

Pycia, Marek; Ünver, M. Utku

Article

Incentive compatible allocation and exchange of discrete resources

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Pycia, Marek; Ünver, M. Utku (2017) : Incentive compatible allocation and exchange of discrete resources, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 12, Iss. 1, pp. 287-329, <https://doi.org/10.3982/TE2201>

This Version is available at:

<https://hdl.handle.net/10419/197190>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>

Incentive compatible allocation and exchange of discrete resources

MAREK PYCIA

Department of Economics, University of California

M. UTKU ÜNVER

Department of Economics, Boston College

The allocation and exchange of discrete resources, such as transplant organs, public housing, dormitory rooms, and many other resources for which agents have single-unit demand, is often conducted via direct mechanisms without monetary transfers. Incentive compatibility and efficiency are primary concerns when designing such mechanisms. We construct the full class of group strategy-proof and Pareto-efficient mechanisms and show that each of them can be implemented by endowing agents with control rights over resources. This new class, which we call trading cycles, contains new mechanisms as well as known mechanisms such as top trading cycles, serial dictatorships, and hierarchical exchange. We illustrate how one can use our construction to show what can and what cannot be achieved in a variety of allocation and exchange problems, and we provide an example in which the new trading-cycles mechanisms are more Lorenz equitable than all previously known mechanisms.

KEYWORDS. Group strategy-proofness, Pareto efficiency, matching, no-transfer allocation and exchange, single-unit demand.

JEL CLASSIFICATION. C78, D78.

1. INTRODUCTION

The central concerns in the design of allocation mechanisms are the incentives of participants and the efficiency of outcomes. In this paper, we construct and characterize

Marek Pycia: pycia@ucla.edu

M. Utku Ünver: unver@bc.edu

First draft: December 2007. For their comments, we thank Andrew Atkeson, Salvador Barberà, Haluk Ergin, Manolis Galenianos, Ed Green, Matthew Jackson, Philippe Jehiel, Onur Kesten, Fuhito Kojima, Sang-Mok Lee, George Mailath, Vikram Manjunath, Szilvia Pápai, Al Roth, Tayfun Sönmez, Qianfeng Tang, William Thomson, Özgür Yılmaz, William Zame, Yongchao Zhang, four anonymous referees, and participants of numerous seminars and conferences in years 2007–2014. Omer Ali, Kenny Mirkin, and Kyle Woodward provided excellent research assistance. Ünver gratefully acknowledges the support of the National Science Foundation through Grants SES 0338619 and SES 0616689, and the Microsoft Research Lab in New England. Earlier drafts of this paper included results on individual strategy-proof and Arrowian-efficient mechanisms, on outside options, and on Lorenz dominance; these results will form the core of companion papers.

Copyright © 2017 The Authors. Theoretical Economics. The Econometric Society. Licensed under the Creative Commons Attribution-NonCommercial License 3.0. Available at <http://econtheory.org>. DOI: 10.3982/TE2201

the entire class of group strategy-proof and Pareto-efficient direct mechanisms in environments without transfers in which agents have single-unit demands.¹ These environments include the allocation and exchange of transplant organs, such as kidneys (see Roth et al. 2004), the allocation of public housing, and the allocation of dormitory rooms (see Abdulkadiroğlu and Sönmez 1999).

We study a finite group of agents, each of whom would like to consume a single indivisible object, which we sometimes refer to as a house, using the terminology coined by Shapley and Scarf (1974). We allow for objects brought to the allocation by individual agents as well as by a social planner. Agents have strict preferences over objects and are indifferent over how objects are allocated to other agents. We focus on mechanisms that match objects to agents.

A mechanism is group strategy-proof if no group of agents can jointly manipulate their reports so that all of them weakly benefit from this manipulation, while at least one agent in the group strictly benefits. Group strategy-proofness matters because coordinated reporting to a mechanism has been documented in relevant environments such as organ allocation and exchange.² Furthermore, coordinated reporting is effectively the only way a group of agents can manipulate the allocation in many of the relevant environments; for instance, without approval from the school, two students cannot trade their dormitory rooms *ex post*. A mechanism is Pareto efficient if its allocations are; an allocation is Pareto efficient if no other allocation makes every agent at least as well off and at least one agent strictly better off.

Our main results show (i) that in a group strategy-proof and Pareto-efficient mechanism, each object is either owned or brokered (in the sense made precise below), and (ii) that all such mechanisms can be implemented by having the owners and brokers recursively swap the objects they control. We refer to the associated algorithm and the resulting class of mechanism as *trading cycles*. Our construction of trading cycles builds on David Gale's top trading cycles (reported by Shapley and Scarf 1974) and on Pápai's (2000) hierarchical exchange.

To get a sense of trading cycles, consider two examples in an environment in which there are as many objects as agents and each agent initially controls one object. One example is Gale's top trading cycles. Gale's algorithm matches agents and objects in a sequence of rounds. In each round, each object points to the agent who controls it, and each agent points to his most preferred unmatched object. Because there is a finite

¹The prior theoretical literature on single-unit-demand allocation without transfers has focused on characterizing mechanisms that are strategy-proof and efficient alongside other properties; see below for examples of such characterizations. In contrast, our main results do not rely on any additional assumptions. The restriction to direct mechanisms is justified by the revelation principle.

²In kidney exchange, both the transplant centers and doctors occasionally game the medical system for the benefit of their patients (see Warmbir 2003, Sönmez and Ünver 2006, 2013, Ashlagi and Roth 2014). Coordinated reporting has also been observed in other related settings, such as school choice (see Pathak and Sönmez 2008). In our setting, group strategy-proofness is equivalent to the lack of manipulation opportunities for groups of two agents. Group strategy-proof mechanisms are immune to manipulation regardless of the information the agents possess, they impose minimal costs of searching for and processing strategic information, and they do not discriminate among agents based on their access to information and ability to strategize.

number of agents, there exists at least one pointing cycle in which an agent, say agent 1, points to an object, say object A; the agent who controls object A points to object B, etc.; and finally the last agent in the cycle points to the object controlled by agent 1.³ The procedure then matches each agent in each pointing cycle with the object to which he points. The pointing cycles thus become cycles of trading. Rounds are repeated until no agents and objects are left unmatched. Roth (1982) showed that the resulting mechanism is group strategy-proof and Pareto efficient. Another example of trading cycles obtains when we take one of the agents—we call him a broker—and change the way he can trade the object he controls—which we call the brokered object or the brokered house. We do so by running the same algorithm as above with the additional requirement that we have the broker point to his most preferred unmatched object that is different from the brokered object. The resulting mechanism remains group strategy-proof and Pareto efficient.⁴

While we described examples of trading cycles in the environment where each agent initially controls one object, the algorithm can be used in more general environments, for instance, when all objects are initially owned by a social planner. In such environments, to run top trading cycles, we need to specify a controlling agent for each object at each round (see Abdulkadiroğlu and Sönmez 1999, Pápai 2000); to run trading cycles, we additionally need to specify for each object whether the agent who controls it is an owner or a broker.

