistent, the same argument would apply. The difficulty in proving non-bossiness is when a broker loses his brokerage right. Condition R5 ensures that cycles of three agents or more are the same under both $\succ$ and $\succ'$, but that cycles of one or two agents can be different. For instance, in the setting of Example 3, consider a preference profile in which agents i_1 and i_3 rank houses $h_1 \succ_{i_1,i_3} h_4 \succ_{i_1,i_3} h_2 \succ_{i_1,i_3} h_3$ and agents i_2 and i_4 rank houses $h_4 \succ_{i_2,i_4} h_2 \succ_{i_2,i_4} h_3 \succ_{i_2,i_4} h_1$. Under this preference profile, $\succ_{\{i_1,i_2,i_3,i_4\}}$, in the first round, broker i_4 obtains object h_2 in a cycle $i_4 \to h_2 \to i_2 \to h_4 \to i_4$. However if i_2 submitted instead preference ranking $\succ'_{i_2}$ identical to $\succ_{i_1,i_3}$, then i_4 and i_2 would not swap houses in first round. They would both be still unmatched in round 2, and i_4 would have lost his brokerage right; house h_4 would now be owned by i_2 (notice that not only it is so in the example but in fact condition R6 requires that i_2 owns h_4 when i_1 becomes matched and i_4 looses the brokerage right). Agent i_2 would then match with h_4 in round 2. In round 3, agent i_4 would become owner of h_2 (again this is so in the example, and, importantly it is guaranteed by the broker-to-heir transition property of R6). Thus, in round 3 agent i_4 would match with h_2. While the cycles are different, the allocations are the same. Looking at requirement R6 (and its broker-to-heir corollary) can give us a sense why – even if one or two agent cycles are different at $\succ$ and $\left(\succ'_{i_2}, \succ_{-i_2}\right)$ – such or a similar scenario is bound to happen.

Beginning of the proof of Theorem 1. Proposition 1 demonstrates Pareto efficiency. By Lemma 1, to prove group strategy-proofness it is enough to show that every TC mechanism is individually strategy-proof and non-bossy. We will prove individual strategy-proofness below, and non-bossiness in Appendix C. Let $\psi^{c,b}$ be a TC mechanism. Let $\succ$ be a preference profile. We fix an agent $i \in I$. We will show that i cannot benefit by submitting $\succ'_i \neq \succ_i$ while the other agents submit $\succ_{-i}$. Let s be the round i leaves (with house h) at $\succ_i$ and s' be the time i leaves (with h') at $\succ'_i$ in the algorithm. We will consider two cases.

Case 1. $s \leq s'$**:** At round s, the same houses and agents are in the market at both $\succ_i$ and $\succ'_i$ by Lemma 3. If i is not a broker at time s under $\succ_i$, then, by submitting $\succ_i$, agent i gets the top-choice house among the remaining ones in round s, implying that he cannot be better off by submitting $\succ'_i$.

Assume now that i is a broker at time s under $\succ_i$. Let e be the brokered house at time s. If e is not agent i's top-choice house remaining under $\succ_i$, then by submitting $\succ_i$, agent i gets the top-choice house among the remaining ones in round s, implying that he cannot be better off by submitting $\succ'_i$.

It remains for us to consider the situation in which e is broker i's top-choice remaining house, and to show that i cannot get e by submitting the profile $\succ'_i$. For an argument by contradiction, assume that under $\succ'_i$ agent i leaves at round s' with house e. Because agent i is a broker when he leaves at $\succ_i$, there is an agent j who is matched with house e at time s. At this time, j is an owner of some owned house h_j, and e is his top-choice house. By Lemma 3, the control rights structure at round s is the same under both $\succ_i$ and $\succ'_i$. Hence, i is also a broker at time s after submitting $\succ'_i$, and j is an owner of h_j. Moreover, j's top choice is still house e. That means that under $\succ'_i$ agent j will stay unmatched until $s' + 1$. Since agent i leaves with e at s', he cannot be the broker of e at this round, because a broker cannot leave with the brokered house, while another owner j is unmatched. Thus, there is a round $s'' \in \{s + 1, ..., s'\}$ at which agent i stops being the broker of e. Since e is still unmatched at this round, there is a broker-to-heir transition between $s'' - 1$ and s'' (by R6). Because j is an owner of h_j at both $s'' - 1$ and s'', he would have inherited e at s'' (by R6). Then, however, j would have left with e at s'', as e is j's top choice among houses left at s (and hence those left at s''). A contradiction.

Case 2. $s > s'$**:** At round s', the same houses and agents are in the market at both $\succ_i$ and $\succ'_i$ by Lemma 3. Consider round s' at both $\succ_i$ and $\succ'_i$. Under $\succ'_i$, agent i points to house $h' = h^1$ that points to agent i^1 that points to ... that points to object h^n that points to agent $i = i^n$ (and this cycle leaves at round s'). If the cycle is trivial ($n = 1$) and h' points back to i, then i owns h'. Since ownership persists by R4, i will own h' at $s > s'$, and thus at round s, agent i would leave with a house at least as good as h'.

In the sequel, assume that there is at least one other agent i^n in the cycle (that is, $n \geq 2$).

If each house h^ℓ is owned by i^ℓ, for all $\ell \in \{1, ..., n\}$, then the chain $h' = h^1 \to i^1 \to h^2 \to ... \to h^n \to i$ will stay in the system as long as i is in the system (by persistence of ownership, implied through R4). Thus, at round s agent i would leave with a house at least as good as h' under $\succ_i$.

If i^ℓ brokers h^ℓ for some $\ell \in \{1, ..., n\}$, then the chain $h' = h^1 \to i^1 \to h^2 \to ... \to h^n \to i$ will stay in the system as long as i^ℓ continues brokering h^ℓ (since there are no other brokerages and ownerships persist by R4). If i^ℓ brokers h^ℓ at round s under $\succ_i$, then we are done, since the same cycle would have formed. Thus suppose that at a round $s'' \in \{s' + 1, ..., s\}$ broker i^ℓ loses his broker status. Because $n \geq 2$, agent $i^{\ell+1}$ is an owner both at rounds $s'' - 1$ and s''. Hence, the loss of brokerage status means that i^ℓ enters a broker-to-heir transition. We must then have $n = 2$ (since by R6, only 1 previous owner can remain unmatched during the broker-to-heir transition). There are two cases: either i^1 owns $h^1 = h'$ and h^2 (and $i^2 = i^\ell$ is the heir) or $i^2 = i$ owns h^1 and h^2. In the former case, i^1 who wants h^2, will leave with it at round s'' under $\succ_i$, and i will inherit $h^1 = h'$ at $s'' + 1$ by R6. In the latter case, i owns $h^1 = h'$ already at round s''. In both cases, at $s \geq s''$ agent i can only leave with a house at least as good as h' under $\succ_i$. **QED**

Let us finish this section with an overview of the proof of Theorem 2 (the full proof is in the Supplementary Appendix D). In the proof we fix a group strategy-proof and Pareto-efficient direct mechanism φ and construct a TC mechanism $\psi^{c,b}$ that is equivalent to φ. We proceed in three steps: we first construct the candidate control rights structure (c, b), then show it satisfies conditions R1-R6, and finally show that the resultant TC mechanism $\psi^{c,b}$ equals φ.

We define a candidate control rights structure in terms of how φ allocates objects for preferences from some special preference classes. To see how this is done, consider the empty submatching and a house h. If φ were a TC and h was owned by an agent then at all preference profiles in which all agents rank h as their most preferred house, φ would allocate h to the same agent – the owner of h at the empty submatching. We thus check whether φ allocates h to the same agent at all above profiles, and if it does, we call this agent the candidate owner of h (in the proof, for brevity, we refer to the candidate owner as owner*). If φ does not allocate h to the same agent at all above profiles, h is a candidate brokered house. Notice that if φ were a TC and h was brokered by an agent then at every profile at which every agent ranks h as his most preferred house and some other house h' as his second-most preferred house, φ would allocate h' to the same agent – the broker of h at the empty submatching. We thus check whether there is an agent who always gets his second-most preferred house at the above profiles, and if there is such an agent we call this agent the candidate broker of h (broker* for short). Finally, we prove the key result that every house h either has a candidate owner or a candidate broker.

The construction of candidate control rights at non-empty submatchings is similar. The only modification is that instead of looking at preferences at which all agents agree on their most preferred house (or two most preferred houses), we impose this commonality only on unmatched agents, and at the same time assume the matched agents rank the houses they are matched with at the top, while all other agents rank matched houses at the bottom. Thanks to the simplifying assumption that $|H| \geq |I|$, the Pareto efficiency of TC mechanisms implies that the above procedure would work well if φ was a TC, and we prove that indeed it works well whenever φ is group strategy-proof and efficient.[15]

The second step of the proof is to show that the above candidate control rights structure indeed satisfies properties R1-R6. We flesh out the argument in several lemmas. With these lemmas proven, we have constructed a TC mechanism $\psi^{c,b}$. The last step of the proof is to show that $\psi^{c,b} = \varphi$. We rely on the recursive structure of TC, and proceed by induction with respect to the rounds of $\psi^{c,b}$.

6 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.

6.1 Model of House Allocation and Exchange

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 $\mathcal{H} = \{H, \emptyset, ..., \emptyset\}$. We may fix H, I and $\mathcal{H}$, and identify the house allocation and

[15] This point in the construction requires more care in the case of $|H| < |I|$; see Section 7.

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

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

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

6.2 Results

Our main characterization result for house allocation and exchange is now an immediate corollary of Theorems 1-2.

Theorem 3. *In house allocation and exchange problems, a mechanism is* individually rational*,* Pareto efficient*, and* group strategy-proof *if and only if it is an individually rational TC mechanism.*

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.*

Proof of Proposition 2. To prove individual rationality of the above subclass of TC mechanisms, consider an agent i and assume that i owns at the empty submatching a house h from his endowment. Then R4 ensures that i owns h throughout the execution of the TC algorithm. Thus, the TC mechanism will allocate to i house h or a house that i prefers to h. Now, let ψ be an individually rational TC mechanism. Recall that ownership* was defined in the proof of Theorem 2. For any agent i and house h from i's endowment, i is owner* of h because individual rationality implies that $\psi[\succ](i) = h$ for any $\succ \in \mathbf{P}[\varnothing, h]$, which is the set of preference profiles that rank h first for all agents. The construction from the proof of Theorem 2 thus yields a control rights structure that assigns to each agent the initial ownership rights over the houses from his endowment, and represents ψ. **QED**

Notice that when one agent is endowed with all houses, there are individually rational mechanisms that might be represented both by a control rights structure that assigns this agent initial ownership rights over all houses, and by an alternative control rights structure that assigns this agent ownership rights over all houses but one. Except for such situations, however, any control rights structure of an individually rational TC mechanism assigns to each agent the initial ownership rights of all houses from his endowment.

As a corollary of the above two results, we obtain a powerful and non-trivial characterization for an important subdomain of allocation and exchange problems:

Theorem 4. *In house allocation and exchange problems where each agent has a nonempty endowment, a mechanism is* individually rational*,* Pareto efficient*, and* group strategy-proof *if and only if it is a TTC mechanism (aka hierarchical exchange) that assigns all agents the initial ownership rights of houses from their endowment.*

Proof of Theorem 4. By Theorem 3, a mechanism φ is individually rational, Pareto efficient and group strategy-proof if and only if there exists an individually rational and consistent control rights structure (c, b) such that $\varphi = \psi^{c,b}$. By Proposition 2 we may assume that each agent has initial ownership rights over the houses from their endowment. By condition R4 of consistency all unmatched agents own a house throughout the mechanism, and hence R3 implies that no agent is a broker. $\psi^{c,b}$ is thus a TTC mechanism. **QED**

This result is a generalization of the result stated by Ma (1994) for the housing market of Shapley and Scarf (1974). A housing market is a house allocation and exchange problem in which $|I| = |H|$ and each agent is endowed with a house. In this environment, Ma characterized TTC (in which agents own their endowments) as the unique mechanism that is individually rational, strategy-proof, and Pareto efficient.

6.3 Market Design Environments

The assumptions of Theorem 3 are satisfied by the *house allocation problem with existing tenants* of Abdulkadiroğlu and Sönmez (1999). Theirs is the subclass of house allocation and exchange problems in which each agent is endowed with one or zero house. In the former case, the agent is referred to as an *existing tenant*. The house allocation problem with existing tenants is modeled after dormitory assignment problems in US college campuses. In each such college, at the beginning of the academic year, there are new senior, junior, and sophomore students, each of whom already occupies a room from the last academic year. There are vacated rooms by the graduating class and there are new freshmen who would like to obtain a room, though they do not currently occupy any.

The assumptions of Theorem 4 are satisfied by the *kidney exchange* with strict preferences (Roth, Sönmez, and Ünver, 2004), and the *kidney exchange problem with good Samaritan donors* (Sönmez and Ünver, 2006). Kidney transplant patients are the agents and live kidney donors are the houses. Each agent is endowed with a live donor who would like to donate a kidney if his paired-donor receives a transplant in return. Thus, all agents have nonempty endowments. The model also allows for unattached donors known as good Samaritan donors who would like to donate a kidney to any patient. In the US, good Samaritan donors have been the driving force behind kidney exchange since 2006. Many 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) have used good Samaritan donors in many of kidney exchanges conducted since 2006 (cf. Rees, Kopke, Pelletier, Segev, Rutter, Fabrega, Rogers, Pankewycz, Hiller, Roth, Sandholm, Ünver, and Montgomery, 2009).

The kidney exchange context underscores the importance of group strategy-proofness. The doctors of patients are the ones who have the information about patients' preferences over kidneys and it is known that doctors (or transplant centers) themselves at times manipulate the system to benefit their patients.[16] An individually strategy-proof mechanism that is not group strategy-proof could thus by manipulated by doctors. Group strategy-proofness guarantees that no doctor is able to manipulate the mechanism on behalf of his or her patients without harming at least one of them.

7 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

[16]Deceased-donor queue procedures are sometimes gamed by physicians acting as advocates for their patients. In particular, in 2003 two Chicago hospitals settled a federal lawsuit alleging that some patients had been fraudulently certified as sicker than they were to move them up on the liver transplant queue (Warmbir, 2003).

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 6, 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 $\sigma : J \to H \cup Y$ 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, $H_\sigma \subseteq H$ denotes the set of houses matched by 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 agent 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 $\varphi : \mathbf{P} \longrightarrow \mathcal{M}$ 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.

In the presence of outside options, the TC class of mechanisms again coincides with the class of Pareto-efficient and group strategy-proof direct mechanisms. The proof resembles the proofs of Theorems 1 and 2; the required modifications are discussed in Supplementary Appendix E.

Theorem 5. *In the environment with outside options, every TC mechanism is group strategy-proof and Pareto efficient. Moreover, every group strategy-proof and Pareto-efficient direct mechanism is TC.*

We are now ready to extend the characterization to the general allocation and exchange setting with outside options. As in Section 6, 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 6 translate to the setting with outside options; the proofs rely on Theorem 5 instead of Theorems 1 and 2, and are otherwise unchanged.

Theorem 6. *In house allocation and exchange problems with outside options, a mechanism is* individually rational*,* Pareto efficient*, and* group strategy-proof *if and only if it is an individually rational TC mechanism.*

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. Our characterization of TTC remains correct.

Theorem 7. *In house allocation and exchange problems with outside options, if each agent has a nonempty endowment, then a mechanism is* individually rational*,* Pareto efficient*, and* group strategy-proof *if and only if it is a TTC mechanism (with outside options) that assigns all agents the initial ownership rights of houses from their endowment.*

A Appendix: Comments on Consistency Requirements

This appendix explains the consistency requirements R3, R5, and R6. The remaining requirements, R1, R2, and R4, were discussed in the main text.

R3 postulates that a broker does not own any houses. Dropping this assumption would violate efficiency. For instance, consider the case of two agents 1 and 2 such that agent 1 brokers house h_1 and owns house h_2 while 2 has no control rights. If agent 1 prefers h_1 over h_2 while agent 2 prefers h_2 over h_1 then running the TC algorithm (with the above inconsistent control rights structure) would allocate h_2 to agent 1 and h_1 to agent 2, which is inefficient.