Surprisingly, adding brokers and brokered objects to the previous constructions of top trading cycles gives us the full class of group strategy-proof and efficient mechanisms. The key insight behind this characterization is that in every group strategy-proof and Pareto-efficient mechanism, each object is either owned or brokered in the following sense: Either there is a unique agent who obtains this object whenever he ranks it as his top choice (we call such an agent an owner) or else there is a unique agent who obtains his second choice whenever all agents rank the object mentioned earlier as their top choice (we call such an agent a broker). Owners and brokers defined in this way correspond to the owners and brokers in the initial round of trading cycles. The analogue of this insight remains true for other rounds. Suppose that some agents are already matched, fix a reference unmatched object, and consider all preference profiles in which matched agents put their outcomes at the top of their ranking, while the unmatched agents put the reference object at the top of their ranking and the matched objects at the bottom of their ranking. We prove that in every group strategy-proof and Pareto-efficient mechanism, either there is a unique unmatched agent who obtains the referenced object at all the above preference profiles, or else there is a unique agent who obtains his second choice at all the above preference profiles. In the former case we call this unique agent an owner; in the latter case we call the unique agent a broker of the reference object at the partial matching considered. In every round of trading cycles, the

³The pointing cycles might be short (agent 1 points to object A, which points back to agent 1) or might involve many agents.

⁴The mechanisms with brokers remain Pareto efficient because the broker is assigned something other than his top choice only if another agent wants this top choice. Intuitively, these mechanisms are strategy-proof because the set of objects any agent can obtain by changing his or her preference ranking is monotonically increasing (in terms of inclusion) with each round of trading cycles.

owners and brokers in the trading-cycles algorithm correspond exactly to the owners and brokers described above at the partial matching obtained in previous rounds. This insight allows us to construct the trading-cycle representation for every group strategy-proof and Pareto-efficient mechanism.⁵

Knowing that all group strategy-proof and Pareto-efficient mechanisms may be represented as trading cycles allows us to reduce many otherwise hard problems to the simple analysis of trading cycles. In this sense, the role trading-cycles mechanisms play in the single-unit demand no-transfer environments we study can be compared to, for instance, the role that the mechanisms of Vickrey (1961), Clarke (1971), and Groves (1973) play in environments with transfers and quasi-linear utilities (see Green and Laffont 1977, and Holmström 1979). Other characterizations of efficient and strategy-proof mechanisms that are nondictatorial have been obtained by Barberà et al. (1997) for sharing a perfectly divisible good among agents with single-peaked preferences over their shares (see Sprumont 1991), and by Barberà et al. (1993) for committee-voting problems with single-peaked preferences (see Moulin 1980).

To illustrate the usefulness of our reduction approach, we use it to obtain further insights into allocation and exchange, and to derive some of the canonical prior results as immediate corollaries. First, we show that the class of group strategy-proof, Pareto efficient, and individually rational mechanisms equals the class of trading cycles that (i) may be represented by a control-rights structure in which each agent is given control rights over all objects from his endowment, and (ii) none of these agents is a broker. In particular, we show that when each agent has a private endowment, then this class of mechanisms can be represented via Pápai's (2000) hierarchical exchange.⁶ Pápai's hierarchical-exchange mechanisms 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 (without brokers) that—for each round of the mechanism and each unmatched object—determines the agent to whom the object points.⁷

Second, we derive Pápai's (2000) and Svensson's (1999) canonical results as immediate corollaries of our characterization. Pápai showed that her hierarchical-exchange mechanisms are precisely these group strategy-proof and Pareto-efficient mechanisms that satisfy an additional technical property that she refers to as reallocation-proofness.⁸ Her result is a corollary of our characterization because trading cycles with brokers do

⁵The trading-cycles mechanism constructed in this way is equivalent to the mechanism being represented, that is the two mechanisms produce the same outcome for every preference profile. We verify this equivalence via a recursive argument on the rounds of trading cycles.

⁶In the special case of our setting in which there are as many objects as agents and each agent is endowed with exactly one object, this corollary of our main results is implied by an earlier result of Ma (1994) that top trading cycles is the unique strategy-proof, Pareto-efficient, and individually rational mechanism. Following our work, Tang and Zhang (2016) extended Ma's result by strengthening individual rationality. See also Pápai (2007) for the analysis of multi-unit demand, and Pycia (2016) for the analysis on networks.

⁷Redefining Pápai's hierarchical-exchange mechanisms in terms of the structure of ownership rights is one of our contributions to the rich literature that follows her seminal work. This redefinition has become standard; see, for instance, Tang and Zhang (2016). We would like to thank a referee for drawing our attention to this contribution.

⁸A mechanism is reallocation-proof in the sense of Pápai 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,

not satisfy Papai's reallocation-proofness property. Svensson showed that a mechanism is neutral and group strategy-proof if and only if it is a serial dictatorship; neutrality means that a mechanism is invariant to any renaming of objects. His result is a corollary because neutral and group strategy-proof mechanisms are Pareto efficient, and because, to be neutral, a trading-cycle mechanism must be a serial dictatorship.⁹

Our paper is the first to recognize the role of brokers in allocation and exchange with no transfers. While the main role played by brokers is to give us the full class of strategy-proof and efficient mechanisms, brokers can also be useful in some mechanism-design settings. We close the paper by providing an example in which the trading cycle with one broker described above is the most equitable allocation mechanism, and in particular it is strictly more equitable—in the sense of Lorenz dominance—than any hierarchical-exchange mechanism.

2. MODEL

Let I be a set of *agents* and let H be a set of *objects* that we often refer to as *houses*; these sets are fixed throughout the analysis. We use 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 by \succ_i .¹⁰ The set of all strict preference relations of agent i is denoted \mathbf{P}_i and the set of *preference profiles* \mathbf{P} is the Cartesian product $\times_{i \in I} \mathbf{P}_i$. For any $J \subseteq I$, $\succ_J = (\succ_i)_{i \in J}$ is the restriction of $\succ \in \mathbf{P}$ to J , and $\mathbf{P}_J = \times_{i \in J} \mathbf{P}_i$ is the set of all such restricted profiles.

These primitives define the model of house allocation. In Section 5, we introduce the possibility that some agents have initial rights over houses; the analysis of the resulting model of allocation and exchange turns out to be straightforward once we understand house allocation. Throughout, we assume that there are no outside options; thus there are at least as many houses as agents.¹¹

We study ways to allocate houses to agents. We call an outcome of such an allocation a *matching*. To define it, we start with a more general concept that we use frequently.

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. As Pápai observes, reallocation-proofness-type properties obtained by dropping condition (ii) or by allowing swaps among more than two agents conflict with group strategy-proofness and Pareto efficiency. We do not use reallocation-proofness in our results.

⁹All allocation papers cited above, and the literature in general, share with our paper the assumption that agents have strict preferences. This is the standard modeling assumption for several reasons. First, 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. Second, and more fundamentally, Ehlers (2002) shows that “one cannot go much beyond strict preferences if one insists on efficiency and group strategy-proofness”; when agents can be indifferent among objects, no mechanism is group strategy-proof and Pareto efficient. See Bogomolnaia et al. (2005) for further exploration of the domain with possible indifferences.

¹⁰By \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 $g \succ_i h$.

¹¹The analysis of outside options from earlier drafts of this paper, including the case of fewer houses than agents, will form the core of a separate paper; for more details, see the Supplement, which is available in a supplementary file on the journal website, <http://econtheory.org/supp/2201/supplement.pdf>. Our insights remain true in the presence of outside options.

A *submatching* on $J \subseteq I$ is a one-to-one function $\sigma : J \rightarrow H$. A submatching allocates a subset of houses to a subset of agents so that no two different agents obtain the same house; $\sigma(i)$ is the house matched with agent $i \in J$, and $\sigma^{-1}(h)$ is the agent matched with house $h \in \sigma(J)$. The set of submatchings is denoted by \mathcal{S} . For each $\sigma \in \mathcal{S}$, the set of matched agents is denoted by I_σ , and the set of matched houses is denoted by H_σ . We write \overline{I}_σ for $I - I_\sigma$ and \overline{H}_σ for $H - H_\sigma$. For any $h \in H$, \mathcal{S}_{-h} is the set of submatchings $\sigma \in \mathcal{S}$ at which house h is unmatched, that is, that belongs to \overline{H}_σ . Identifying a function with its graph, we can talk about submatchings as sets of agent–house pairs; these sets are ordered by inclusion. A *matching* is a maximal submatching; that is, $\mu \in \mathcal{S}$ is a matching if $I_\mu = I$. The set of matchings is denoted by \mathcal{M} , and $\overline{\mathcal{M}} = \mathcal{S} - \mathcal{M}$.

A *mechanism* is a mapping $\varphi : \mathbf{P} \rightarrow \mathcal{M}$ that assigns a matching for each preference profile.¹²

2.1 Strategy-proofness and Pareto efficiency

A matching is Pareto efficient if no other matching makes each agent 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 assigns a Pareto-efficient matching to every preference profile.

A mechanism is *individually strategy-proof* if truthful revelation of preferences is a weakly dominant strategy for any agent, that is, for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ'_i \in \mathbf{P}_i$ such that $\varphi[\succ'_i, \succ_{-i}](i) \succ_i \varphi[\succ](i)$. A mechanism is *group strategy-proof* if there is no group of agents who can misreport their preferences in such a way that each agent in the group gets a weakly better house and at least one agent in the group gets a strictly better house irrespective of the preference reports of the agents not in the group. 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 $\varphi[\succ'_J, \succ_{-J}](i) \succeq_i \varphi[\succ](i)$ for all $i \in J$ and $\varphi[\succ'_J, \succ_{-J}](j) \succ_j \varphi[\succ](j)$ for at least one $j \in J$.

Group strategy-proofness has two properties that we use throughout. First, in the environment we study, group strategy-proofness is equivalent to the conjunction of two noncooperative properties: individual strategy-proofness and nonbossiness.¹³ Nonbossiness (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 *nonbossy* if, for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ'_i \in \mathbf{P}_i$ such that $\varphi[\succ'_i, \succ_{-i}](i) = \varphi[\succ](i)$ and $\varphi[\succ'_i, \succ_{-i}] \neq \varphi[\succ]$. The following lemma is due to Pápai (2000).

LEMMA 1 (Pápai 2000). *A mechanism is group strategy-proof if and only if it is individually strategy-proof and nonbossy.*

¹²We study direct mechanisms. By the revelation principle, this is without loss of generality.

¹³Both of these properties are noncooperative in the sense that they relate a mechanism's outcomes under two scenarios when a single agent makes unilateral preference revelation deviations.

Second, in the environment we study, group strategy-proofness is equivalent to Maskin monotonicity (Maskin 1999). A preference profile $\succ' \in \mathbf{P}$ is a φ -monotonic transformation of $\succ \in \mathbf{P}$ if $h \succeq'_i \varphi[\succ](i)$ implies $h \succeq_i \varphi[\succ](i)$ for all $i \in I$ and $h \in H$. 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. A mechanism φ is *Maskin monotonic* if $\varphi[\succ'] = \varphi[\succ]$ whenever $\succ' \in \mathbf{P}$ is a φ -monotonic transformation of $\succ \in \mathbf{P}$. The following lemma was proven by Takamiya (2001) for a subset of the problems we study; his proof can be extended to our more general setting.

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

3. PRELIMINARIES

To set the stage for our trading-cycles mechanism, we redefine Pápai's (2000) hierarchical-exchange, or top-trading-cycles, mechanisms, and discuss an example of a group strategy-proof and Pareto-efficient mechanism that is different from all top-trading-cycles mechanisms. We then introduce the concept of control rights and define trading cycles (TC).

3.1 Top trading cycles

The well known top-trading-cycles (TTC) algorithm has been extended to house allocation by Abdulkadiroğlu and Sönmez (1999) and generalized by Pápai (2000). The class of mechanisms presented in this subsection is identical to Pápai's hierarchical exchange. Our presentation, however, is novel, and it aims to simplify Pápai's construction and to introduce some of the terminology we later use to describe trading cycles.¹⁴

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 are not changed in subsequent rounds. 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 agent who owns it, and each unmatched agent points to his most preferred unmatched house. 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 $k = 1, 2, \dots$, 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. When all agents are matched, the algorithm terminates. Because at least one agent–house pair is matched in each round, the algorithm terminates after finitely many rounds.

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

¹⁴Our reformulation has become standard in the literature; see Tang and Zhang (2016) for a recent example.

DEFINITION 1. A *structure of ownership rights* is a collection of mappings $\{c_\sigma : \overline{H}_\sigma \rightarrow \overline{I}_\sigma\}_{\sigma \in \overline{\mathcal{M}}}$. The structure of ownership rights $\{c_\sigma\}_{\sigma \in \overline{\mathcal{M}}}$ is *consistent* if, for all submatchings $\sigma \subseteq \sigma' \in \overline{\mathcal{M}}$, if an agent $i \in \overline{I}_{\sigma'}$ owns a house $h \in \overline{H}_{\sigma'}$ at σ , then i owns h at σ' .

For each submatching σ and unmatched house h , the structure of ownership rights specifies the unmatched agent $c_\sigma(h)$ who owns it. 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 the agent and the house are unmatched.