R5 might be called limited persistence of brokerage, and is the counterpart of R4 for brokers. R5 states that a brokerage right persists when we move from smaller to larger submatching provided two or more owners from the smaller submatching remain unmatched at the larger submatching. The following example illustrates why we need this requirement to keep TC individually strategy-proof:

Example 4. ***Why do we need R5 to prevent individual manipulation?*** Consider four agents $i_1, ..., i_4$. Assume that at the empty submatching agent i_2 brokers a house and other agents own one house each. Denote by h_k the house controlled by agent i_k. Let us maintain R1-R3, R4, and R6, and violate R5 by assuming that h_2 is owned by i_4 at submatching $\{(i_1, h_1)\}$. Now, there are two previous owners unmatched at $\{(i_1, h_1)\}$, i_3 and i_4. Moreover, i_2 is no longer a broker. Consider now a preference profile such that h_1 is i_1's and i_2's mutual first-choice house, h_2 is the first choice of the other agents, and h_3 is the second choice of i_2 and i_3. Under this preference profile and control rights structure, i_2 would benefit by misrepresenting his preferences and declaring h_3 to be his first choice.

R6 refers to the case where a broker loses his right at a submatching at which only a single previous owner is unmatched. In this case, the broker requires some protection against losing his right. That is to say, when the previous owner gets matched with the ex-brokered house, the ex-broker owns the houses of this owner. This is the *broker-to-heir transition* of the ex-broker.

The following two examples illustrate why we need R6 to keep TC both individually strategy-proof and non-bossy. The first one is similar to the above one:

Example 5. ***Why do we need R6 to prevent individual manipulation?*** Consider four agents $i_1, ..., i_4$. Assume that at the empty submatching agent i_2 brokers h_2, i_1 owns h_1, h_4, and i_3 owns h_3. At submatching $\{(i_1, h_1)\}$, assume that i_3 owns h_2 as well, and i_2 loses his brokerage right. Now, i_4 inherits h_4 as an owner. We assume R1-R5, and violate R6. R5 is not violated, as there is a single previous owner unmatched at $\{(i_1, h_1)\}$, and he is i_3. However, R6 is violated, as at the submatching $\{(i_1, h_1), (i_3, h_2\}$, i_2 is not the heir to i_3. That is, i_2 does not own the ex-owned house h_3 of i_3, but i_4 does. Consider the preference profile at which agents i_1 and i_2 have house h_1, i_2 has h_2 and i_4 has h_3 as their first choices; and agent i_2's second choice is h_3. Then, i_2 would benefit by ranking h_3 first.

Example 6. ***Why do we need R6 to prevent bossiness?*** Consider the same control rights structure as in Example 5. Consider the preference profile at which i_1 and i_3's first choices are h_1, and i_2 and i_4's first choice is h_3; second choice of i_3 is h_2. Agent i_3 will be bossy by ranking h_2 first. In both cases he receives house h_2. However, in the first case, i_2 receives h_4, while in the latter, he receives h_3.

B Appendix: Further Illustrative Applications

In house allocation problems, we characterized the set of group strategy-proof and efficient mechanisms through a new class, *trading cycles.* The results in Sections 6 and 7 used the characterization of Theorems 1 and 2 to draw new characterizations in environments with outside options and private property rights, respectively. In this appendix we give three further illustrative examples of how our main results radically simplify the analysis of allocation and exchange problems.

B.1 Neutrality

Neutrality and group strategy-proofness were characterized through serial dictatorships by Svensson (1999) when there are no outside options. In a serial dictatorship agents are ordered, first agent in the ordering gets his most preferred house, the second agent in the ordering gets her most preferred among houses unassigned to agents higher in the ordering, etc. Svensson's result is implied by Theorem 2as illustrated below.

A mechanism is neutral if whenever the house names are relabeled in the problem, the mechanism outcome assigns agents the house that carries the relabeled name of the house that was assigned in the original problem. Formally, a relabeling of houses is a bijection $\pi : H \to H$. For any preference profile $\succ \in \mathbf{P}$, and relabeling π, let $\succ^{\pi} \in \mathbf{P}$ be such that $g \succ_i^{\pi} h \Leftrightarrow \pi^{-1}(g) \succ_i \pi^{-1}(h)$ for all $i \in I$ and $g, h \in H$. A mechanism φ is **neutral** if for all relabelings π, all $\succ \in \mathbf{P}$, and all $i \in I$, we have $\varphi[\succ^{\pi}](i) = \pi(\varphi[\succ](i))$.

Corollary 1. *A mechanism is group strategy-proof and neutral if and only if it is a serial dictatorship.*

.

Proof of Corollary 1. Let φ be a group strategy-proof and neutral mechanism. Neutrality implies that ϕ has full range that is $\phi[\mathbf{P}] = \mathcal{M}$. Indeed, for any $\mu \in \mathcal{M}$, we can take an arbitrary $\succ \in \mathbf{P}$, define relabeling π so that $\pi(\varphi[\succ](i)) = \mu(i)$ for all $i \in I$, and conclude from neutrality that $\varphi[\succ^\pi] = \mu$. As observed in Section 2, full range and group strategy-proofness imply Pareto efficiency. Thus, φ is a trading cycles mechanism $\psi^{c,b}$ by Theorem 2.

It remains to show that any neutral trading cycles mechanism $\psi^{c,b}$ is equivalent to a serial dictatorship. Let $\sigma \in \overline{\mathcal{M}}$. By R1 and R3 there is an agent $i \in \bar{I}_\sigma$ who owns some house $h \in \overline{H_\sigma}$ at σ. In particular, for any $\succ \in \mathbf{P}[\sigma; h]$, $\psi^{c,b}[\succ](i) = h$. Let $\sigma' \in \overline{\mathcal{M}}$ with $I_{\sigma'} = I_\sigma$. Let $g \in \overline{H_{\sigma'}}$. Take a relabeling π such that $\pi(h) = g$ and $\pi(\sigma(j)) = \sigma'(j)$ for all $j \in I_\sigma$. Now, $\succ^\pi \in \mathbf{P}[\sigma'; g]$ and by neutrality $\psi^{c,b}[\succ^\pi](i) = \pi(h) = g$. Maskin monotonicity implies that i is allocated the best unmatched house at σ' as long as $I_{\sigma'} = I_\sigma$. The mechanism$\psi^{c,b}$ is thus equivalent to a serial dictatorship. **QED**

B.2 Reallocation-proofness

Reallocation-proofness, group strategy-proofness, and efficiency were characterized by Pápai (2000) through hierarchical exchange (i.e., TTC mechanisms). A mechanism φ is **reallocation-proof if** there exists no pair of agents $i, j \in I$ such that for some $\succ \in \mathbf{P}$, $\succ'_i \in \mathbf{P}_i$, and $\succ'_j \in \mathbf{P}_j$ with $\varphi[\succ'_i, \succ_{-i}] = \varphi[\succ'_j, \succ_{-j}] = \varphi[\succ]$, we have $\varphi[\succ'_{\{i,j\}}, \succ_{-\{i,j\}}](j) \succ_i \varphi[\succ](i)$ and $\varphi[\succ'_{\{i,j\}}, \succ_{-\{i,j\}}](i) \succsim_j \varphi[\succ](j)$. We can derive the key insight of Pápai (2000) as a corollary of 2:

Corollary 2. *If a mechanism is group strategy-proof, Pareto efficient, and reallocation-proof then it is a hierarchical exchange mechanism.*

Proof of Corollary 2. Let φ be a group strategy-proof, efficient, and reallocation-proof mechanism. By Theorem 2, it is equivalent to a reallocation-proof TC mechanism $\psi^{c,b}$. It remains to show that the control right structure (c, b) can be chosen in such a way that there are no brokers. Take any submatching $\sigma \in \overline{\mathcal{M}}$. First notice that if there are two owners, j and k at σ then no house is σ-brokered. Indeed, by way of contradiction assume that some house h is σ-brokered by an agent i, and let h_j be a house owned by j and h_k be a house owned by k. Consider a preference profile $\succ \in \mathbf{P}[\sigma]$ and such that $\succ_i \in \mathbf{P}_i[\sigma; h, h_k]$, $\succ_j \in \mathbf{P}_j[\sigma; h_j]$, and $\succ_k \in \mathbf{P}_k[\sigma; h]$. Then, the deviation to $\succ'_i \in \mathbf{P}_i[\sigma; h, h_j, h_k]$ and $\succ'_j \in \mathbf{P}_j[\sigma; h, h_j]$ violates the reallocation-proofness condition. Hence, (c, b) can allow brokers only at submatchings with a unique owner. But then $\psi^{c,b}$ is equivalent to $\psi^{c',b'}$ such that (c', b') is identical to (c, b) except that at any submatching σ at which (c, b) gives brokerage right over a house h to an agent i, the primed control right structure (c', b') gives ownership of h to the unique (c, b) owner j at σ, and gives i the ownership of all unmatched houses at $\sigma \cup \{(j, h)\}$. **QED**

B.3 Invariance

Our main results allows one to easily find and prove new characterization results. As an example let us propose an alternative characterization of hierarchical exchange of Pápai (2000). A mechanism is **invariant** if for any agent $i \in I$ and any object $h \in H$ if for all $g \succsim_i h$

$$\phi\left(\succ_i, \succ_{-i}\right)(i) = g \Longleftrightarrow \phi\left(\succ_i', \succ_{-i}\right)(i) = g$$

then for all $g \succsim_i h$ and all $j \in I$

$$\phi\left(\succ_i, \succ_{-i}\right)(j) = g \Longleftrightarrow \phi\left(\succ_i', \succ_{-i}\right)(j) = g.$$

We then get the following corollary of Theorems 1 and 2,

Corollary 3. *A mechanism is strategy-proof, efficient, and invariant if, and only if, it is a hierarchical exchange mechanism.*

For the proof, notice that invariance implies non-bossiness, and hence the class of mechanisms in the lemma is group strategy-proof. The rest of the proof resembles the proof of Corollary 2.

References

Abdulkadiroğlu, A., and Y.-K. Che (2010): "The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles," Working paper.

Abdulkadiroğlu, A., Y.-K. Che, and Y. Yasuda (2011): "Expanding Choice in School Choice," Working paper.

Abdulkadiroğlu, A., and T. Sönmez (1999): "House Allocation with Existing Tenants," *Journal of Economic Theory*, 88, 233–260.

——— (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Arrow, K. J. (1950): "A Difficulty in the Concept of Social Welfare," *Journal of Political Economy*, 58, 328–346.

Ashlagi, I., and A. E. Roth (2011): "Individual rationality and participation in large scale, multi-hospital kidney exchange," Working paper.

Balinski, M., and T. Sönmez (1999): "A Tale of Two Mechanisms: Student Placement," *Journal of Economic Theory*, 84, 73–94.

Barberà, S., F. Gül, and E. Stacchetti (1993): "Generalized Median Voter Schemes and Committees," *Journal of Economic Theory*, 61, 262–289.

Barberà, S., and M. O. Jackson (1995): "Strategy-proof Exchange," *Econometrica*, 63, 51–87.

Barberà, S., M. O. Jackson, and A. Neme (1997): "Strategy-proof Allotment Rules," *Games and Economic Behavior*, 18, 1–21.

Bennett, E., and D. Conn (1976): "The group incentive properties of mechanisms for the provision of public goods," *Public Choice*, 29, 95–102.

Bogomolnaia, A., R. Deb, and L. Ehlers (2005): "Strategy-Proof Assignment on the Full Preference Domain," *Journal of Economic Theory*, 123, 161–186.

Clarke, E. H. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 11, 17–33.

Dasgupta, P., P. Hammond, and E. Maskin (1979): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46, 185–216.

Ehlers, L. (2002): "Coalitional Strategy-Proof House Allocation," *Journal of Economic Theory*, 105, 298–317.

Ehlers, L., and B. Klaus (2006): "Efficient Priority Rules," *Games and Economic Behavior*, 55, 372–384.

Erdil, A., and H. Ergin (2008): "What's the matter with tie-breaking? Improving efficiency in school choice," *American Economic Review*, 98, 669–689.

Ergin, H. (2000): "Consistency in House Allocation Problems," *Journal of Mathematical Economics*, 34, 77–97.

——— (2002): "Efficient Resource Allocation on the Basis of Priorities," *Econometrica*, 70, 2489–2498.

Gibbard, A. (1973): "Manipulation of Voting Schemes: A General Result," *Econometrica*, 41, 587–601.

Green, J., and J.-J. Laffont (1977): "Characterization of Satisfactory Mechanisms for Revelation of Preferences for Public Goods," *Econometrica*, 45, 427–438.

——— (1979): "On Coalition Incentive Compatibility," *The Review of Economic Studies*, 46(2), 243–254.

Groves, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617–631.

Hatfield, J. W. (2009): "Strategy-Proof, Efficient, and Nonbossy Quota Allocations," *Social Choice and Welfare*, 33 No. 3, 505–515.

Hylland, A., and R. Zeckhauser (1979): "The Efficient Allocation of Individuals to Positions," *Journal of Political Economy*, 87, 293–314.

Kesten, O. (2006): "On Two Competing Mechanisms for Priority-based Allocation Problems," *Journal of Economic Theory*, 127, 155–171.

Kojima, F., and M. Manea (2010): "Axioms for Deferred Acceptance," *Econometrica*, 78, 633–653.

Ma, J. (1994): "Strategy-proofness and the strict core in a market with indivisibilities," *International Journal of Game Theory*, 23, 75–83.

Maskin, E. (1999): "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23–38.

Moulin, H. (1980): "On Strategy-proofness and Single Peakedness," *Public Choice*, 35, 437–455.

Pápai, S. (2000): "Strategyproof Assignment by Hierarchical Exchange," *Econometrica*, 68, 1403–1433.

——— (2007): "Exchange in a General Market with Indivisible Goods," *Journal of Economic Theory*, 132, 208–235.

Pathak, P. A., and T. Sönmez (2008): "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism," *American Economic Review*, 98, 1636–1652.

Rees, M. A., J. E. Kopke, R. P. Pelletier, D. L. Segev, M. E. Rutter, A. J. Fabrega, J. Rogers, O. G. Pankewycz, J. Hiller, A. E. Roth, T. Sandholm, M. U. Ünver, and R. A. Montgomery (2009): "A Non-simultaneous Extended Altruistic Donor Chain," *The New England Journal of Medicine*, 360, 1096–1101.

Roberts, K. (1979): "Welfare Considerations of Nonlinear Pricing," *Economic Journal*, 89, 66–83.

Roth, A. E. (1982): "Incentive Compatibility in a Market with Indivisibilities," *Economics Letters*, 9, 127–132.

Roth, A. E., and A. Postlewaite (1977): "Weak Versus Strong Domination in a Market with Indivisible Goods," *Journal of Mathematical Economics*, 4, 131–137.

Roth, A. E., T. Sönmez, and M. U. Ünver (2004): "Kidney Exchange," *Quarterly Journal of Economics*, 119, 457–488.

Satterthwaite, M. (1975): "Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions," *Journal of Economic Theory*, 10, 187–216.

Satterthwaite, M., and H. Sonnenschein (1981): "Strategy-Proof Allocation Mechanisms at Differentiable Points," *Review of Economic Studies*, 48, 587–597.

Shapley, L., and H. Scarf (1974): "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1, 23–37.

Sönmez, T. (1999): "Strategy-Proofness and Essentially Single-Valued Cores," *Econometrica*, 67, 677–690.

Sönmez, T., and M. U. Ünver (2006): "Kidney Exchange with Good Samaritan Donors: A Characterization," Working paper.

——— (2010): "Market Design for Kidney Exchange," Zvika Neeman, Muriel Niederle, and Nir Vulkan (eds.) *Oxford Handbook of Market Design*, forthcoming, OUP.

Sprumont, Y. (1991): "The Division Problem with Single-Peaked Preferences : A Characterization of the Uniform Allocation Rule," *Econometrica*, 59, 509–519.

Svensson, L.-G. (1994): "Queue Allocation of Indivisible Goods," *Social Choice and Welfare*, 11, 223–230.

——— (1999): "Strategy-proof Allocation of Indivisible Goods," *Social Choice and Welfare*, 16, 557–567.

Takamiya, K. (2001): "Coalition Strategy-Proofness and Monotonicity in Shapley-Scarf Housing Markets," *Mathematical Social Sciences*, 41, 201–213.

Vickrey, W. (1961): "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Warmbir, S. (2003): "UIC Hospital Sued for Medicare Fraud," *Chicago Sun Times*, July 29.

Zhou, L. (1991): "Inefficiency of Strategy-Proof Allocation Mechanisms in Pure Exchange Economies," *Social Choice and Welfare*, 8, 247–257.

Supplement to "Incentive Compatible Allocation and Exchange of Discrete Resources"

by
Marek Pycia and M. Utku Ünver

C Supplementary Appendix: Proof of Theorem 1

Proof of Theorem 1. In the main text we showed that TC is Pareto efficient and individually strategy-proof. By Papai's Lemma 1, it is sufficient to show that TC is non-bossy. Let $\psi^{c,b}$ be a TC mechanism. Fix an agent $i_* \in I$ and two preference profiles $\succ = [\succ_{i_*}, \succ_{-i_*}]$ and $\succ' = [\succ'_{i_*}, \succ_{-i_*}]$ such that

$$h_* = \psi^{c,b}[\succ'](i_*) = \psi^{c,b}[\succ](i_*).$$

Let s be the round i_* leaves (with house h_*) submitting $\succ_{i_*}$ and s' be the time i_* leaves (with h_*) submitting $\succ'_{i_*}$. By symmetry, it is enough to consider the case $s \leq s'$. In order to show that

$$\psi^{c,b}[\succ](i) = \psi^{c,b}[\succ'](i) \qquad \forall\, i \in I,$$

we will prove the following stronger statement:

Hypothesis: If a cycle $h^1 \to i^1 \to h^2 \to ... \to h^n \to i^n \to h^1$ of length $n \in \{1, 2, ...\}$ forms and is removed at round r under preference profile $\succ$, then under preference profile $\succ'$ one of the following three (non-exclusive) cases obtains:

1. the same cycle $h^1 \to i^1 \to h^2 \to ... \to h^n \to i^n \to h^1$ forms; or

2. $n = 2$ and two cycles form:

   - cycle $h^1 \to i^2 \to h^1$ or cycle $g \to i^2 \to h^1 \to i \to g$ for some agent i and some house g, and
   - cycle $h^2 \to i^1 \to h^2$ or cycle $h \to i^1 \to h^2 \to j \to h$ for some agent j and some house h;

   or

3. $n = 1$ and there exists an agent $j \neq i^1$ and a house $h \neq h^1$ such that the cycle $h \to i^1 \to h^1 \to j \to h$ forms.

Whenever in the proof we encounter cycles of length n, the superscripts on houses and agents will be understood to be modulo n, that is $i^{n+1} = i^1$ and $h^{n+1} = h^1$. By $\sigma^{s-1}[\succ]$ we denote the submatching of agents and houses matched before round s of $\psi^{c,b}$ when agents submitted preference profile $\succ$. We refer to cycles formed under $\succ$ as $\succ$-cycles, and to cycles formed under $\succ'$ as $\succ'$-cycles.

By Lemma 3, the above hypothesis is true for all $r < s$. The proof for $r \geq s$ proceeds by induction over the round r.

Initial step. Consider $r = s$. Under $\succ$, house h_*^1 points to agent $i_* = i_*^1$ points to house $h_* = h_*^2$ that points to agent i_*^2 that points to ... that agent i_*^n that points to house h_*^1, and the cycle

$$h_*^1 \to i_*^1 \to h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1$$

is removed at round s. Lemma 3 implies that the same houses and agents are in the market at time s under both $\succ$ and $\succ'$ and that all agents from $I_{\sigma^s[\succ]} - \{i_*^1, ..., i_*^n\}$ are matched by $\sigma^s[\succ']$ in the same way as in $\sigma^s[\succ]$. Lemma 3 also implies that the chain $h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1 \to i_*^1$ forms at round s under preferences $\succ'$.

If all pairs (i_*^ℓ, h_*^ℓ), for all $\ell \in \{2, ..., n\}$, consist of an owner and an owned house at $\sigma^s[\succ]$, then they consist of an owner and an owned house at $\sigma^s[\succ']$ and the chain $h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1 \to i_*^1$ will stay in the system as long as i_*^1 is in the system (by R4). Thus, at s' all agents $i_*^1, ..., i_*^n$ would leave in the same cycle as under $\succ$. Notice that this argument fully covers the case $n = 1$.

If $n > 1$ and i_*^ℓ brokers h_*^ℓ for some $\ell \in \{2, ..., n\}$, then the chain $h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1 \to i_*^1$ will stay in the system as long as i_*^ℓ continues to broker h_*^ℓ. If i_*^ℓ continues to broker h_*^ℓ until round s' under $\succsim'$, then the initial step is proved. Otherwise, there is a round $s'' \in \{s+1, ..., s'\}$ such that agent i_*^ℓ has the brokerage right over h_*^ℓ at rounds $s, ..., s''-1$ but not at round s''. By R6's broker-to-heir transition property, $n = 2$ and $i_*^{\ell+1}$ owns h_*^ℓ at $\sigma^{s''}[\succ']$ because he owns $h_*^{\ell+1}$at both $\sigma^{s''-1}[\succ']$ and $\sigma^{s''}[\succ']$. As $i_*^{\ell+1}$'s top preference is then h_*^ℓ, he will leave with it at s''. By R6's broker-to-heir transition property, agent i_*^ℓ will inherit $h_*^{\ell+1}$ at $s''+1$ and will be matched with it. This case ends the proof of the inductive hypothesis for $r = s$.

Inductive step. Now, take any round $r > s$ such that $\sigma^r[\succ] - \sigma^{r-1}[\succ]$ is non-empty, and assume that the inductive hypothesis is true for all rounds up to $r-1$. Consider agents and houses

$$h^1 \to i^1 \to h^2 \to ... \to h^n \to i^n \to h^1$$

that form a cycle of length $n \geq 1$ at round r under $\succ$. Since all agents but i^* (who is matched before round r) have same preferences in both profiles $\succ$ and $\succ'$, so do agents $i^1, ..., i^n$. We start with two preparatory claims.

Claim 1. If agent j and house h are unmatched at submatchings σ, σ', and j controls h at σ but not at $\sigma \cup \sigma'$, then j brokers h at σ. If, additionally, agent j' and house h' are unmatched at submatchings σ, σ', and, at σ', agent j controls h and agent j' owns h', then $j \neq j'$, j' owns h and h' at $\sigma \cup \sigma'$, and j brokers h at σ' and owns h' at $\sigma \cup \sigma' \cup \{(j', h)\}$.

Notice that $j \neq j'$ in the claim, but we allow $h = h'$.

Proof of Claim 1: The first statement follows from R4. To prove the second statement, first notice that R4 implies that j brokers h at σ', and hence $j \neq j'$. R4 furthermore implies that j' owns h' at all submatchings between σ' and $\sigma \cup \sigma'$. Since j stops brokering h at a submatching between σ' and $\sigma \cup \sigma'$, assumption R6 implies that j' owns h at $\sigma \cup \sigma'$, and j owns h' at $\sigma \cup \sigma' \cup \{(j', h)\}$. QED

Claim 2. Under $\succ'$, all houses i^ℓ prefers over $h^{\ell+1}$, except possibly h^ℓ, are matched with agents other than i^ℓ. If i^ℓ is a $\sigma^{r-1}(\succ)$-owner then there is no exception: under $\succ'$, all houses i^ℓ prefers over $h^{\ell+1}$ are matched with agents other than i^ℓ.

Proof of Claim 2: Consider the run of algorithm under $\succ'$. If i^ℓ is $\sigma^{r-1}[\succ]$-owner then all houses i^ℓ prefers over $h^{\ell+1}$ are matched before round r under $\succ$. The inductive assumption thus implies that they are also matched with agents other than i^ℓ under $\succ'$. Similarly, if i^ℓ is $\sigma^{r-1}[\succ]$-broker then all houses i^ℓ prefers over $h^{\ell+1}$, except possibly h^ℓ, are matched before round r under $\succ$, and the inductive assumption yields the claim. QED

To conclude the proof, let us introduce the following notation:
t is the earliest round one of the houses $h^1, ..., h^n$ is matched under $\succ'$;
$h^{\ell+1}$ is a house matched in round t under $\succ'$,
$j^{\ell+1}$ is the agent controlling house $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$, and
$\nu = \sigma^{r-1}[\succ] \cup \sigma^{t-1}[\succ']$.

If $j^{\ell+1} = i^{\ell+1}$, then agent $i^{\ell+1}$ controls $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$. Three cases are possible:

- If $n = 1$, then $i^{\ell+1}$ owns $h^{\ell+1}$ at $\sigma^{r-1}[\succ]$, and by R4 at ν, as well. First, let us now show that $i^{\ell+1}$ cannot broker $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$. If he does then there exists some agent j that owns a house h at $\sigma^{t-1}[\succ']$ so that $h \to j \to h^{\ell+1} \to i^{\ell+1}$ is part of the cycle occurring in round t under $\succ'$. Moreover, $i^{\ell+1}$ loses brokerage of $h^{\ell+1}$ between $\sigma^{t-1}[\succ']$ and ν, then by R6, j owns $h^{\ell+1}$ at ν contradicting $i^{\ell+1}$ owns $h^{\ell+1}$ at ν. Thus, $i^{\ell+1}$ owns $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$, and, by Claim 2, he points to $h^{\ell+1}$ and is matched with it. The inductive hypothesis is correct for the cycle.

- If $n \geq 2$ and $i^{\ell+1}$ prefers $h^{\ell+2}$ over $h^{\ell+1}$, then by Claim 2, he points to it (he cannot broker it since he controls $h^{\ell+1}$), and house $h^{\ell+1}$ is matched in a cycle that contains $h^{\ell+1} \to i^{\ell+1} \to h^{\ell+2} \to ....$

- If $n \geq 2$ and $i^{\ell+1}$ prefers $h^{\ell+1}$ over $h^{\ell+2}$, then $i^{\ell+1}$ is the broker of $h^{\ell+1}$ at $\sigma^{r-1}[\succ]$. First, let us show that $i^{\ell+1}$ cannot own own $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$. If he does then he owns it at ν, as well. Then $i^{\ell+1}$ loses brokerage of $h^{\ell+1}$ between $\sigma^{r-1}[\succ]$ and ν. R5 implies that there is at most one other $\sigma^{r-1}[\succ]$-owner left unmatched at ν, and i^{ℓ} is that owner. R6 implies that i^{ℓ} should own $h^{\ell+1}$ at ν, a contradiction. Hence, $i^{\ell+1}$ brokers $h^{\ell+1}$ at $\sigma^{t-1}[\succ']$. By Claim 2, $i^{\ell+1}$ points to $h^{\ell+2}$ (as he cannot point to $h^{\ell+1}$), and house $h^{\ell+1}$ is matched in a cycle that contains $h^{\ell+1} \to i^{\ell+1} \to h^{\ell+2} \to ....$

We can conclude that $j^{\ell+1} \neq i^{\ell+1}$, or the inductive hypothesis is true for the cycle of $i^1, ..., i^n$, or that the cycle of $h^{\ell+1}$ at $\succ'$ contains $h^{\ell+1} \to i^{\ell+1} \to h^{\ell+2} \to ....$

In the last of this three possibilities, let us define $j^{\ell+2}$ to be the agent controlling $h^{\ell+2}$ at $\sigma^{t-1}[\succ']$, and repeat the above procedure for $h^{\ell+2}$. In this way, repeating this procedure, we either show that the cycle

$$h^{\ell+1} \to i^{\ell+1} \to h^{\ell+2} \to ... \to i^{\ell} \to h^{\ell+1}$$

leaves at round t under $\succ'$ (and the inductive hypothesis (part 1) is true for this cycle), or there is k such that $i^{\ell+k} \neq j^{\ell+k}$. In the sequel, we consider the latter case.

Without loss of generality, we may assume that $\ell + k = 1$ (modulo n), $j^1 \neq i^1$ controls h^1 at $\sigma^{t-1}[\succ']$, and h^1 is matched in round t under σ'. By Claim 2, all agents $i^1, ..., i^n$ are unmatched at $\sigma^{t-1}[\succ']$ because all houses $h^1, ..., h^n$ are. Thus, all these agents and houses are unmatched at ν. Consider three cases.

***Case* 1.** *$n = 1$:* Agent i^1 then owns h^1 at $\sigma^{r-1}[\succ]$, and by the inductive assumption, he gets at most house h^1 under $\succ'$. Thus, i^1 is unmatched in round t. Consider two subcases depending on whether j^1 is matched at $\sigma^{r-1}[\succ]$ or not.

- If $j^1 \in I_{\sigma^{r-1}[\succ]}$: Then $j^1 \neq i^1$ and the inductive assumption and $h^1 \notin H_{\sigma^{r-1}[\succ]}$ imply that

* ⋆ there exists house $h \neq h^1$ such that j^1 is matched before round r in a cycle $h \to j^1 \to h$ under $\succ$, and
* ⋆ there exists agent i such that the cycle $h \to i \to h^1 \to j^1 \to h$ is matched at round t under $\succ'$.

Notice that $i \notin I_{\sigma^{r-1}[\succ]}$, as otherwise the inductive assumption implies that i is matched with h^1 under $\succ$, contrary to $h^1 \notin H_{\sigma^{r-1}[\succ]}$. Thus both i and i^1 are unmatched at ν, and Claim 1 implies that either $i = i^1$, or i is a broker of h at $\sigma^{t-1}[\succ']$. In the latter case, R1 implies that j^1 is an owner of h^1 at $\sigma^{t-1}[\succ']$, and R4 implies that j^1 owns h^1 at ν contrary to $j^1 \neq i^1$ and i^1 owning h^1 at $\sigma^{r-1}[\succ]$, and hence at ν. This contradiction shows that $i^1 = i$, and hence that the inductive hypothesis (part 3) is true for i^1.

* If $j^1 \notin I_{\sigma^{r-1}[\succ]}$: Then, agents $j^1 \neq i^1$ are unmatched at the submatching ν, and Claim 1 implies that j^1 is a broker of h^1 at $\sigma^{t-1}[\succ']$. Let $j' \neq j$ be an agent matched in the same cycle as h^1 at t under $\succ'$. Then j' is an owner of a house h' at $\sigma^{t-1}[\succ']$.

  If j' is matched at $\sigma^{r-1}[\succ]$ then the inductive assumption implies that, under $\succ'$, h^1 is matched to i^1 in a two-agent cycle, and hence $j' = i^1$, and is unmatched at $\sigma^{r-1}[\succ]$; a contradiction. An analogous argument shows that h' is unmatched at $\sigma^{r-1}[\succ]$. Thus j' and h' are unmatched at $\sigma^{r-1}[\succ]$, and hence at ν. Then, R4, R5, and R6 imply that j' owns h^1 at ν; and thus $j' = i^1$. Since i^1 points to h^1, they are matched in the cycle $h' \to i^1 \to h^1 \to j \to h'$, and the inductive hypothesis (part 3) is true for i^1.

Case 2. *$n > 1$ and $i^1 \neq j^1$ brokers h^1 at $\sigma^{r-1}[\succ]$*: Then i^2 is the $\sigma^{r-1}[\succ]$-owner of h^2. We show that in this case $n = 2$, and the inductive hypothesis (part 2) holds for both h^1 (and hence i^2) and for h^2 (and hence i^1).