Each consistent structure of ownership rights $\{c_\sigma\}_{\sigma \in \overline{\mathcal{M}}}$ determines a *hierarchical-exchange* mechanism of Pápai (2000), and each hierarchical-exchange mechanism can be implemented as TTC with a consistent structure of ownership rights. We thus also refer to hierarchical exchanges as TTC mechanisms. Pápai showed that all TTC mechanisms are group strategy-proof and Pareto efficient, extending an earlier insight of Roth (1982).¹⁵

3.2 Example: Beyond top trading cycles

What might a group-strategy-proof and efficient non-TTC mechanism look like? Consider the following example that builds on the TTC idea.

EXAMPLE 1. Consider three agents i_1, i_2, i_3 , three houses h_1, h_2, h_3 , and an ownership structure that allocates initial ownership of house h_ℓ to agent i_ℓ for $\ell = 1, 2, 3$. Given this structure, we 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 points to his most preferred house among unmatched houses other than h_1 .¹⁶

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 point to h_1 , but agent i_1 points to his second-choice house, h_2 . An exchange cycle forms 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 shows 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, 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.

¹⁵To appreciate the generality of 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. Also, Pápai's fixed endowment mechanisms are TTC mechanisms such that $c_\sigma(h) = c_{\sigma'}(h)$ for all σ and σ' such that $I_\sigma = I_{\sigma'}$.

¹⁶Pápai (2000) gives an example of a non-TTC mechanism. Her construction is different from ours, though the resulting mechanisms are identical. As we show in the next section, the advantage of our construction lies in it being generalizable to the whole class of group strategy-proof and efficient mechanisms.

Roth's (1982) approach to the group strategy-proofness of TTC can be extended to the above mechanism. Alternatively, for each agent and round of the algorithm, define the set of houses *obtainable* by the agent to be the set of houses the agent could obtain in this round by either submitting his true preference ranking or by changing the ranking of unmatched objects. As long as an agent is unmatched, the agent's set of obtainable objects stays the same or becomes larger with each round. Similarly, the set of houses that a group of agents could obtain in a given round by varying the ranking of unmatched objects weakly increases with each round. This monotonicity drives the strategy-proofness of the mechanism. \diamond

The above mechanism is different from all TTC mechanisms. Indeed, 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 when either of these two preference profiles is submitted. The reason is that the TTC ownership structure uniquely determines which agent owns h_1 at the empty submatching, implying that 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, observe that agent i_1 does not have full ownership rights 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 rights over h_1 , however: he can trade h_1 for houses owned by other agents. In our general trading-cycles algorithm, we refer to such weak control rights as brokerage.

Ownership and brokerage can be identified by the mechanism's outcomes. Whenever agent i_2 ranks house h_2 as his top choice, the mechanism in the example assigns him this house irrespective of other agents' rankings. For any mechanism, not necessarily group strategy-proof or efficient, we call an agent who obtains a house whenever he ranks it first an owner of this house (or, an *owner** of this house, when we distinguish between this notion of ownership and the ownership in the algorithms). In the current example i_2 is an owner* of h_2 and i_3 is an owner* of h_3 . Furthermore, whenever all agents rank house h_1 as their top choice, the mechanism in the example assigns agent i_1 his second choice, irrespective of how other houses are ranked. For any mechanism, not necessarily group strategy-proof or efficient, we call an agent who obtains his second choice across all preference profiles in which all agents rank the same house first, a broker of the house that everyone ranks first (or a *broker** of this house when we distinguish between this notion of brokerage and the brokerage in the algorithms). In the current example i_1 is a broker* of h_1 .

In the current example, each house is either owned* or brokered*. In Section 4, we prove that this property holds true for any group strategy-proof and Pareto-efficient mechanism. This allows us to represent any such mechanism through an algorithm similar to the algorithm in the example. In our representation, owners in the algorithm correspond to owners* defined in terms of outcomes and brokers in the algorithm correspond to brokers* defined in terms of outcomes.

3.3 Trading cycles

The TC algorithm matches agents and houses in exchange cycles recursively over a sequence of rounds. TC is more flexible than TTC as it allows two types of intra-round control rights over houses: ownership and brokerage. We add brokerage rights to our definition of the structure of ownership rights as follows.

DEFINITION 2. A *structure of control rights* is a collection of mappings

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

At submatching σ , the function c_σ maps each unmatched house h to the unique unmatched agent $c_\sigma(h)$ who *controls* it. The type of control is determined by b_σ . The agent $c_\sigma(h)$ *owns* h at σ if $b_\sigma(h) = \text{ownership}$ and $c_\sigma(h)$ *brokers* h at σ if $b_\sigma(h) = \text{brokerage}$. In the former case, we call the agent an *owner* and the controlled house an *owned house*. In the latter case, we use the terms *broker* and *brokered house*.

We impose two types of consistency requirements on TC control-rights structures: constraints on brokerage at any given submatching (the *within-round* requirements) and constraints on how the control rights are related across different submatchings (the *across-rounds* requirements). For simplicity, we initially formulate the consistency requirements and the TC algorithm assuming that $|H| > |I|$; we relax this assumption in the Supplement.¹⁷

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 control any other houses at σ .

With requirements R1–R3 in place, we are ready to describe the basic TC algorithm.¹⁸

The TC algorithm matches agents with houses over a finite sequence of rounds $r = 1, 2, \dots$. We set $\sigma^0 = \emptyset$ and recursively construct submatchings σ^{r-1} of agents and houses matched before round r . Each round r consists of the following steps.

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 is a broker at σ^{r-1} , then he points to his most preferred house among houses owned at σ^{r-1} . Every other agent $i \in \overline{I}_{\sigma^{r-1}}$ points to his most preferred house in $\overline{H}_{\sigma^{r-1}}$.

¹⁷This appendix relaxes the assumption that there are more houses than agents and extends the consistency conditions, the TC algorithm, and all our results accordingly.

¹⁸R1 allows for different houses to be brokered at different submatchings, even though there is at most one brokered house at any given submatching. The lemmas of Appendix C.2 show why R1 and R3 are implied by Pareto efficiency and group strategy-proofness. R2 is imposed so as to simplify the basic TC algorithm presented here (it is not needed in the version of the TC algorithm presented in the Supplement). We define R1–R3 (and, later, R4–R6) for all submatchings whether or not they can appear in a round of the TC algorithm. This simplifies the definition of the consistency conditions, but we could equivalently define the control rights only on submatchings that can appear in a round of the TC algorithm.

Step 2. Matching Trading Cycles. There exists $n \in \{1, 2, \dots\}$ and an exchange cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$ in which agents $i^\ell \in \overline{I_{\sigma^{r-1}}}$ point to houses $h^{\ell+1} \in \overline{H_{\sigma^{r-1}}}$ and houses h^ℓ point to agents i^ℓ (here $\ell = 1, \dots, n$ and superscripts are added modulo n). Each agent in each trading cycle is matched with the house he is pointing to.

Step 3. Submatching σ^r is defined as the union of σ^{r-1} and the set of newly matched agent–house pairs. When all agents or all houses are matched, the algorithm terminates and gives matching $\sigma^r \in \mathcal{M}$ as its outcome.

The TC algorithm builds on Gale’s, Abdulkadiroğlu and Sönmez’s, and Pápai’s top-trading-cycles ideas but allows more general trading cycles than top cycles.¹⁹ There exists a trading cycle in each round because the number of agents is finite, each agent points to a unique house, and each house points to a unique agent. This also implies that trading cycles cannot intersect and, hence, which houses are matched in a round is well defined.²⁰ Because we match at least one agent–house pair in every round, and because there are finitely many agents and houses, the algorithm terminates after finitely many rounds. For instance, the algorithm described in the previous section is TC with agent i_1 brokering house h_1 and other agents owning their respective houses.

The TC algorithm with a control-rights structure satisfying R1–R3 determines a mechanism that maps profiles from \mathbf{P} to Pareto-efficient matchings in \mathcal{M} .²¹ To guarantee that the resulting mechanism is group strategy-proof, we impose the following across-round consistency requirements on the control-rights structure.

ACROSS-ROUND REQUIREMENTS. Consider submatchings σ, σ' such that $\sigma \subset \sigma' \in \overline{\mathcal{M}}$ and an agent $i \in \overline{I_{\sigma'}}$ who owns a house $h \in \overline{H_{\sigma'}}$ at σ . Then the following statements hold:

- (R4) Agent i owns h at σ' .
- (R5) If i' brokers house h' at σ and $i' \in \overline{I_{\sigma'}}, h' \in \overline{H_{\sigma'}}$, then either i' brokers h' at σ' or i owns h' at σ' .²²
- (R6) If agent $i' \in \overline{I_{\sigma'}}$ controls $h' \in \overline{H_{\sigma'}}$ at σ , then i' owns h at $\sigma \cup \{(i, h')\}$.

Requirement R4 (*persistence of ownership*) postulates that ownership rights persist: agents retain control rights as they move from smaller to larger submatchings or through the rounds of the algorithm. This is the only requirement that we imposed on control-rights structures without brokers to obtain Pápai’s (2000) hierarchical-exchange mechanisms; the other requirements R1–R3, R5, and R6 have no counterpart in Pápai’s construction.

¹⁹Previous developments of Gale’s idea—e.g., Abdulkadiroğlu and Sönmez (1999, 2003), Pápai (2000), Ehlers et al. (2002), Ehlers and Klaus (2004), and Roth et al. (2004)—required each agent to point to his or her top choice among unmatched houses and hence allowed only top trading cycles. For an extension of trading cycles to the school choice environment, see Pycia and Ünver (2011).

²⁰The order in which we clear cycles in the TC algorithm does not matter for its outcome (see the Supplement). In the version of the algorithm introduced here, we clear cycles as soon as possible.

²¹The recursive argument for the efficiency of the mechanism from Section 3.2 applies.

²²Notice that the latter case can only happen if i is the only agent in $\overline{I_{\sigma'}}$ who owns a house at σ .

Requirement R5 (*persistence of brokerage*) is a counterpart of R4 for brokers. It implies that the brokerage right persists whenever there are at least two owners.²³ R5 allows a broker to lose brokerage rights when there is exactly one owner at the base submatching σ ; at the larger submatching σ' , this unique owner must then own the formerly brokered house. If no owner at the base submatching remains unmatched at σ' , then R5 imposes no restrictions on the control rights over h' ; we refer to such situations as a complete restart.

Requirement R6 (*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 owned by i . A key implication of R4, R5, and 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 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*.

The examples in the Supplement illustrate the role of requirements R4–R6 in ensuring that trading cycles are group strategy-proof. A good way to understand these requirements is to consider the concept of obtainable objects introduced in Section 3.2: the requirements ensure that for each agent the set of obtainable objects stays the same or becomes larger with each round. The complexity of R5 and R6 notwithstanding, a priori one could expect the conditions needed to define the class of group strategy-proof and efficient mechanisms to be much more complex than they turned out to be.²⁴

We are now ready to define the TC mechanisms.

DEFINITION 3. A control-rights structure is *consistent* if it satisfies requirements R1–R6. A mechanism is *TC (trading cycles)* if its outcomes may be determined by running the TC algorithm with some consistent control-rights structure.

We denote by $\psi^{c,b}$ the TC mechanism obtained from a consistent control-rights structure $\{(c_{\sigma}, b_{\sigma})\}_{\sigma \in \overline{\mathcal{M}}}$. The TTC mechanisms of Section 3.1 and the non-TTC mechanism of Section 3.2 are examples of TC.²⁵

4. MAIN RESULTS

Every group strategy-proof and Pareto-efficient mechanism can be described in terms of owners and brokers and implemented as trading cycles. We first state and prove these

²³Indeed, the assumption postulates that if a broker i' loses the brokerage right, then the house is owned by any agent i who owns some house at σ and remains unmatched. Hence, if there are two agents owning houses at σ , then the brokerage right cannot be lost.

²⁴Sections 5 and 6 illustrate the use of brokers and restrictions R5 and R6. We think that nearly all of the flexibility of brokers is captured by the simpler class of mechanisms in which R5 and R6 are replaced with the restriction that if $|\sigma'| < |I| - 1$ and i brokers h at σ , then i brokers h at σ' . This simpler restriction results in a smaller class of mechanisms (see the Supplement).

²⁵A natural extension of Pápai's (2000) fixed-endowment mechanisms (see Section 3.1 for a discussion) is the TC mechanisms in which $(c_{\sigma}, b_{\sigma})(h) = (c_{\sigma'}, b_{\sigma'})(h)$ for all σ and σ' such that $I_{\sigma} = I_{\sigma'}$. In addition to Pápai's fixed-endowment mechanisms, this class contains, for instance, the TC mechanism from Section 3.2.

results—Theorems 1 and 2—for the model of allocation in which there are more objects than agents, and we later relax these assumptions.²⁶

THEOREM 1. *A mechanism is group strategy-proof and Pareto efficient if and only if it is a trading-cycles mechanism.*

The proof that trading-cycles mechanisms are Pareto efficient follows the same recursive steps as the argument for the efficiency of the non-TTC mechanism in Section 3.2.²⁷ By the lemma of Pápai (2000) (Lemma 1 above), a mechanism is group strategy-proof if and only if it is individually strategy-proof and nonbossy. The proof of these two properties builds on the following straightforward lemma.