First, consider how h^1 is matched under $\succ'$. Suppose $j^0 \to h^1 \to j^1$ is the part of the cycle of h^1 in round t under $\succ'$. By the inductive assumption, j^0 is unmatched at $\sigma^{r-1}[\succ]$, and one of the two subcases obtains:

(a) all other houses and agents in the cycle of h^1 under $\succ'$ are unmatched at $\sigma^{r-1}[\succ]$, or

(b) the cycle $h^0 \to j^0 \to h^1 \to j^1 \to h^0$ occurs under $\succ'$ and j^1 is matched with h^0 in $\sigma^{r-1}[\succ]$, i.e., $\{(j^1, h^0)\} \subseteq \sigma^{r-1}[\succ]$.

We handle these two subcases separately:

* Subcase (a): Two further subcases are possible depending on whether j^1 brokers h^1 at $\sigma^{t-1}[\succ]$ or not:

- ⋆ Assume j^1 brokers h^1 at $\sigma^{t-1}[\succ']$. Then, there exists owner-owned house pair (j^0, h^0) at $\sigma^{t-1}[\succ']$ such that $h^0 \to j^0 \to h^1 \to j^1$ is part of the cycle of h^1 in round t under $\succ'$. Either, j^1 or i^1 exits brokerage between $\sigma^{t-1}[\succ']$ or $\sigma^{r-1}[\succ]$, respectively, and ν, as both of them cannot broker it at ν. Depending on whether j^1 or i^1 loses brokerage right, by R5 there are only two agents in the cycle of h^1 under $\succ'$ or $\succ$, respectively; moreover, by R6, j^0 or i^n owns h^1 at ν, respectively. However, then neither j^1 nor i^1 can broker h^1 at ν, implying that both lose brokerage rights, and hence, h^1 is owned by both j^0 and i^n at ν. Thus, $j^0 = i^n$, and $n = 2$. We conclude that i^2 is matched with h^1 under $\succ'$ and the cycle he gets matched in has two agents, i.e., $h^0 \to i^2 \to h^1 \to j^1 \to h^0$.
- ⋆ Assume j^1 owns h^1 at $\sigma^{t-1}[\succ']$. Then, R4 implies that j^1 owns h^1 at ν. Moreover, again by R4, i^1 brokers h^1 at $\sigma^{r-1}[\succ]$ and loses his brokerage right between $\sigma^{r-1}[\succ]$ and ν. By R5, there could be at most one $\sigma^{r-1}[\succ]$- owner still not matched at ν. Hence, $n = 2$ and, $i^n = i^2$ is the remaining owner (of house h^2). By R6, at ν, i^2 also owns h^1. Since also j^1 owns h^1 at ν, $j^1 = i^2$. By Claim 2, h^1 or h^2 is the best house that i^2 can get. Since i^2 is an owner at $\sigma^{r-1}[\succ]$ and he points to h^1 rather than h^2, he prefers h^1 over h^2. Therefore, the cycle of h^1 in round t under $\succ'$ is $h^1 \to i^2 \to h^1$.

- Subcase (b): R6 implies that at $\sigma^{t-1}[\succ'] \cup \{(j^1, h^0)\} \subset \nu$, j^0 owns h^1. Then i^1 leaves brokerage of h^1 between $\sigma^{r-1}[\succ]$ and ν. As i^2 owns h^2 at $\sigma^{r-1}[\succ]$, by R5, i^2 is the only previous owner unmatched at the submatching i^1 leaves brokerage. Thus, the cycle of h^1 under $\succ$ includes only two agents i^1 and i^2, i.e., $n = 2$. Moreover, R6 and R4 imply that i^2 owns h^1 at ν. Thus, $i^2 = j^0$, and the cycle of h^1 at $\succ'$ is $h^0 \to i^2 \to h^1 \to j^1 \to h^0$.

Either subcase proves that there are two agents in the cycle of h^1 under $\succ$, i.e., $n = 2$, and the inductive hypothesis (part 2) holds for i^1 and h^2.

Next, consider how h^2 is matched under $\succ'$. Since, at $\sigma^{r-1}(\succ)$, i^1 controls h^1 and i^2 controls h^2, R6 implies that i^1 owns h^2 at $\nu \cup \{(i^2, h^1)\}$. Let t^1 be the round in which i^1 is matched and t^2 be the round in which h^2 is matched under $\succ'$. Since h^1 is matched with i^2 and not i^1 under $\succ'$, Claim 2 implies that $t^1 \geq t^2$. Moreover, $t^2 \geq t$. Suppose $j' \to h^2 \to j^2$ is part of the cycle of h^2 in round t^2 under $\succ'$. Let

$$\nu^2 = \sigma^{r-1}[\succ] \cup \sigma^{t^2-1}[\succ'] \cup \{(i^2, h^1)\}.$$

We have $\nu \subseteq \nu^2$. Thus, by R4, i^1 owns h^2 at ν^2. We consider two subcases: $i^1 = j^2$ and $i^1 \neq j^2$.

- Assume $i^1 = j^2$. First consider the case i^1 brokers h^2 at $\sigma^{t^2-1}[\succ']$. Then $i^1 = j^2 \neq j'$, and j' is an owner at $\sigma^{t^2-1}[\succ']$. Agent j' is not matched at $\sigma^{r-1}[\succ]$, as otherwise the inductive assumption would imply that h^2 is matched at $\sigma^{r-1}[\succ]$, a contradiction. Hence, j' is not matched at ν^2. Moreover, i^1 loses brokerage right of h^2 between $\sigma^{t^2-1}[\succ']$ and ν^2, as he owns it at ν^2. By R6, j' owns h^2 at ν^2, contradicting i^1 owns it at ν^2. We can conclude that i^1 owns h^2 at $\sigma^{t^2-1}[\succ']$. Since h^1 is matched with i^2 and not i^1 under $\succ'$, Claim 2 implies that $h^2 \to i^1 \to h^2$ is the cycle under $\succ'$, showing that the inductive hypothesis (part 2) holds true for i^1 and h^2.

- Assume $i^1 \neq j^2$. By the inductive assumption, j' is unmatched at ν^2, and either

  (a) all other houses and agents in the cycle of h^2 under $\succ'$ are unmatched at $\sigma^{r-1}[\succ]$, or

  (b) the cycle $h' \to j' \to h^2 \to j^2 \to h'$ occurs under $\succ'$ and j^2 is matched with h' in $\sigma^{r-1}[\succ]$, i.e., $\{(j^2, h')\} \subseteq \sigma^{r-1}[\succ]$.

  ⋆ Subcase (a): As i^1 owns h^2 at ν^2, R4 implies that j^2 is the broker of h^2 at $\sigma^{t^2-1}[\succ]$, and he loses this brokerage right between $\sigma^{t^2-1}[\succ']$ and ν^2. Hence, $j' \neq j^2$ and by R5, there are no other agents than j^2 and j' in the cycle of h^2 under $\succ'$. By R6 and R4 j' owns h^2 at ν^2. Thus, $j' = i^1$. Hence, the cycle of h^2 under $\succ'$ is $h' \to i^1 \to h^2 \to j^2 \to h'$.

  ⋆ Subcase (b): R6 implies that at $\sigma^{t^2-1}[\succ'] \cup \{(j^2, h')\} \subset \nu^2$, j' owns h^2. By R4, j' owns h^2 at ν^2. Recall that i^1 owns h^2 at ν^2. Then $i^1 = j'$, and the cycle of h^2 under $\succ'$ is $h' \to i^1 \to h^2 \to j^2 \to h'$.

  Either subcase proves that the inductive hypothesis (part 2) holds for i^1 and h^2.

Case 3. *$n > 1$ and i^1 owns h^1 at $\sigma^{r-1}[\succ]$* (in particular, R4 implies that i^1 owns h^1 at ν). We will show that this case cannot happen. By the inductive assumption either (a) all agents in the $\succ'$-cycle of h^1 are unmatched at $\sigma^{r-1}[\succ]$ or (b) $n = 2$ and the cycle $h^0 \to j^0 \to h^1 \to j^1 \to h^0$ occurs under $\succ'$ for some house h^0, and j^1 is matched with h^0 in $\sigma^{r-1}[\succ]$, that is $\{(j^1, h^0)\} \subseteq \sigma^{r-1}[\succ]$.

- Subcase (a): By R4, agent j^1 is the broker of h^1 at $\sigma^{t-1}[\succ']$ and loses this right between $\sigma^{t-1}[\succ']$ and ν. Hence, $j^0 \neq j^1$ owns a house h^0 such that $h^0 \to j^0 \to h^1 \to j^1$ is part of the cycle of h^1 under $\succ'$. As j^1 loses brokerage right of h^1, by R5 there can be at most one other agent in this cycle, and hence the cycle is $h^0 \to j^0 \to h^1 \to j^1 \to h^0$. R6 implies that j^0 owns h^1 at ν, hence $j^0 = i^1$. Then i^1 gets h^1 under $\succ'$ in round t. However, as i^1 is both a $\sigma^{r-1}[\succ]$-owner and a $\sigma^{t-1}[\succ']$-owner, Claim 2 implies that he would point to h^2 not h^1 under $\succ'$ in round t, a contradiction.

- Subcase (b): R6 implies that j^0 owns h^1 at $\sigma^{t-1}[\succ'] \cup \{(j^1, h^0)\} \subset \nu$. Furthermore j^0 is unmatched at $\sigma^{r-1}[\succ]$, as otherwise the inductive assumption would imply that h^1 is matched at this submatching contrary to h^1 being unmatched at $\sigma^{r-1}[\succ]$. Thus, j^0 is unmatched at ν, and, by R4, he owns h^1 at ν. As i^1 also owns h^1 at ν, we have $j^0 = i^1$. Thus, $h^0 \to i^1 \to h^1 \to j^1 \to h^0$ is the cycle of h^1 under $\succ'$. But we know that i^1 (a $\sigma^{r-1}[\succ]$-owner) prefers h^2 over h^1. Because h^2 is unmatched at $\sigma^{t-1}[\succ']$, it must be that i^1 brokers h^2 at this submatching. Thus, $h^0 = h^2$. This contradicts the fact that h^0 is matched and h^2 is unmatched at $\sigma^{r-1}[\succ]$.

Either subcase leads to a contradiction showing that Case 3 cannot happen. This completes the proof of the inductive hypothesis. **QED**

D Supplementary Appendix: Proof of Theorem 2 (Implementation Result)

Let φ be a group strategy-proof and Pareto-efficient mechanism (fixed throughout the proof). We are to prove that φ may be represented as a TC mechanism. We will first construct the candidate control rights structure (c, b) and then show that the induced TC mechanism $\psi^{c,b}$ is equivalent to φ.

Let us start by introducing some useful terms and notation. Let $\sigma \in \overline{\mathcal{M}}$, $n \geq 0$ and $h^1, h^2, ..., h^n \in \overline{H_\sigma}$, and $i \in I$.

$\mathbf{P}_i[\sigma, h^1, ..., h^n]$ is the set of preferences $\succ_i$ of agent i such that

- if $i \in I_\sigma$, then

$$\sigma(i) \succ_i g \quad \text{for all } g \in H - \{\sigma(i)\},$$

- if $i \in \overline{I_\sigma}$, then

$$h^1 \succeq_i h^2 \succeq ... \succeq_i h^n \succ_i g \succ_i g' \text{ for all } g \in \overline{H_\sigma} - \left\{h^1, ..., h^n\right\} \text{ and all } g' \in H_\sigma.$$

That is, if i is not matched in submatching σ, $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ is the set of preferences that rank h^1,...,h^n in order over the remaining houses unmatched under σ, and rank those over the houses matched under σ; otherwise, $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ is the set of preferences that rank agent i's match under σ over all other houses (observe that $\mathbf{P}_i[\varnothing] \equiv \mathbf{P}_i$).

$\mathbf{P}[\sigma, h^1, ..., h^n] \subseteq \mathbf{P}$ is the Cartesian product of $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ over all $i \in I$. We define

$$\mathbf{P}^*[\sigma, h] = \cup_{h' \in \overline{H_\sigma} - \{h\}} \mathbf{P}[\sigma, h, h'],$$

i.e., the set of preference profiles generated by $\mathbf{P}[\sigma, h]$ that rank the same house as the second choice across all agents unmatched under σ.

When σ is fixed, we will occasionally write $\langle h^1, ..., h^n, ...\rangle$ instead of $\mathbf{P}_i[\sigma, h^1, ..., h^n]$.

We are ready to introduce some new terminology for the mechanism φ that is similar to the control rights structure terminology of the TC mechanisms. To distinguish the two classes defined for TC and φ, we will suffix these new definitions with *.

A house $h \in \overline{H_\sigma}$ is an **owned* house at** $\sigma \in \overline{\mathcal{M}}$ if $\varphi[\succ]^{-1}(h) = i$ for all $\succ \in \mathbf{P}[\sigma, h]$ for some $i \in \overline{I_\sigma}$; we refer to i as the **owner* of** h **at** σ.

A house $e \in \overline{H_\sigma}$ is a **brokered* house at** $\sigma \in \overline{\mathcal{M}}$ if there exist some $\succ$ and $\succ' \in \mathbf{P}^*[\sigma, e]$ such that $\varphi[\succ]^{-1}(e) \neq \varphi[\succ']^{-1}(e)$. Agent k is the **broker* of** e **at** σ if e is a brokered* house at σ and for all $\succ \in \mathbf{P}^*[\sigma, e]$ house $\varphi[\succ](k)$ is the second choice of k in $\succ_k$. Observe that a house cannot be both owned* and brokered* at the same submatching. [17]

Notice that if φ is a TC mechanism and i is an owner at σ then i is an owner* at σ, and similarly for the broker*. Thus, owners* and brokers* are the *candidate* owners and brokers in the TC mechanism that we will construct. We will show that the starred terms can be used to determine a consistent control rights structure (c, b) and a TC mechanism $\psi^{c,b}$. The proof of Theorem 2 will be finished after we show that $\varphi = \psi^{c,b}$.

Two lemmas proved in Pápai (2000) will be useful. Following her definition, we say that j **envies** i **at** $\succ$ if

$$\varphi[\succ](i) \succ_j \varphi[\succ](j).$$

Lemma 4. *(Pápai 2000) For all $i, j \in I$, all $\succ \in \mathbf{P}$, and all $\succ_j^* \in \mathbf{P}_j$, if j envies i at $\succ$ and $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$, then*

$$\varphi[\succ](i) \succ_i \varphi[\succ_j^*, \succ_{-j}](i).$$

[17] It may appear from the definitions that there is a third option for an unmatched house besides being owned* and brokered* at a submatching. Proposition However, 4 below shows that these are the only two options.

Lemma 5. *(Pápai 2000) For all $i, j \in I$, all $\succ \in \mathbf{P}$, and all $\succ_j^* \in \mathbf{P}_j$, if j envies i at $\succ$ and $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$, then there exists $\succ_i^* \in \mathbf{P}_i$ such that*

$$\varphi[\succ_i^*, \succ_j^*, \succ_{-\{i,j\}}](i) = \varphi[\succ](j).$$

This last lemma allows us to prove

Lemma 6. *For all $i, j \in I$, all $\succ \in \mathbf{P}$, and all $\succ_j^* \in \mathbf{P}_j$, if j envies i at $\succ$, then*

$$\varphi[\succ_j^*, \succ_{-j}](i) \succeq_i \varphi[\succ](j).$$

Proof of Lemma 6. If $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$ then the lemma follows from Lemma 5 and strategy-proofness of i. If $\varphi[\succ_j^*, \succ_{-j}](i) = \varphi[\succ](i)$ then Pareto efficiency of $\varphi(\succ)$ implies that i cannot envy j when j envies i and hence the claim of the lemma follows. QED

D.1 The Starred Control Rights Structure is Well Defined

The lemma below shows that if a house does not have a well-defined owner*, then it has a well-defined broker*. Thus the starred (candidate) control rights structure is well defined. All lemmas in this section are formulated and proven at a fixed submatching $\sigma \in \overline{\mathcal{M}}$.