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 same submatching forms before round r under $[\succ_i, \succ_{-i}]$ and $[\succ'_i, \succ_{-i}]$, and hence the control rights structure at round r is the same under $[\succ_i, \succ_{-i}]$ and $[\succ'_i, \succ_{-i}]$.*

This lemma is key in proving individual strategy-proofness. It tells us that as long as an agent is unmatched, he cannot influence when he becomes an owner, becomes a broker, or enters the broker-to-heir transition (see the discussion of R6) by choosing which preferences to submit. Owners cannot benefit from waiting because they get the best available house at the time they match under \succ . Checking that brokers cannot benefit from waiting is only slightly more subtle, as is checking that no agent can improve his match by matching early. Appendix A provides the details.

To get a sense of why the TC mechanisms are nonbossy, consider a TC mechanism without brokers and an agent i who gets the same object whether he submits preferences \succ_i or \succ'_i . A standard inductive argument then shows that the algorithm goes through 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 differ. The same argument applies in the presence of brokers unless one of them loses the brokerage right.²⁸ A broker can lose the brokerage right only in a limited set of circumstances delineated in conditions R5 and R6; these conditions ensure that the mechanism is nonbossy even if some brokers lose their brokerage rights. Appendix B provides the details.

The main part of the proof of Theorem 1 is to show that if a mechanism φ is group strategy-proof and Pareto efficient, then we can construct a TC mechanism $\psi^{c,b}$ that is equivalent to φ . The construction proceeds in three steps: we construct the candidate control-rights structure (c, b) , we check that it satisfies conditions R1–R6, and we verify by an induction on the rounds of the resulting TC mechanism $\psi^{c,b}$ that $\psi^{c,b} = \varphi$. Appendix C provides the details of all three steps.

The keystone of the proof is the construction of the candidate control-rights structure. We define it in terms of how φ allocates objects for preferences from some special

²⁶See the end of this section for more details.

²⁷For completeness, we present the proof of Pareto efficiency in the Supplement.

²⁸When a broker loses his brokerage right, cycles of one or two agents can be different under \succ and \succ' . The Supplement contains an example illustrating this possibility.

preference classes, using the approach we introduced in Section 3.2. Consider first the empty submatching and a house h . Note that if φ were a TC, h was owned, and all agents ranked h at the top, then φ would allocate h to its owner. We thus check whether φ allocates h to the same agent when all agents put h at the top, and if it does, we say that this agent *owns** h at the empty submatching.²⁹ If φ does not allocate h to the same agent at all the above profiles, we call h a *brokered** house. Note that if φ were a TC and h was brokered, 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 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 say that this agent *brokers** h at the empty submatching. The construction of candidate control rights at nonempty submatchings is similar. 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. When matched agents rank the houses they are matched with at the top while all other agents rank matched houses at the bottom, the Pareto efficiency of TC mechanisms implies that the above procedure recovers the ownership and brokerage rights of any TC mechanism. We thus use this procedure to define candidate control rights at nonempty submatchings.³⁰

To show that this construction delivers the entire candidate structure of control rights, we need to check whether each unmatched house is either *owned** or *brokered**. The surprising result establishing this property is the keystone of the entire paper (and its proof is the hardest).

THEOREM 2. *For any group strategy-proof and Pareto-efficient mechanism, for any submatching σ , and for any house unmatched at σ , there is either a unique agent who owns* this house at σ or else there is a unique agent who brokers* this house at σ .*

To get a sense of the proof, consider the empty submatching. We first show that if all agents rank some house h as their first choice and some other house as their second choice, then the agent who obtains h does not depend on the rest of the preference profile (Lemma 8). Then we show that this house is either *owned** or *brokered**: the house is *owned** if the same agent obtains it when all agents rank it first, irrespective of what house is ranked second (Lemma 10); otherwise, the house is *brokered** (Lemma 9). Appendix C.1 provides the details.

The next section extends our analysis to allocation and exchange, and the Supplement relaxes the assumption that there are more objects than agents. The latter assumption simplifies the exposition, but our results do not hinge on it. In particular, with

²⁹By Maskin monotonicity, a group strategy-proof mechanism φ allocates h to an agent whenever he ranks it at the top if and only if φ allocates h to this agent whenever all agents put h at the top. Thus, the formulation of *ownership** in the introduction and Section 3.2 is equivalent to the one above.

³⁰Recovering the control-rights structure in this way ensures that there are no brokers at a submatching unless there are at least two owners at this submatching. For flexibility we also allow control-rights structures in which a broker coexists with a single owner at some submatchings. Following Theorem 6 in the Appendix, we discuss the equivalence of these two approaches.

this assumption relaxed, the key insight that group strategy-proof and Pareto-efficient mechanisms can be described in terms of owners* and brokers* (Theorem 2) remains true with no change in its statement or its proof. The resulting owner-and-broker mechanisms can then again be described through a recursive trading-cycles algorithm.³¹

5. INDIVIDUALLY RATIONAL HOUSE ALLOCATION AND EXCHANGE

Our results extend to the setting in which some agents have private endowments.³²

5.1 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 $\bigcup_{i \in \{0\} \cup I} H_i = H$. 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 . Because we allow some of the agents to have an empty endowment, the allocation model of Section 2 is contained as a special case with $\mathcal{H} = \{H, \emptyset, \dots, \emptyset\}$. Matchings and mechanisms are defined as in the allocation model of Section 2.

The concepts of Pareto efficiency and group strategy-proofness are the same as before. In particular, the equivalence between group strategy-proofness and the conjunction of individual strategy-proofness and nonbossiness continues 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 his endowment: $\mu(i) \succeq_i h$ for all $i \in I$ and $h \in H_i$. For agents with empty endowments, $H_i = \emptyset$, this condition is tautologically true.

5.2 Results

The analogue of Theorem 2 remains true with no change in the statement or the proof, and the following characterization is an immediate corollary of Theorem 1.³³

³¹In pre-2015 drafts, we erroneously claimed that no mechanism with multiple simultaneous brokers* is group strategy-proof whether or not $|H| > |I|$: the last paragraph of the proof of Lemma 11 was incorrect in the case of three houses and three agents. This error had no impact on the results in the main part of our paper; its impact on our analysis was limited to the last step of our recursive construction and proofs in the environment of the current Supplement. Bade (2014) recognized this error, implemented the special case of our owner*-and-broker* construction with exactly three houses that all are brokered* via an algorithm that avoids certain matchings, and verified its efficiency and group strategy-proofness.

³²Our results allow one to easily obtain further new insights into allocation and exchange without transfers. For instance, they imply that a mechanism is individually strategy-proof, Pareto efficient, and invariant if and only if it is a hierarchical-exchange mechanism of Pápai (2000). This and other results from earlier drafts will form the core of a separate paper.