Lemma 7. *Let $\sigma \in \overline{\mathcal{M}}$. For all $i \in I_\sigma$ and all $h \in \overline{H_\sigma}$,*

$$\varphi[\succ](i) = \sigma(i) \text{ for all } \succ \in \mathbf{P}[\sigma, h].$$

Proof of Lemma 7. Suppose that an agent in $i \in I_\sigma$ does not get $\sigma(i)$ at $\varphi[\succ]$. Then we can create a new matching by assigning all agents in $\overline{I_\sigma}$ that get a house in H_σ a house in $\overline{H_\sigma}$ that was assigned to an agent in I_σ, all other agents j in $\overline{I_\sigma}$ the house $\varphi[\succ](j)$, and all agents j in I_σ the house $\sigma(j)$. Since each agent in $\overline{I_\sigma}$ ranks houses in H_σ lower than houses in $\overline{H_\sigma}$ and each agent in I_σ ranks his σ-house as his first choice, this new matching Pareto dominates $\varphi[\succ]$, contradicting that φ is Pareto efficient. **QED**

Lemma 8. *Let $\sigma \in \overline{\mathcal{M}}$ and $e, h \in \overline{H_\sigma}$. Then there exists some agent $i \in \overline{I_\sigma}$ such that $\varphi[\succ](i) = e$ for all $\succ \in \mathbf{P}[\sigma, e, h]$.*

Proof of Lemma 8. By way of contradiction suppose that $\succ, \succ' \in \mathbf{P}[\sigma, e, h]$ are such that $\varphi[\succ](i) = e \neq \varphi[\succ'](i)$. Without loss of generality, we assume that $\succ$ and $\succ'$ differ only in preferences of a single agent $j \in \overline{I_\sigma}$. Let $g = \varphi[\succ](j)$. By strategy-proofness for j, we have $j \neq i$ and $g \neq e$. Moreover, by Maskin monotonicity, if it were true that $g = h$, then $\varphi[\succ'] = \varphi[\succ]$ would be true, contradicting that $\varphi[\succ'] \neq \varphi[\succ]$. Thus, $g \neq h$. We may further assume that

$$\succ_i \in \langle e, h, g, ... \rangle ,$$

as Maskin monotonicity for i implies that $\varphi(\succ)$ does not depend on how i ranks houses below e, and strategy-proofness for i implies that we still have $e \neq \varphi[\succ_i, \succ'_{-i}](i) = \varphi[\succ'](i)$.

Let $g' = \varphi[\succ'](j)$. By non-bossiness, $g' \neq g$ and by strategy-proofness $g' \neq e, h$. Maskin monotonicity for j allows us also to assume that

$$\succ_j \in \langle e, h, g, g', ... \rangle \text{ and } \succ'_j \in \langle e, h, g', g, ... \rangle .$$

Let $i' \in \overline{I_\sigma}$ be the agent who gets e at $\succ'$, and $k \in \overline{I_\sigma}$ be the agent who gets h at $\succ$. Notice that such agents exist because of Pareto efficiency. Because neither i nor j gets e at $\succ'$, we have $i' \neq i, j$. Furthermore, we saw above that j does not get h at $\succ$, and Lemma 4 implies that neither i nor i' gets h at $\succ$. Thus $k \neq i, i', j$.

Claim 1. (1) Under $\succ$, agents i, j, k are matched with houses as follows

$$\varphi[\succ](i) = e, \varphi[\succ](j) = g, \text{ and } \varphi[\succ](k) = h$$

(2) Under $\succ'$, agents i', j, k are matched with houses as follows

$$\varphi[\succ'](i') = e, \varphi[\succ'](j) = g', \varphi[\succ'](i) = g, \text{ and } \varphi[\succ'](k) = h.$$

Proof of Claim 1. The first five equalities were proved or assumed above and are listed for convenience only. The last two equalities require an argument. First, let us show that $\varphi[\succ'](i) = g$. Since agent j envies i at $\succ$ and $\varphi[\succ](j) = g$, Lemma 6 implies that i gets at least g at $\succ' = \left(\succ_{-j}, \succ'_j\right)$. Hence, $\varphi[\succ'](i) \in \{h, g\}$. Furthermore, Lemma 4 tells us that j cannot envy i at $\succ'$. Hence, $\varphi[\succ'](i) = g$.

Second, let us show that $\varphi[\succ'](k) = h$. Consider an auxiliary preference ranking $\tilde{\succ}_k \in \langle e, h, g, ... \rangle$ that agrees with $\succ_k$ except possibly for the relative ranking of g. Maskin monotonicity implies that

$$\varphi\left(\tilde{\succ}_k, \succ_{-k}\right) = \varphi(\succ).$$

Thus, agent j envies k at $(\tilde{\succ}_k, \succ_{-k})$ and $\varphi[\tilde{\succ}_k, \succ_{-k}](j) = g$, and thus Lemma 6 implies that $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j](k) \succeq_k g$. Strategy-proofness for k implies that k cannot get e at

$\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$. To prove that k gets h it is thus enough to show that i gets g at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$. The proof is analogous to the proof equality $\varphi[\succ'](i) = g$ above: i gets at least g at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$, and because j cannot envy i at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$ (by Lemma 4), we must have $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j](i) = g$. We have thus shown that $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j](k) = h$, and by Maskin monotonicity it must be that $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j] = \varphi[\succ_k, \succ_{-k,j}, \succ'_j] = \varphi[\succ']$. Thus, $\varphi[\succ'](k) = h$, and the claim is proved. QED

The above claim and Maskin monotonicity, allows us to assume in the sequel that

$$\succ_k \in \langle e, h, g, ... \rangle .$$

Let us also fix three auxiliary preference rankings for use in the subsequent analysis:

$$\begin{aligned} &\succ^*_i \in \langle h, e, g, ... \rangle , \\ &\succ^*_{i'} \in \langle h, e, ... \rangle , \text{ and} \\ &\succ^*_k \in \langle e, g, h, ... \rangle . \end{aligned}$$

We will prove a number of claims.

Claim 2. (1) $\varphi[\succ^*_i, \succ_{-i}](i) = h$ and $\varphi[\succ^*_{i'}, \succ'_{-i'}](i') = h$.
(2) $\varphi[\succ^*_i, \succ_{-i}](j) = g$.

Proof of Claim 2.

(1) By strategy-proofness for i, $\varphi[\succ^*_i, \succ_{-i}](i) \succeq^*_i e$. Everybody else in $\overline{I_\sigma}$ ranks e over h. Thus, by Lemma 7 and Pareto efficiency, i should get h at $[\succ^*_i, \succ_{-i}]$. The symmetric argument yields $\varphi[\succ^*_{i'}, \succ'_{-i'}](i') = h$.

(2) By Maskin monotonicity for i, $\varphi[\succ^*_i, \succ'_{-i}] = \varphi[\succ']$. Thus, j gets g' at $[\succ^*_i, \succ'_{-i}]$. By strategy-proofness for j, agent j gets at least g' and no house better than g at $[\succ^*_i, \succ_{-i}]$ (recall that between $\succ_{-i}$ and $\succ'_{-i}$ only j changes preferences). Thus, in order to prove the claim that j gets g at $[\succ^*_i, \succ_{-i}]$ it is enough to show that he cannot get g' at $[\succ^*_i, \succ_{-i}]$. Assume to the contrary that j gets g' at $[\succ^*_i, \succ_{-i}]$. Then, non-bossiness would imply that i gets h at $[\succ^*_i, \succ'_{-i}]$. By strategy-proofness for i, he gets at least h at $\succ'$. But then j envies i both at $\succ$ and $\succ' = \left(\succ'_j, \succ_{-j}\right)$ and by Lemma 4, i must get same house at these two profiles. This contradiction proves Claim 2. QED

Claim 3. (1) $\varphi[\succ^*_k, \succ_{-k}](k) = g$.
(2) $\varphi[\succ^*_k, \succ'_{-k}] = \varphi[\succ^*_k, \succ_{-k}]$.

Proof of Claim 3. (1) Because k gets h at $\succ$, strategy-proofness implies that k cannot get e and gets at least h at $[\succ_k^*, \succ_{-k}]$. Thus, k gets h or g at $[\succ_k^*, \succ_{-k}]$. Everybody else in $\overline{I_\sigma}$ ranks h over g. Thus, by Lemma 7 and Pareto efficiency, agent k should get g at $[\succ_k^*, \succ_{-k}]$.

(2) Profiles,$[\succ_k^*, \succ'_{-k}]$ and $[\succ_k^*, \succ_{-k}]$ differ only in preferences of agent j who ranks g above g' at $\succ_j$ and the other way at $\succ'_j$. We established in part (1) that j does not get g at $[\succ_k^*, \succ_{-k}]$. Maskin monotonicity for j implies $\varphi[\succ_k^*, \succ'_{-k}] = \varphi[\succ_k^*, \succ_{-k}]$. QED

Claim 4. $\varphi[\succ_k^*, \succ_{-k}](i) = e$ and $\varphi[\succ_k^*, \succ_{-k}](i') = h$.

Proof of Claim 4. Because agent k envies agent i at $\succ$, Lemma 6 implies that i gets at least $h = \varphi[\succ](k)$ at $[\succ_k^*, \succ_{-k}]$. Hence $\varphi[\succ_k^*, \succ_{-k}](i) \in \{e, h\}$. Analogously, because agent k envies agent i' at $\succ'$, Lemma 6 implies that i' gets at least $h = \varphi[\succ'](k)$ at $[\succ_k^*, \succ'_{-k}]$. Hence $\varphi[\succ_k^*, \succ'_{-k}](i') \in \{e, h\}$. By Claim 3(2), $\varphi[\succ_k^*, \succ_{-k}](i') \in \{e, h\}$. Thus,

$$\{\varphi[\succ_k^*, \succ_{-k}](i), \varphi[\succ_k^*, \succ_{-k}](i')\} = \{e, h\}.$$

This equality implies that to prove the claim it is enough to show that $\varphi[\succ_k^*, \succ_{-k}](i) = h$ and $\varphi[\succ_k^*, \succ_{-k}](i') = e$ cannot both be true. Suppose they are. By Maskin monotonicity for i, $\varphi[\succ_k^*, \succ_{-k}] = \varphi[\succ_k^*, \succ_i^*, \succ_{-\{k,i\}}]$. This equivalence and Claim 3(1) give $\varphi[\succ_k^*, \succ_i^*, \succ_{-\{k,i\}}](k) = g$. By strategy-proofness, agent k gets at least g and not e at $[\succ_i^*, \succ_{-i}]$. By Claim 2(1), we must thus have $\varphi[\succ_i^*, \succ_{-i}](k) = g$. But this contradicts Claim 2(2). QED

Claim 5. (1) $\varphi[\succ_k^*, \succ_{-k}] = \varphi[\succ_k^*, \succ_{i'}^*, \succ_{-\{k,i'\}}] = \varphi[\succ_k^*, \succ_{i'}^*, \succ'_j, \succ_{-\{k,i',j\}}]$.
(2) $\varphi[\succ_k^*, \succ_{i'}^* \succ'_{-\{k,i'\}}](k) = g$.

Proof of Claim 5. The first equality of part (1) follows from Maskin monotonicity for i' and Claim 4. To prove the second equality of part (1), notice that at preference profile $(\succ_k^*, \succ_{-k})$ agent j does not get e or h (by Claim 4), and he does not get g by Claim 3(1). Thus the second equality follows from Maskin monotonicity for j. Now, part (2) of the claim follows from part (1) and Claim 3(1). QED

Claim 6. $\varphi[\succ_{i'}^*, \succ'_{-i'}](i) = e$.
Proof of Claim 6. Strategy-proofness for k and Claim 5(2) imply that agent k gets at least g at $\left(\succ_{i'}^*, \succ'_{-i'}\right)$ but does not get e. By Claim 2(1), k gets g at $\left(\succ_{i'}^*, \succ'_{-i'}\right)$. By non-bossiness for k and part (2) of Claim 5,

$$\varphi[\succ_{i'}^*, \succ'_{-i'}] = \varphi[\succ_k^*, \succ_{i'}^* \succ'_{-\{k,i'\}}].$$

This equality and part (1) of Claim 5 imply that

$$\varphi[\succ^*_{i'}, \succ'_{-i'}] = \varphi[\succ^*_k, \succ_{-k}].$$

This equation and Claim 4 give us $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) = e$. QED

Claim 7. $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) \neq e$.

Proof of Claim 7. Let us first prove that $\varphi[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}](i) \neq h$. Suppose not. Then, Maskin monotonicity for i' gives $\varphi[\succ^*_i, \succ'_{-i}] = \varphi[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}]$, and in particular, $\varphi[\succ^*_i, \succ'_{-i}](i) = h$. By strategy-proofness for i, $\varphi[\succ'](i) \succeq_i h$, contradicting that $\varphi[\succ'](i') = e$ and $\varphi[\succ'](k) = h$, and proving the required inequality.

Since $\succ_i$ pushes down the ranking of h in $\succ^*_i$, the just-proven inequality and Maskin monotonicity for i give:

$$\varphi[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}] = \varphi[\succ^*_{i'}, \succ'_{-i'}].$$

A symmetric argument implies that $\varphi[\succ^*_{\{i,i'\}}, \succ_{-\{i,i'\}}](i') \neq h$ and

$$\varphi[\succ^*_{\{i,i'\}}, \succ_{-\{i,i'\}}] = \varphi[\succ^*_i, \succ_{-i}].$$

Contrary to the claim we are proving, suppose that $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) = e$. Then, the first of the above-displayed equalities implies $\varphi[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}](i) = e$ and, hence, j envies i at $[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}] = [\succ^*_{\{i,i'\}}, \succ_{-\{i,i',j\}}, \succ'_j]$. This, however, leads to a contradiction with Lemma 4, because Claim 2 and the second above-displayed equality implies that $\varphi[\succ^*_{\{i,i'\}}, \succ_{-\{i,i',j\}}, \succ_j](i) = h$. Thus, we have shown that $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) \neq e$. QED

The contradiction between Claims 6 and 7 shows that the initial assumption $\varphi[\succ](i) = e \neq \varphi[\succ'](i)$ cannot be correct. **QED**

Lemma 9. ***(Existence and uniqueness of a broker* for each brokered* house)*** *Let $\sigma \in \overline{\mathcal{M}}$ and e be a brokered* house at σ. Then there exists an agent $k \in \overline{I_\sigma}$ who is the unique broker* of e at σ.*

Proof of Lemma 9. Let $\sigma \in \overline{\mathcal{M}}$ and e be a brokered* house at σ. We start with the following preparatory claim:

Claim 1. If h, h' are two different houses in $\overline{H_\sigma} - \{e\}$, and $\succ, \succ' \in \mathbf{P}[\sigma, e, h, h']$, then $\varphi[\succ']^{-1}(h) = \varphi[\succ]^{-1}(h)$.

Proof of Claim 1. By Lemma 8, $\varphi[\succ']^{-1}(e) = \varphi[\succ]^{-1}(e)$. Let $i = \varphi[\succ]^{-1}(e)$. Also let $\succ^*$ and $\succ'^*$be monotonic transformations of $\succ$ and $\succ'$, respectively, such that i ranks e first, all agents in $\overline{I_\sigma}$ rank e below all houses in $\overline{H_\sigma} - \{e\}$, and the relative rankings of all other houses at $\succ^*$, $\succ$ and $\succ'^*$, $\succ'$ are respectively the same. By Maskin monotonicity, $\varphi[\succ^{*\prime}] = \varphi[\succ']$ and $\varphi[\succ^*] = \varphi[\succ]$. Also $\succ^*, \succ'^* \in \mathbf{P}[\sigma \cup \{(i,e)\}, h, h']$. Thus, by Lemma 8, $\varphi[\succ^*]^{-1}(h) = \varphi[\succ^{*\prime}]^{-1}(h)$. Hence, $\varphi[\succ']^{-1}(h) = \varphi[\succ'^*]^{-1}(h) = \varphi[\succ^*]^{-1}(h) = \varphi[\succ]^{-1}(h)$. QED

Claim 2. If h, h' are two different houses in $\overline{H_\sigma} - \{e\}$, and profiles $\succ \in \mathbf{P}[\sigma, e, h, h']$ and $\succ' \in \mathbf{P}[\sigma, e, h']$ are such that $\varphi[\succ']^{-1}(e) \neq \varphi[\succ]^{-1}(e)$, then $\varphi[\succ']^{-1}(h') = \varphi[\succ]^{-1}(h)$.