³³The assumptions of Theorem 3 are satisfied, for instance, in the *house allocation with existing tenants* of Abdulkadiroğlu and Sönmez (1999). Their subject matter is the subclass of house allocation and exchange in which each agent is endowed with one or zero houses—a situation common, for instance, in the allocation of college dormitory rooms.

THEOREM 3. *A mechanism is individually rational, group strategy-proof, and Pareto efficient if and only if this mechanism is an individually rational TC mechanism.*

It is straightforward to identify individually rational TC mechanisms.

THEOREM 4. *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 over all houses from his endowment.*³⁴

PROOF. To prove individual rationality of the above subclass of TC mechanisms, consider an agent i and assume that, at the empty submatching, i owns some house h from his endowment. Then R4 ensures that i owns h throughout the execution of the TC algorithm. Thus, the TC mechanism allocates 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 1](#). For any agent i and house h from i 's endowment, i is the owner* of h because individual rationality implies that $\psi[\succ](i) = h$ for any $\succ \in \mathbf{P}[\emptyset, h]$, which is the set of preference profiles that rank h first for all agents. The construction from the proof of [Theorem 1](#), therefore, yields a control-rights structure that assigns to each agent the initial ownership rights over the houses from his endowment, and represents ψ . \square

As a corollary, we obtain the following characterization of top trading cycles in an important subdomain of allocation and exchange.³⁵

THEOREM 5. *Suppose each agent has a nonempty endowment. A mechanism is individually rational, group strategy-proof, and Pareto efficient if and only if this mechanism is a TTC mechanism (aka hierarchical exchange) that assigns all agents the initial ownership rights over houses from their endowment.*

PROOF. 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 [Theorem 4](#) we may assume that each agent has initial ownership rights over the houses from their endowment. By condition R4 (regarding consistency), all unmatched agents own a house throughout the mechanism and, hence, R3 implies that no agent is a broker. Thus $\psi^{c,b}$ is a TTC mechanism. \square

³⁴The initial control rights are the control rights at the empty submatching. 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, which is brokered by a broker. 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.

³⁵The assumptions of [Theorem 5](#) are satisfied by the *kidney exchange* with strict preferences ([Roth et al. 2004](#)), and the *kidney exchange with good Samaritan donors* ([Sönmez and Ünver 2006](#)). In this interpretation, 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.

6. FURTHER ILLUSTRATIVE COROLLARIES

Some of the canonical prior insights of the allocation literature easily follow from our main results.³⁶

6.1 Neutrality

Svensson (1999) showed that all neutral and group strategy-proof mechanisms are serial dictatorships. A mechanism is *neutral* if, whenever the house names are relabeled, the mechanism assigns each agent the house that carries the relabeled name of the house that was assigned before the relabeling.³⁷ In a serial dictatorship, agents are ordered, and the first agent in the ordering gets his most preferred house, the second agent in the ordering gets her most preferred house among those unassigned to agents higher in the ordering, etc.

Svensson’s result is a corollary of Theorem 1, as illustrated below.

COROLLARY 1 (Svensson 1999). *A mechanism is group strategy-proof and neutral if and only if it is a serial dictatorship.*

PROOF. Let φ be a group strategy-proof and neutral mechanism. Neutrality implies that φ has full range, that is, $\varphi[\mathbf{P}] = \mathcal{M}$. 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 $\varphi[\succ^\pi] = \mu$. Full range and group strategy-proofness imply Pareto efficiency. Thus, φ is a trading-cycles mechanism $\psi^{c,b}$ by Theorem 1. Consider this TC mechanism. For any $\sigma \in \overline{\mathcal{M}}$, conditions R1 and R3 imply that there is an agent $i \in \overline{I}_\sigma$ 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.³⁸ \square

6.2 Reallocation-proofness

Pápai (2000) showed that group strategy-proof and Pareto-efficient mechanisms that satisfy an additional condition can be implemented as hierarchical exchange. The condition she relies on is that a mechanism φ is *reallocation-proof* if there exists no pair

³⁶Our results may also be used, for instance, to facilitate the analysis of consistent, individually strategy-proof, and Pareto-efficient mechanisms (Sönmez and Ünver 2006, Ehlers and Klaus 2007, and Velez 2014). Indeed, individually strategy-proof and consistent mechanisms are group strategy-proof. See also Ergin (2000).

³⁷More formally, a relabeling of houses is a bijection $\pi : H \rightarrow H$. 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 $\varphi[\succ^\pi](i) = \pi(\varphi[\succ](i))$.

³⁸We prove here Svensson’s result for the case when there are more houses than agents, but neither Svensson’s result nor our argument hinges on this assumption; see the Supplement for details. Interestingly, in the presence of outside options, Svensson’s result is no longer true, as shown by Pycia and Ünver (2014).

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 $\varphi[\succ'_{\{i,j\}}, \succ_{-\{i,j\}}](i) \succ_j \varphi[\succ](j)$. We can derive the key insight of Pápai (2000) as follows.

COROLLARY 2 (Pápai 2000). *If a mechanism is group strategy-proof, Pareto efficient, and reallocation-proof, then it is a hierarchical-exchange mechanism.*

PROOF. Let φ be a group strategy-proof, Pareto -efficient, and reallocation-proof mechanism. By Theorem 1, φ is equivalent to a reallocation-proof TC mechanism $\psi^{c,b}$. The unique control-rights structure we construct has the property that there are no brokers at submatchings with exactly one owner (see Theorem 6). To prove that there are no brokers at other submatchings, let $\sigma \in \overline{\mathcal{M}}$, let $j, k \in \overline{I}_\sigma$ be two different owners, and let h_j be a house owned by j and h_k be a house owned by k at σ . By way of contradiction, assume that some house h is brokered by some agent i at σ . Consider a preference profile $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \mathbf{P}_i[\sigma; h, h_k, h_j]$, $\succ_j \in \mathbf{P}_j[\sigma; h_j]$, and $\succ_k \in \mathbf{P}_k[\sigma; h]$ (these sets of preference rankings are defined in the proof of Theorem 1). 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.³⁹ \square

7. CONCLUSION

We study allocation and exchange in environments without transfers and with single-unit demands. Addressing the central concern of both market-design practice and its theory, we construct the full class of group strategy-proof and Pareto-efficient mechanisms.⁴⁰ Our construction relies on the introduction of brokers to the analysis of allocation and exchange. Suppose some (possibly none) agents and objects are already matched, and consider the unmatched agents and objects: a broker of an object is the agent who obtains his second choice whenever all agents rank the brokered object as their top choice. Our key insight is that every unmatched house either has a uniquely determined broker or a uniquely determined owner (that is, an agent who obtains this house whenever he prefers it over other outcomes).