Proof of Claim 2. Let $k' = \varphi[\succ']^{-1}(h')$ and $\succ^* \in \mathbf{P}[\sigma, e, h', h]$ be such that the only difference between $\succ^*$ and $\succ$ is the relative ranking of house h'. Since by Claim 1 $\varphi[\succ^*]^{-1}(h') = \varphi[\succ']^{-1}(h') = k'$ and since we lower house h' in everybody's preferences except k' at $[\succ^*_{k'}, \succ_{-k'}]$, by Maskin monotonicity

$$\varphi[\succ^*_{k'}, \succ_{-k'}] = \varphi[\succ^*].$$

In particular, $\varphi[\succ^*_{k'}, \succ_{-k'}](k') = h'$. By strategy-proofness for k', we have $\varphi[\succ](k') \in \{h, h'\}$. On the other hand, by Lemma 8,

$$\varphi[\succ^*]^{-1}(e) = \varphi[\succ']^{-1}(e).$$

The two above displayed equalities imply that $\varphi[\succ^*_{k'}, \succ_{-k'}]^{-1}(e) = \varphi[\succ']^{-1}(e)$. By assumption of the claim, $\varphi[\succ]^{-1}(e) \neq \varphi[\succ']^{-1}(e) = \varphi[\succ^*_{k'}, \succ_{-k'}]^{-1}(e)$. By non-bossiness, agent k' changes his own allocation while switching between the two profiles $\succ$ and $[\succ^*_{k'}, \succ_{-k'}]$, implying that $\varphi[\succ](k') = h$. QED

Claim 3. If h, h' are two different houses in $\overline{H_\sigma} - \{e\}$, and $\succ \in \mathbf{P}[\sigma, e, h]$, and $\succ' \in \mathbf{P}[\sigma, e, h', h]$, then $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h')$.

Proof of Claim 3. If $\varphi[\succ]^{-1}(e) \neq \varphi[\succ']^{-1}(e)$, then Claim 3 reduces to Claim 2. Assume that $\varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e)$. Because e is brokered* at σ, there exists some $h'' \in \overline{H_\sigma} - \{e\}$ such that for some $\succ'' \in \mathbf{P}[\sigma, e, h'']$,

$$\varphi[\succ'']^{-1}(e) \neq \varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e).$$

By Lemma 8, $h'' \neq h$. By the same lemma, we assume that $\succ'' \in \mathbf{P}[\sigma, e, h'', h]$.

By Claim 2, $\varphi[\succ'']^{-1}(h'') = \varphi[\succ]^{-1}(h)$ and $\varphi[\succ'']^{-1}(h'') = \varphi[\succ']^{-1}(h')$, implying that $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h')$. QED

Claim 4. If $h \in \overline{H_\sigma} - \{e\}$ and $\succ, \succ' \in \mathbf{P}[\sigma, e, h]$, then $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h)$.

Proof of Claim 4. By Lemma 8, $\varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e)$. Because e is brokered* at σ, there exists some $h'' \in \overline{H_\sigma} - \{e\}$ such that for some $\succ'' \in \mathbf{P}[\sigma, e, h'']$,

$$\varphi[\succ'']^{-1}(e) \neq \varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e).$$

By Lemma 8, $h'' \neq h$ and, by the same lemma, we may assume $\succ'' \in \mathbf{P}[\sigma, e, h'', h]$. By Claim 3, $\varphi[\succ'']^{-1}(h'') = \varphi[\succ]^{-1}(h)$ and $\varphi[\succ'']^{-1}(h'') = \varphi[\succ']^{-1}(h)$, implying that $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h)$. QED

To complete the proof of the lemma notice that e being brokered implies there is at least one house in $\overline{H_\sigma} - \{e\}$. Let h and $h' \in \overline{H_\sigma} - \{e\}$, $\succ \in \mathbf{P}[\sigma, e, h]$, $\succ' \in \mathbf{P}[\sigma, e, h']$. If $h = h'$, then $\varphi[\succ']^{-1}(h) = \varphi[\succ]^{-1}(h)$ by Claim 4. Consider the case $h \neq h'$, and fix $\succ^* \in \mathbf{P}[\sigma, e, h, h']$. By Claim 3, $\varphi[\succ']^{-1}(h) = \varphi[\succ^*]^{-1}(h')$ and by Claim 4 $\varphi[\succ^*]^{-1}(h) = \varphi[\succ]^{-1}(h)$, implying that $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h')$. Thus, the agent $\varphi[\succ]^{-1}(h)$ is the unique broker* of e at σ. **QED**

Lemma 10. *Let $\sigma \in \overline{\mathcal{M}}$, $i \in \overline{I_\sigma}$, and $h \in \overline{H_\sigma}$. If $\varphi[\succ](i) = h$ for all $\succ \in \mathbf{P}^*[\sigma, h]$ then i owns* h at σ.*

Proof of Lemma 10. Let us start with two preparatory claims:

Claim 1. Suppose $\sigma \in \overline{\mathcal{M}}$, houses g and $h \in \overline{H_\sigma}$ are different, and agent $i \in \overline{I_\sigma}$. If $\varphi[\succ'](i) = h$ for all $\succ' \in \mathbf{P}[\sigma, g, h]$, then $\varphi[\succ_i^*, \succ_{-i}](i) = g$ for all $\succ_i^* \in \langle g, ...\rangle$ and all $\succ_{-i} \in \mathbf{P}_{-i}[\sigma, h]$.

Proof of Claim 1. Let $\succ_{-i} \in \mathbf{P}_{-i}[\sigma, h]$. Take any $\succ_i \in \langle h, g, ...\rangle$. If $\varphi[\succ](i) = h$, then Pareto efficiency and strategy-proofness imply that $\varphi[\succ_i^*, \succ_{-i}](i) = g$ for all $\succ_i^* \in \langle g, h, ...\rangle$, and furthermore, by strategy-proofness, for all $\succ_i^* \in \langle g, ...\rangle$. It remains to consider the case $\varphi[\succ](i) \neq h$.

Take $\succ' \in \mathbf{P}[\sigma, h, g]$ such that $\succ'$ and $\succ$ coincide other than unmatched agents' ranking of house g. We have $\varphi[\succ'](i) = h$ by the hypothesis of the claim. Two cases are possible: $\varphi[\succ](i) = g$ and $\varphi[\succ](i) \neq g$. If $\varphi[\succ](i) = g$, then by strategy-proofness, $\varphi[\succ_i^*, \succ_{-i}](i) = g$ and we are done. Thus, in the remainder assume that there exists some agent $k = \varphi[\succ]^{-1}(g) \neq i$. By Maskin monotonicity, $\varphi[\succ'_{\{i,k\}}, \succ_{-\{i,k\}}](i) = h$ and $\varphi[\succ'_{\{i,k\}}, \succ_{-\{i,k\}}](k) = g$.

Let $\succ_i^* \in \langle g, h, ...\rangle$. By strategy-proofness, agent i gets at least h at $[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}]$; and by Pareto efficiency, agent i gets g. Also recall that $\varphi[\succ](i) \prec_i g$ and $\varphi[\succ](k) = g$. Thus,

$\varphi[\succ_i^*, \succ_k', \succ_{-\{i,k\}}](k) \neq h$ because otherwise agents i and k could jointly improve upon their $\varphi[\succ]$ allocation by submitting $[\succ_i^*, \succ_k']$ at $\succ$, contradicting group strategy-proofness. Thus, $g \succ_k' \varphi[\succ_i^*, \succ_k', \succ_{-\{i,k\}}](k)$, and furthermore, Maskin monotonicity implies $\varphi[\succ_i^*, \succ_k', \succ_{-\{i,k\}}] = \varphi[\succ_i^*, \succ_{-i}]$. In particular, $\varphi[\succ_i^*, \succ_{-i}](i) = g$. QED

Claim 2. Suppose $\sigma \in \overline{\mathcal{M}}$, houses g and $h \in \overline{H_\sigma}$ are different, and $\varphi[\succ']^{-1}(h) = i \in \overline{I_\sigma}$ for all $\succ' \in \mathbf{P}[\sigma, g, h]$. If $\succ \in \mathbf{P}[\sigma, h]$ and there is some $\succ' \in \mathbf{P}[\sigma, h, g]$ such that $\succ_k \in \langle h, g, ...\rangle$ for $k = \varphi[\succ']^{-1}(g)$, then $\varphi[\succ](i) = h$.

Proof of Claim 2. By way of contradiction, assume that i is the owner* of h at σ, that $\succ' \in \mathbf{P}[\sigma, h, g]$, and that $k = \varphi[\succ']^{-1}(g)$, but there is some $\succ \in \mathbf{P}[\sigma, h]$ such that $\succ_k \in \langle h, g, ...\rangle$ and $\varphi[\succ]^{-1}(h) \neq i$. By strategy-proofness, we can choose $\succ_i \in \langle h, g, ...\rangle$. Furthermore, we can choose $\succ$ such that $\succ$ and $\succ'$ differ only in the preferences of a single agent $j \in \overline{I_\sigma}$ and in how house g is ranked by the agents.

Let $\succ^* \in \mathbf{P}[\sigma, h]$ be the unique profile, such that $\succ^*$ and $\succ$ differ only in the preferences of agent j, and $\succ^*$ and $\succ'$ differ only in how house g is ranked by the agents. Notice that $j \neq k$ as otherwise Maskin monotonicity would imply that i gets h at $\succ$. Thus, $\succ_k^* \in \langle h, g, ...\rangle$, and Maskin monotonicity implies that $\varphi[\succ^*](i) = h$.

Let h' be the house that j gets at $\succ$ and let $\succ''$ be the unique profile in $\mathbf{P}[\sigma, h, g]$ such that $\succ''$ and $\succ$ differ only in how house g is ranked by agents. By Maskin monotonicity, we may assume that $\succ_j'' \in \langle h, g, h', ...\rangle$.

By Claim 1 and strategy-proofness, $\varphi[\succ_j'', \succ_{-j}](i)$ equals either h or g. At the same time, strategy-proofness implies that $\varphi[\succ_j'', \succ_{-j}](j)$ equals either g or h'. In either case, agent j prefers the allocation of agent i at $[\succ_j'', \succ_{-j}]$. If $\varphi[\succ_j'', \succ_{-j}](i) = g$, this would be a contradiction with Lemma 3, as j could improve the allocation of i by switching from $[\succ_j'', \succ_{-j}]$ to $[\succ_j^*, \succ_{-j}] = \succ^*$. Hence, $\varphi[\succ_j'', \succ_{-j}](i) = h$, and by non-bossiness $\varphi[\succ_j'', \succ_{-j}](j) = g$. However, $k \neq j$ gets g at $\succ'$ and by strategy-proofness j cannot get it at $[\succ_j'', \succ_{-j}']$. This is a contradiction because $[\succ_j'', \succ_{-j}] = [\succ_j'', \succ_{-j}']$. QED

We are ready to finish the proof of the lemma. Fix $\sigma \in \overline{\mathcal{M}}$. We proceed by way of contradiction. Let $i \in \overline{I_\sigma}$ be such that $\varphi[\succ'](i) = h$ for all $\succ' \in \mathbf{P}^*[\sigma, h]$. Let $\succ \in \mathbf{P}[\sigma, h]$ be such that $\varphi[\succ]^{-1}(h) = j \neq i$. For all unmatched houses $g \neq h$ at σ, define $\succ^g$ to be the unique profile in $\mathbf{P}[\sigma, h, g]$ that differs from $\succ$ only in how agents rank g.

Take a house $g_1 \neq h$ unmatched at σ, and let k_1 be the agent that gets g_1 at $\succ^{g_1}$. By Claim 2, agent i gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_1 ranks g_1 second. Hence, by Maskin monotonicity i also gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_1 gets g_1.

Let $g_2 = \varphi[\succ](k_1)$ and let k_2 be the agent that gets g_2 at $\succ^{g_2}$. Because i does not get h at $\succ$, the previous paragraph yields $g_2 \neq g_1$ and $k_2 \neq k_1$. As in the previous paragraph, Claim 2 and Maskin monotonicity imply that i gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_2 gets g_2 or ranks g_2 second.

Furthermore, we will show that i gets h at any profile $\succ' \in \mathbf{P}[\sigma, h]$ at which k_2 ranks g_1 second. Indeed, suppose $\succ'_{k_2} \in \langle h, g_1, ... \rangle$ and i does not get h at $\succ'$. Let $\succ''_i \in \langle h, g_1, ... \rangle$. By Claim 1 and strategy-proofness, agent i gets g_1 at $[\succ''_i, \succ'_{-i}]$. By the previous paragraph and strategy-proofness, k_2 does not get h at $[\succ''_i, \succ'_{-i}]$, and thus k_2 envies i at $[\succ''_i, \succ'_{-i}]$. However, by the previous paragraph k_2 can improve the outcome of agent i, contrary to Lemma 4. Thus, i gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_2 ranks g_1 second.

Let g_3 be the house that k_2 gets at $\succ$ and let k_3 be the agent that gets g_3 at $\succ^{g_3}$. As above, we can show that i gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_3 ranks g_3 or g_2 or g_1 second.

Since the number of agents is finite, by repeating the procedure we arrive at an agent k_n who ranks one of the houses $g_1, ..., g_n$ second at $\succ$. That means that i gets h at $\succ$, a contradiction that concludes the proof. **QED**

Lemmas 9 and 10 and the definitions of owned* and brokered* houses give us the key result of this subsection:

Proposition 4. ***(Houses are either brokered* or owned*)*** *For any $\sigma \in \overline{\mathcal{M}}$, any house $h \in \overline{H_\sigma}$ is owned* or brokered* at σ, but not both. In particular, if there is a single owner* at σ then h is not a brokered* house.*

Although the starred control rights do not allow having a broker* when there is a single owner*, R1-R6 do not eliminate this possibility from consistent control rights structures. However, it turns out that for any control rights obeying R1-R6, it is trivial to construct an equivalent one in which control rights are set equal to the original ones at all submatchings except possibly submatchings with single owners, where all houses are now owned by this original owner. In particular, if there is a broker in a submatching with a single owner in the original control rights structure, then in the superior submatching that matches the owner with the originally brokered house, the new control rights are set such that the original broker owns the remaining houses.

D.2 The Starred Control Rights Structure Satisfies R1-R6

Before proving R1-R6 let us state and prove one more auxiliary result.

Lemma 11. ***(Relationship between brokerage* and ownership*)***. *Let $\sigma \in \overline{\mathcal{M}}$, agent k be a broker* of house e at σ, and $\succ'' \in \mathbf{P}^*[\sigma, e]$. Then agent $\varphi[\succ'']^{-1}(e)$ is the owner* of house $\varphi[\succ''](k)$ at σ.*

Proof of Lemma 11. Let $\succ'' \in \mathbf{P}^*[\sigma, e]$ and $h = \varphi[\succ''](k)$. Because k is a broker* at σ, Lemma 9 implies that house h is agent k's second choice. Since $\succ'' \in \mathbf{P}^*[\sigma, e]$, house h is the second choice of all agents in $\overline{I_\sigma}$ at $\succ''$, and thus,

$$\succ'' \in \mathbf{P}[\sigma, e, h].$$

There exists an agent $i \in (\overline{I_\sigma}) - \{k\}$ such that $\varphi[\succ'']^{-1}(e) = i$. By Lemma 8, for all $\succ \in \mathbf{P}[\sigma, e, h]$, agent i gets e at $\succ$. We are to show that i is the owner* of h at σ.

Claim 1. If $\succ \in \mathbf{P}[\sigma, e, h]$, then $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$.

Proof of Claim 1. The first claim follows from Lemma 8, and the second from Lemma 9. QED

Claim 2. $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$.

Proof of Claim 2. Let preference profile $\succ$ be such that $\succ_{i'} = \succ''_{i'}$ for all $i' \in \{k, i\} \cup I_\sigma$ and all houses in $\overline{H_\sigma}$ are ranked above the houses in H_σ by $i' \in \overline{I_\sigma}$. By Claim 1 and Maskin monotonicity, $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$. QED.

Claim 3. $\varphi[\succ_i^*, \succ_{-i}](i) = h$.

Proof of Claim 3. Let $\succ_i^* \in \langle h, e, ...\rangle$. By the strategy-proofness of φ, since $\varphi[\succ](i) = e$, agent i gets at least e at $[\succ_i^*, \succ_{-i}]$, and since all other agents in $\overline{I_\sigma}$ prefer e over h, the Pareto efficiency of φ implies that $\varphi[\succ_i^*, \succ_{-i}](i) = h$.

Claim 4. $\varphi[\succ_k^*, \succ_{-k}] = \varphi[\succ]$.

Proof of Claim 4. Let $\succ_k^* \in \langle h, e, ...\rangle$. Since $\varphi[\succ](k) = h$, profile $[\succ_k^*, \succ_{-k}]$ is a monotonic transformation of $\succ$ and by the Maskin monotonicity of φ, we have $\varphi[\succ_k^*, \succ_{-k}] = \varphi[\succ]$.

Claim 5. $\varphi[\succ_{\{i,k\}}^*, \succ_{-\{i,k\}}](i) = h$.

Proof of Claim 5. By Claim 4, $\varphi[\succ_k^*, \succ_{-k}](i) = \varphi[\succ](i) = e$, and, by the strategy-proofness of φ, i gets at least e at $[\succ_{\{i,k\}}^*, \succ_{-\{i,k\}}]$. Thus, if i does not get h at $[\succ_{\{i,k\}}^*, \succ_{-\{i,k\}}]$ then one of the following two cases would have to obtain.

Case 1. An agent $j \notin \{i, k\}$ gets h at $[\succ_{\{i,k\}}^*, \succ_{-\{i,k\}}]$: Then i gets e, and k gets some house worse than e. But then jointly i and k can report $\succ_{\{i,k\}}$ instead of $\succ_{\{i,k\}}^*$ and they would jointly improve at $\succ_{\{i,k\}}^*$, i.e., $\varphi[\succ](i) = e = \varphi[\succ_{i,k}^*, \succ_{-i,k}](i)$ and $\varphi[\succ](k) = h \succ_k^* \varphi[\succ_{i,k}^*, \succ_{-i,k}](k)$, contradicting that φ is group strategy-proof.

Case 2. Agent k gets h at $[\succ^*_{i,k}, \succ_{-i,k}]$: By the strategy-proofness of φ, agent k should at least get h at $[\succ^*_i, \succ_{-i}]$. But we know by Step 2 that $\varphi[\succ^*_i, \succ_{-i}](i) = h$, thus we should have $\varphi[\succ^*_i, \succ_{-i}](k) = e$. Then by the Maskin monotonicity of φ, we have $\varphi[\succ^*_{i,k}, \succ_{-i,k}](i) = \varphi[\succ^*_i, \succ_{-i}](i) = h$ where the last equality follows by Step 2, a contradiction that proves the claim. QED

Claim 6. If $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](i) = h$, then $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](k) \neq e$.

Proof of Claim 6. For an indirect argument, suppose that $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](i) = h$ and $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](k) = e$. Then, $\varphi[\succ^*_i, \succ_{-i}](k) = e$ by the strategy-proofness of φ. Since e is a brokered* house at σ, there exist some house $g \notin \{e, h\}$ and some preference profile $\succ' \in \mathbf{P}[\sigma, e, g]$ such that $\varphi[\succ']^{-1}(e) = j$ for some agent $j \notin \{i, k\}$. By Lemma 8, we may assume that each agent $i' \in \overline{I_\sigma}$ ranks houses other than g and h in the same way at $\succ'_{i'}$ and $\succ_{i'}$ and that $\succ'_{i'} \in \langle e, g, h, ...\rangle$. Since k is the broker* of e at σ, we have $\varphi[\succ'](k) = g$. By Maskin monotonicity,

$$\varphi[\succ'] = \varphi[\succ'_{\{i,k\}}, \succ_{-\{i,k\}}].$$

Now i gets a house weakly worse than h at $[\succ'_{\{i,k\}}, \succ_{-\{i,k\}}]$. However, if i and k manipulated and submitted $\succ^*_{\{i,k\}}$ instead of $\succ'_{\{i,k\}}$, they would get h and e respectively at $[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}]$. Both agents weakly improve, while k strictly improves. This contradicts the fact that φ is group strategy-proof. QED

Now, Claims 5 and 6 imply that $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](i) = h$ and $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](k) \neq e$. By Maskin monotonicity, we can drop the ranking of e in $\succ^*_i$ and $\succ^*_k$, and yet, the outcome of φ will not change. Recall that $\succ_{-\{i,k\}}$ was an arbitrary profile in which all houses in $\overline{H_\sigma}$ are ranked above the houses in H_σ by $i' \in \overline{I_\sigma} - \{i, k\}$. Thus, i gets h at all profiles of $\mathbf{P}[\sigma, h]$.
QED

The following six lemmas show that the starred control rights structure satisfies R1-R6 (respectively).

Lemma 12. ***(R1; Uniqueness of a brokered* house)****. Let $\sigma \in \overline{\mathcal{M}}$. If e is a brokered* house at σ, then no other house is a brokered* house at σ (and all other unmatched houses are owned* houses).*

Proof of Lemma 12. Let e be a brokered* house at σ. By Lemma 9, there is a broker* of e at σ; let us denote him as k. Consider a house $h \in \overline{I_\sigma} - \{e\}$. By Lemma 8, there is an agent i who gets e at all profiles in $\mathbf{P}[\sigma, e, h]$. By Lemma 10, i is the owner* of h. Thus h is not a brokered* house at σ. **QED**

Lemma 13. ***(R2; Last unmatched agent is an owner).*** *Let $\sigma \in \overline{\mathcal{M}}$, such that there exists a unique agent i unmatched at σ. Then i owns* all unmatched houses at $\sigma \in \overline{I_\sigma}$.*

Proof of Lemma 13. Let $\succ \in \mathbf{P}[\sigma, h]$ for $h \in \overline{H_\sigma}$. By Pareto efficiency of φ, $\varphi[\succ](i) = h$, implying that i owns* h at σ. **QED**

Lemma 14. ***(R3; Broker* does not own*).*** *Let $\sigma \in \overline{\mathcal{M}}$. If agent k is the broker* of house e at σ, then he cannot own* any houses at σ.*

Proof of Lemma 14. Suppose that k owns* a house $h \neq e$ at σ. By Lemma 8, there exists some agent $i \neq k$ who gets e at all profiles in $\mathbf{P}[\sigma, e, h]$. Thus, i gets h at all $\succ \in \mathbf{P}^*[\sigma, h]$, contradicting that k owns* h. **QED**

Lemma 15. ***(R4; Persistence of ownership*).*** *Let i own* h at some $\sigma \in \overline{\mathcal{M}}$. If $\sigma' \supsetneq \sigma$, and i and h are unmatched at σ', then i owns* h at σ'.*

Proof of Lemma 15. Imagine to the contrary that i gets h at all $\succ \in \mathbf{P}[\sigma, h]$, but there is some $\succ' \in \mathbf{P}[\sigma', h]$ such that some agent $j \in I_{\sigma'} - I_\sigma$, such that $j \neq i$, gets h at $\succ'$. Take $\succ \in \mathbf{P}[\sigma, h]$ such that

- for each agent $k \notin I_{\sigma'} - I_\sigma$, $\succ_k = \succ'_k$, and
- each agent $k \in I_{\sigma'} - I_\sigma$ ranks $\sigma'(k)$ as his second choice (just behind h) in $\succ_k$.

Each $k \in I_{\sigma'} - I_\sigma$ is indifferent between $\succ'$ and $\succ$ because:

- at $\succ'$ agent k gets $\sigma'(k)$ by Lemma 7,
- at $\succ$ agent k gets $\sigma'(k)$ by the Pareto efficiency of φ and the fact that $\varphi[\succ](i) = h$.

The only difference between the profiles $\succ'$ and $\succ$ are the preferences of the agents in $I_{\sigma'} - I_\sigma$. Thus, agents $I_{\sigma'} - I_\sigma$ are indifferent between $\succ$ and $\succ'$, while agent j is strictly better off at $\succ'$. This contradicts the fact that φ is group strategy-proof. **QED**

Lemma 16. ***(R5; Limited persistence of brokerage*)*** *Let $\sigma, \sigma' \in \overline{\mathcal{M}}$ be such that $\sigma' \supsetneq \sigma$. Suppose that agent k is the broker* of house e at σ, agent i is the owner* of house h at σ, and agent $i' \neq i$ is the owner* of house h' at σ. If k, i, i', e, h, h' are unmatched at σ', then k brokers* e at σ'.*

Proof of Lemma 16. First, notice that i gets e at all $\succ \in \mathbf{P}[\sigma, e, h]$ and i' gets e at all $\succ \in \mathbf{P}[\sigma, e, h']$, and k gets h and h', respectively by Lemma 11. Take $\succ^h \in \mathbf{P}[\sigma, e, h]$ and $\succ^{h'} \in \mathbf{P}[\sigma, e, h']$ such that each agent $j \in I_{\sigma'} - I_\sigma$ has $\sigma'(j)$ as his third choice and each agent $j \in I - I_{\sigma'}$ ranks each house unmatched at σ' above all houses matched at σ' at both preference profiles. Let profile $\succ'^h$ be obtained from $\succ^h$ by moving $\sigma'(j)$ for all $j \in I_{\sigma'} - I_\sigma$ up to be the first choice of j. Let $\succ'^{h'}$ be obtained analogously from $\succ^{h'}$. By Maskin monotonicity, $\varphi[\succ'^h]^{-1}(e) = i \neq i' = \varphi[\succ'^{h'}]^{-1}(e)$. Since $\succ'^h$ and $\succ'^{h'} \in \mathbf{P}^*[\sigma', e]$, house e is a brokered* house at σ'.

For an indirect argument for the second part of the proof, suppose that k is not the broker* of e at σ'. Then, by Lemma 9 there exists some other agent $k' \neq k$ who brokers* e at σ'.

Let $\succ' \in \mathbf{P}[\sigma', e, h]$ be arbitrary and $\succ \in \mathbf{P}[\sigma, e, h]$ be such that each agent j in $I_{\sigma'} - I_\sigma$ lists $\sigma'(j)$ as his third choice at $\succ$, each agent in $I - I_{\sigma'}$ lists houses in $H_{\sigma'}$ lower than houses in $H_{\sigma'} - H_\sigma$ at $\succ$, and the rest of the relative rankings of the houses are the same between $\succ$ and $\succ'$. Since k brokers* e at σ and i owns* h at σ, by Lemma 11 $\varphi[\succ](k) = g$ and $\varphi[\succ](i) = e$. Then, by Pareto efficiency, $\varphi[\succ'](j) = \sigma'(j)$ for all $j \in I_{\sigma'} - I_\sigma$, and thus, by Maskin monotonicity, $\varphi[\succ'] = \varphi[\succ]$. Now, $\varphi[\succ'](k) = h$, however, this contradicts the fact that agent $k' \neq k$ brokers* e at σ' and thus, $\varphi[\succ'](k') = h$. Therefore, k brokers* e at σ', as well. **QED**

Lemma 17. ***(R6; Consolation for lost control rights*)*** *Let $\sigma \in \overline{\mathcal{M}}$, $i, j \in \overline{I_\sigma}$, and $g, h \in \overline{H_\sigma}$ be such that $i \neq j$ and $g \neq h$, i controls* h and j controls* g at σ. Then i owns* g at $\sigma' = \sigma \cup \{(j, h)\}$.*

Proof of Lemma 17. First consider the case i brokers* h and j owns* g at σ. By Lemmas 10 and 11 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle h, g, ...\rangle$, $\succ_j \in \langle h, ...\rangle$, we have $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Then, by Maskin monotonicity, for any $\succ' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ](i) = g$, i.e., i owns* g at σ'.

Next consider the case i owns* h and j owns* g at σ. For all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle g, h...\rangle$, $\succ_j \in \langle h, g, ...\rangle$, strategy-proofness for i implies $\varphi[\succ](i) \succeq_i h$ as otherwise $\succ'_i \in \langle h, ...\rangle$, $\varphi[\succ'_i, \succ_{-i}](i) = h$. Similarly, $\varphi[\succ](j) \succeq_j g$. Pareto efficiency implies $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Hence, by Maskin monotonicity, for all $\succ'' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ''](i) = g$, i.e., i owns* g at σ'.

Finally consider the case i owns* h and j brokers* g at σ. By Lemmas 10 and 11 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle g, ...\rangle$, $\succ_j \in \langle g, h...\rangle$, we have $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Then, by Maskin monotonicity, for any $\succ' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ](i) = g$, i.e., i owns* g at σ'. **QED**

Lemma 18. ***(R6; Brokered*-to-Owned* House Transition)*** *Let $\sigma \in \overline{\mathcal{M}}$, $k, j, i \in \overline{I_\sigma}$, and $e, g, h \in \overline{H_\sigma}$ be such that $k \neq j$ and $e \neq g$, k brokers* e at σ but not at $\sigma' = \sigma \cup \{(j, g)\}$, and i owns* h at σ. Then i owns* e at σ'.*

Proof of Lemma 18. By Lemmas 10 and 11 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle e, ...\rangle$, $\succ_k \in \langle e, h, ...\rangle$, we have $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$. Since $\mathbf{P}[\sigma'] \subset \mathbf{P}[\sigma]$, Proposition 4 implies that either i owns* e at σ' or k brokers* e at σ'. The latter is not true, by an assumption made in the lemma; hence i owns* e at σ'. **QED**

D.3 The TC Mechanism Defined by the Starred Control Rights Structure Equals φ

We showed above that the starred control rights structure (c, b) is well defined and consistent (satisfies R1-R6). We will now close the proof of Theorem 2 by showing that the resulting TC mechanism, $\psi^{c,b}$, maps preferences to outcomes in the same way as φ does.

Fix $\succ \in \mathbf{P}$. We will show that $\varphi[\succ] = \psi^{c,b}[\succ]$ proceeding by induction on rounds of $\psi^{c,b}$. Let I^r be the set of agents removed in round r of $\psi^{c,b}$. For each agent $i \in I^r$, there is a unique house that points to him and is removed in the same cycle as i; let us denote this house h_i. Let us construct the following preference profile $\succ^*$ by modifying $\succ$.

- If $\psi^{c,b}[\succ](i) = h_i$, then $\succ_i^* = \succ_i$.
- If $\psi^{c,b}[\succ](i) \neq h_i$ and if no brokered house was removed in the same cycle as i or the brokered house was assigned to i, then we construct $\succ_i^*$ from $\succ_i$ by moving h_i just after $\psi^{c,b}[\succ](i)$ (we do not change the ranking of other houses).
- If i is removed as owner and a brokered house e^r was removed in the same cycle as i but not assigned to i, then we construct $\succ_i^*$ from $\succ_i$ by moving e^r just after $\psi^{c,b}[\succ](i)$ and moving h_i just after e^r.
- If a broker k^r was removed in a cycle

$$h_{i^1} \rightarrow i^1 \rightarrow h_{i^2} \rightarrow i^2 \rightarrow ...h_{i^n} \rightarrow i^n \rightarrow e^r \rightarrow k^r \rightarrow h_{i^1},$$

  then we construct $\succ_{k^r}^*$ from $\succ_{k^r}$ by moving h_{i^n} just below h_{i^1}.

We will show that

$$\varphi[\succ^*](i) = \psi^{c,b}[\succ^*](i) \qquad \forall i \in \cup_{s \leq r} I^s, \qquad \forall r = 0, 1, 2, ... \tag{1}$$

by induction over the round r of $\psi^{c,b}$. The claim is trivially true for $r = 0$. Fix round $r \geq 1$ and let σ^{r-1} be the matching fixed before round r (in particular, $\sigma^0 = \varnothing$). For the inductive step, assume that

$$\varphi[\succ^*](i) = \psi^{c,b}[\succ^*](i) \qquad \forall i \in \cup_{s \leq r-1} I^s = I_{\sigma^{r-1}}$$

We will prove that the same expression holds for agents in I^r using the following three claims.

Claim 1. $\varphi[\succ^*](i) \succeq_i^* h_i$ for all owners $i \in I^r$.

Proof of Claim 1. Let $\succ' \in \mathbf{P}[\sigma^{r-1}, h_i]$ be a preference profile such that the relative ranking of all houses in $H - H_{\sigma^{r-1}} - \{h_i\}$ in $\succ_j'$ is the same as in $\succ_j^*$ for all $j \in (I - I_{\sigma^{r-1}}) - \{i\}$, and let $\succ'' \in \mathbf{P}[\sigma^{r-1}]$ be a preference profile such that the relative ranking of all houses in $H - H_{\sigma^{r-1}}$ in $\succ_j''$ is the same as in $\succ_j^*$ for all $j \in (I - I_{\sigma^{r-1}}) - \{i\}$.

By Maskin monotonicity,

$$\varphi[\succ^*] = \varphi[\succ''_{(I - I_{\sigma^{r-1}}) - \{i\}}, \succ'_{I_{\sigma^{r-1}}}, \succ_i^*].$$

Furthermore, by definition h_i is owned by i at σ^{r-1} under $\psi^{c,b}$ and the construction of the control right structure (c, b) from φ means that h_i is owned* by i in φ. Thus,

$$\varphi[\succ'](i) = h_i,$$

and no agent $j \in (I - I_{\sigma^{r-1}}) - \{i\}$ gets h_i at $\varphi[\succ']$. These agents also do not get houses from $H_{\sigma^{r-1}}$ at $\varphi[\succ']$. Maskin monotonicity thus implies that

$$\varphi[\succ'] = \varphi[\succ''_{(I - I_{\sigma^{r-1}}) - \{i\}}, \succ'_{I_{\sigma^{r-1}} \cup \{i\}}].$$

Taken together the first above-displayed equation, the strategy-proofness of φ, the third and second above-displayed equation give us

$$\varphi[\succ^*](i) = \varphi[\succ''_{(I - I_{\sigma^{r-1}}) - \{i\}}, \succ'_{I_{\sigma^{r-1}}}, \succ_i^*](i) \ \succeq_i^* \ \varphi[\succ''_{(I - I_{\sigma^{r-1}}) - \{i\}}, \succ'_{I_{\sigma^{r-1}} \cup \{i\}}](i) = \varphi[\succ'](i) = h_i.$$

QED

Claim 2. If $i \in I^r$ and no brokered house was removed in the cycle of i, then $\varphi[\succ^*](i) = \psi^{c,b}[\succ^*](i)$.

Proof of Claim 2. The inductive assumption implies that all houses better than $\psi^{c,b}[\succ^*](i)$

are already given to other agents; hence

$$\psi^{c,b}[\succ^*](i) \succeq_i^* \varphi[\succ^*](i).$$

For an indirect argument, suppose $\varphi[\succ^*](i) \neq \psi^{c,b}[\succ^*](i)$. Then, Claim 1 and the construction of $\succ^*$ imply that

$$\varphi[\succ^*](i) = h_i.$$

Let

$$h_i \to i \to h_{i^2} \to i^2 \to ... \to h_{i^n} \to i^n \to h_i$$

be the cycle in which i is removed under $\psi^{c,b}[\succ^*]$. From

$$\varphi[\succ^*](i) = h_i = \psi^{c,b}[\succ^*](i^n),$$

we conclude that $\varphi[\succ^*](i^n) \neq \psi^{c,b}[\succ^*](i^n)$, and Claim 1 and the construction of $\succ^*$ imply that

$$\varphi[\succ^*](i^n) = h_{i^n} = \psi^{c,b}[\succ^*](i^{n-1}).$$

As we continue iteratively, we obtain that

$$\varphi[\succ^*](j) = h_j$$

for all $j \in \{i, i^2, ..., i^n\}$. Hence, the matching obtained by assigning $\psi^{c,b}[\succ^*](j)$ to each agent $j \in \{i, i^2, ..., i^n\}$ and $\varphi[\succ^*](j)$ to each agent $j \in I - \{i, i^2, ..., i^n\}$ Pareto dominates $\varphi[\succ^*]$ at $\succ^*$, contradicting that $\varphi[\succ^*]$ is Pareto efficient. QED

Claim 3. If $i \in I^r$ and a brokered house was removed in the cycle of i, then $\varphi[\succ^*](i) = \psi^{c,b}[\succ^*](i)$.

Proof of Claim 3. Let $e \equiv h_{i^0}$ be the brokered house and $k \equiv i^0$ be the broker at σ^{r-1}. Let

$$h_{i^1} \to i^1 \to h_{i^2} \to ... \to i^n \to e \to k \to h_{i^1}$$

be the cycle in which they are removed in round r of $\psi^{c,b}$. By the inductive assumption, for each i^ℓ, $\ell = 1, ..., n$, all houses better than $h_{i^{\ell+1}}$ are given to other agents before round r. Hence, Claim 1 implies that

$$\varphi[\succ^*](i^\ell) \in \{h_{i^{\ell+1}}, e, h_{i^\ell}\}, \qquad \ell = 1, ..., n \tag{2}$$

Recall that $h^*_{i^{\ell+1}} \succeq^*_{i^\ell} e \succ_{i^\ell} h_{i^\ell}$. We prove Claim 3 in two steps:

Step 1. Let us show that $\varphi[\succ^*](i^n) = e = \psi^{c,b}[\succ^*](i^n)$. Suppose not. Then, $\varphi[\succ^*](i^n) \neq e$. Since $e = h_{i^{n+1}}$, the above displayed inclusion gives us $\varphi[\succ^*](i^n) = h_{i^n}$. Thus, the above displayed inclusion tells us that $\varphi[\succ^*](i^\ell) \in \{e, h_{i^\ell}\}$ for $\ell = n-1$. We cannot have $\varphi[\succ^*](i^\ell) = e$ as it would not be Pareto efficient because agents i^n and i^{n-1} would be better off by swapping their allocations. Thus, $\varphi[\succ^*](i^\ell) = h_{i^\ell}$. Iterating this last argument we show that

$$\varphi[\succ^*](i^\ell) = h_{i^\ell}, \qquad \ell = n, n-1, ..., 1.$$

Let us construct an auxiliary preference profile $\succ' \in \mathbf{P}[\sigma^{r-1}]$ from $\succ^*$ by pushing up $\sigma^{r-1}(i)$ in preferences of agents $i \in I_{\sigma^{r-1}}$, pushing down houses matched at σ^{r-1} in preferences of agents $i \in I - I_{\sigma^{r-1}}$, and pushing down h^2 in preferences of i^1 while preserving the relative ranking of houses otherwise. By above observations, $\succ'$ is a φ-Maskin-monotone transformation of $\succ^*$, and hence $\varphi[\succ^*] = \varphi[\succ']$. Notice that agent i^1 owns* h^1 at σ^{r-1} in φ and agent k brokers* e at σ^{r-1} in φ (by construction of $\psi^{c,b}$ in which i^1 is the owner of h^1 and agent k is the broker of e at σ^{r-1}). Because $\succ'_k \in \mathbf{P}_k[\sigma^{r-1}, h_{i^1}, ...] \cup \mathbf{P}_k[\sigma^{r-1}, e, h_{i^1}, ...]$, and $\succ'_{i^1} \in \mathbf{P}_k[\sigma^{r-1}, e, h_{i^1}, ...]$, we get $\varphi[\succ'](i^1) = e$ and thus $\varphi[\succ^*](i^1) = e$ contrary to the above displayed equations. This contradiction concludes Step 1.

Step 2. Let us show that

$$\varphi[\succ^*](i^\ell) = h_{i^{\ell+1}} = \psi^{c,b}[\succ^*](i^\ell) \qquad \forall\, \ell \in \{0, ..., n-1\}\,.$$

By way of contradiction, suppose there exists some $\ell \in \{0, ..., n-1\}$ such that $\varphi[\succ^*](i^\ell) \neq h_{i^{\ell+1}}$. Then, inclusion 2 and Step 1 imply that $\varphi[\succ^*](i^\ell) = h_{i^\ell}$. Thus, $\varphi[\succ^*](i^{\ell-1}) \neq h_{i^{(\ell-1)+1}}$. Iterating this argument we show

$$\varphi[\succ^*](i^m) = h_{i^m} \quad m = \ell-1, \ell-2, ..., 1.$$

Let us construct an auxiliary preference profile $\succ' \in \mathbf{P}[\sigma^{r-1}]$ from $\succ^*$ by pushing up $\sigma^{r-1}(i)$ in preferences of agents $i \in I_{\sigma^{r-1}}$, pushing down houses matched at σ^{r-1} in preferences of agents $i \in I - I_{\sigma^{r-1}}$, and pushing down h^1 in preferences of $i^0 \equiv k$ while preserving the relative ranking of houses otherwise.

The above-displayed equations imply $\varphi[\succ^*](k) \neq h_{i^1}$, and thus $\succ'$ is a φ-Maskin-monotone transformation of $\succ^*$, and hence $\varphi[\succ^*] = \varphi[\succ']$. Notice that agent i^n owns* h^n at σ^{r-1} in φ and agent k brokers* e at σ^{r-1} in φ (by construction of $\psi^{c,b}$ in which i^1 is the owner of h^1 and agent k is the broker of e at σ^{r-1}). Because $\succ'_k \in \mathbf{P}_k[\sigma^{r-1}, h_{i^n}, ...] \cup \mathbf{P}_k[\sigma^{r-1}, e, h_{i^n}, ...]$, and $\succ'_{i^1} \in \mathbf{P}_k[\sigma^{r-1}, e, h_{i^n}, ...]$, we get $\varphi[\succ'](k) = h_{i^n}$ and thus

$$\varphi[\succ^*](k) = h_{i^n}.$$

In consequence, inclusion 2 and Step 1 imply that $\varphi[\succ^*](i^{n-1}) = h_{i^{n-1}}$. Thus, $\varphi[\succ^*](i^{n-2}) \neq h_{i^{n-1}}$. Iterating this argument we show

$$\varphi[\succ^*](i^m) = h_{i^m} \quad m = n-1, n-2, ..., 1.$$

The above-displayed equations and Step 1 imply that $\varphi[\succ^*]$ is Pareto dominated by the allocation in which each agent i^m, $m = 0, ..., n-1$, gets house h^{m+1}, and all other agents get their $\varphi[\succ^*]$ houses. This contradiction concludes Step 2, and proves Claim 3. QED

Claims 2 and 3 show that $\varphi[\succ^*](i) = \psi^{c,b}[\succ^*](i)$ for all $i \in I^r$. This completes the inductive proof of equations (1). Now, the theorem follows from

$$\psi^{c,b}[\succ] = \psi^{c,b}[\succ^*], \qquad \psi^{c,b}[\succ^*] = \varphi[\succ^*], \qquad \text{and} \qquad \varphi[\succ^*] = \varphi[\succ].$$

The first of these equations follows directly from the construction of $\succ^*$. The second equation is equivalent to equations 1. To prove the third equation, observe that for every agent $i \in I$,

$$\left\{ h \in H : h \succeq_i \underbrace{\psi^{c,b}[\succ](i)}_{=\psi^{c,b}[\succ^*](i) = \varphi[\succ^*]} \right\} = \left\{ h \in H : h \succeq_i^* \underbrace{\psi^{c,b}[\succ](i)}_{=\psi^{c,b}[\succ^*](i) = \varphi[\succ^*]} \right\}.$$

In particular,

$$\{h \in H : h \succeq_i \varphi[\succ^*](i)\} = \{h \in H : h \succeq_i^* \varphi[\succ^*](i)\} \text{ for all } i \in I,$$

and hence $\succ$ is a $\varphi-$monotonic transformation of $\succ^*$. The third equation thus follows from Maskin monotonicity of φ. **QED**

E Supplementary Appendix: Proof of Theorem 5

The argument for the Pareto efficiency of TC remains the same as in the TC example of Section 3.2. As before, group strategy-proofness is equivalent to individual strategy-proofness and non-bossiness.

Lemma 19. *In the environment with outside options, a mechanism is group strategy-proof if and only if it is individually strategy-proof and non-bossy.*

The proof follows word-by-word the proof of Lemma 1 in Pápai (2000). **QED**

Our arguments for individual strategy-proofness and non-bossiness go through with two modifications. First, when in the proof of Theorem 1 we assume that an agent is matched with a house, we should now substitute "a house or the agent's outside option." If the agent is matched in a cycle of a length above 1, we can then conclude that the agent is indeed matched with a house. Second, in some steps of the proof we consider separately the case when a broker is matched with his outside option. We handle these cases below. This allows us to assume this case away in the relevant parts of the original proof.

Consider the proof of individual strategy-proofness. In Case 1: $s \leq s'$, let i be a broker of house e and under $\succ_i$ leaves with his outside option in round s. Since the same houses are matched under $\succ_i$ and $\succ_i'$, under $\succ_i'$ the best the broker can do is to leave either with his outside option, or – if he prefers the brokered house e to his outside option – to leave with the brokered house e. We need to prove that the latter cannot happen. By Lemma 3, in round s of TC under $\succ_i'$, agent i is a broker of e and there is an owner j whose first preference is e. For i to be matched with e, he would need to lose the brokerage right but by R5-R6 if this happens then j becomes the owner of e, and is then matched with it, ending the argument for Case 1. In Case 2: $s > s'$, if i be a broker of house e matched with his outside option under $\succ_i'$, then submitting this preference profile cannot be better than submitting the true profile $\succ_i$, as under any profile agent i is matched at least with his outside option.

Consider the proof of non-bossiness. We run the same induction as in the proof without outside options. In the initial step of the induction, consider the additional case when i_* is a broker and is matched with his outside option at time s under $\succ$. By assumption i_* is matched with his outside option under $\succ'$ and the inductive hypothesis is true. In the inductive step, consider the additional case in which i^1 is a broker and is matched with his outside option at time $r > s$ under $\succ$ (handling this case separately allows us to assume this case away in all claims of the inductive step). By the inductive assumption, there is an r^* such that $\sigma^{r-1}[\succ] \subseteq \sigma^{r^*}[\succ']$. At $\sigma^{r-1}[\succ]$, i^1 brokers a house h and all houses other than h that i^1 prefers to his outside option are matched. Since i^1 gets at least his outside option, he either gets his outside option (and the inductive step is true) or he gets h. In the latter case, as in the proof of individual strategy-proofness, at $\sigma^{r-1}[\succ]$, there is an owner j whose top preference is h. He remains unmatched as long as h is unmatched. Since for i^1 to obtain h he would need to lose his brokerage right, conditions R5-R6 imply that j would get ownership over h, and would match with h. Hence i^1 cannot be matched with h and is matched with his outside option.

To prove that any group strategy-proof and efficient mechanism is TC we follow the same steps as in the proof of Theorem 2 with one important modification. For $\sigma \in \overline{\mathcal{M}}$, $n \geq 0$ and $h^1, h^2, ..., h^n \in \overline{H_\sigma}$, and $i \in I$, we redefine $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ to be the set of preferences $\succ_i$ of agent i such that

- if $i \in I_\sigma$, then
$$\sigma(i) \succ_i g \text{ for all } g \in H - \{\sigma(i)\},$$
- if $i \in \overline{I_\sigma}$, then
$$h^1 \succeq_i h^2 \succeq ... \succeq_i h^n \succ_i y_i \succ g \text{ for all } g \in H_\sigma.$$

In particular, the definitions of ownership* and brokerage* are repeated word-by-word, but the meaning of $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ is changed as above. With this modification, the proof goes through. **QED**