Brokers can be useful beyond allowing us to construct the full class of strategy-proof and efficient mechanisms. In asymmetric settings, we may want to use brokered trading cycles instead of top trading cycles because of equity or other considerations. For instance, consider a manager who assigns n tasks h_1, \dots, h_n to n employees i_1, \dots, i_n with strict preferences over the tasks. 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 h_n to employee i_n (for instance, this employee may be known to be less proficient than the other employees at this task). That is, if there is a Pareto-efficient

³⁹As in the case of Svensson's result, here we prove Pápai's result for the case when there are more houses than agents, but neither Pápai's result nor our argument hinges on this assumption; see the Supplement for details.

⁴⁰One of our corollaries characterizes TTC; for later such characterizations, see Abdulkadiroğlu and Che (2010), Morrill (2013), Dur (2013), Dur and Ünver (2015), Tang and Zhang (2016), and Pycia (2016).

matching that avoids assigning h_n to i_n , she would like to choose it. Because she does not know employees' preferences, she wants to use a group strategy-proof mechanism.

The use of brokerage rights allows the manager to improve the outcome of agent i_n over any top-trading-cycles mechanism. Indeed, to use one of the top-trading-cycles mechanisms of Pápai (2000), the manager would need to initially endow employees i_1, \dots, i_{n-1} with all the tasks, and the manager can allow i_n to inherit some task only after either (i) all other employees are already matched, or (ii) task h_n is matched. The trading-cycles mechanism that modifies such a top-trading-cycles mechanism by endowing employee i_n with initial brokerage rights over h_n improves i_n 's allocation and satisfies the constraints of the manager's problem.

The use of brokerage rights also allows the manager to treat the employees more equitably than is possible under top trading cycles. In an earlier draft, we showed that a trading-cycles mechanism in which i_n is the permanent broker of h_n and every other agent is an owner of one of the remaining tasks is more equitable—in the sense of strict Lorenz dominance (Lorenz 1905)—than any top-trading-cycles mechanism.⁴¹

Trading cycles with brokers are useful in maximizing many other social welfare criteria. In fact, as discussed in Pycia and Ünver (2016), the class of trading-cycles mechanisms is exactly the class of mechanisms that are individually strategy-proof and always select the efficient outcome with respect to a social ranking satisfying Arrovian postulates.

APPENDIX A: PROOF OF INDIVIDUAL STRATEGY-PROOFNESS OF TC

The first part of the following lemma is used in the proof of individual strategy-proofness, and the second part is used in the proof of non-bossiness of TC.

LEMMA 4. *Let $\psi^{c,b}$ be a TC mechanism, $\succ \in \mathbf{P}$, and $n \geq 2$. If a chain $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n$ forms in some round of $\psi^{c,b}$, then $\psi^{c,b}(i^n) \succsim_{i^n} h^1$. If, additionally, $\psi^{c,b}(i^n) = h^1$, then $\psi^{c,b}(i^\ell) = h^{\ell+1}$ for $\ell = 1, \dots, n - 1$ and either the cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$ forms in some round of $\psi^{c,b}$ or else $n = 2$ and the cycles $h^2 \rightarrow i^1 \rightarrow h^2$ and $h^1 \rightarrow i^2 \rightarrow h^1$ form in two subsequent rounds of $\psi^{c,b}$.*

PROOF. First suppose that in round r for each $\ell \in \{1, \dots, n\}$ agent i^ℓ owns house h^ℓ . R4 implies that the chain $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n$ continues forming in subsequent rounds as long as i^n is unmatched. Thus, at the round agent i leaves, he is matched with a house at least as good as h^1 , and the first part of the lemma obtains. In addition, if i^n is matched with h^1 , then he leaves in the cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$, and the remainder of the lemma obtains.

Second, suppose there is $\ell \in \{1, \dots, n\}$ such that i^ℓ brokers h^ℓ in round r . Because all agents in the chain other than i^ℓ have ownership rights, R4 implies that the chain $h' = h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n$ forms in subsequent rounds as long as i^ℓ continues

⁴¹For further development of our Lorenzian approach to equity in deterministic allocation without transfers, see Harless and Manjunath (2016) and Harless (2015); for other approaches, see Hakimov and Kesten (2014) and Morrill (2015a).

brokering h^ℓ and i^{n+1} is unmatched. If i^ℓ continues to broker h^ℓ until the round i^n is matched, then in this round the cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$ is matched, and the lemma obtains. Suppose, thus, that i^ℓ brokers h^ℓ between round r and some round $s - 1 \geq r$ but not at round s , and that i^n is still unmatched at s . Other agents in the cycle of i^ℓ are owners at both rounds $s'' - 1$ and s'' . By R5, each of these owners becomes the owner of h^ℓ in round s'' , implying that there is only one such owner and that $n = 2$. By R1 and R5, in round s'' the unique previous owner owns both houses: his previously owned house and the ex-brokered house. There are two subcases: in round s either i^1 owns both $h^1 = h'$ and h^2 or i^2 does. In the former subcase, i^1 leaves with h^2 in round s in a single-agent cycle, and R6 implies that i^2 owns h^1 in round $s + 1$ (broker-to-heir transition) and, hence, R1 yields the lemma. In the latter subcase, i^2 owns h^1 in round s'' and, by R1, leaves with a house at least as good as h^1 , and the first part of the lemma obtains. If, now, i^2 leaves with h^1 , then he leaves in a single-agent cycle; furthermore R6 then implies that i^1 owns h^2 after i^1 left. Since h^1 is already matched, we conclude that i^1 is matched with h^2 in a single-agent cycle, and the remainder of the lemma obtains. \square

PROOF OF THE INDIVIDUAL STRATEGY-PROOFNESS OF TC. Let $\psi^{c,b}$ be a TC mechanism and let \succ be a preference profile. 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 .

Fix agent $i \in I$. We need to show that i cannot benefit from submitting $\succ'_i \neq \succ_i$ while the other agents submit \succ_{-i} . Let s be the round i leaves submitting \succ_i and let h be the house he is matched to. Let s' be the round i leaves submitting \succ'_i and let h' be the house he is matched to. By Lemma 3, at round $\min(s, s')$ the same houses and agents are in the market and the control rights are the same at both \succ_i and \succ'_i . We consider two cases.

Case $s > s'$. In round $s' = \min(s, s')$ under \succ'_i , agent i points to house $h' = h^1$ that points to some agent i^1 that points to \dots that points to some house h^n that points to agent $i = i^n$, and this cycle leaves. If $n = 1$, then i is the owner of h' in round s' and R1 yields the individual strategy-proofness. Suppose $n \geq 2$. The chain $h' = h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n = i$ forms in round s' under \succ , and Lemma 4 implies that when agent i reports the true ranking, he is matched with a house he weakly prefers over h' . Thus, the individual strategy-proofness obtains.

Case $s \leq s'$. If i is not a broker at time $s = \min(s, s')$ under \succ_i or else he prefers some other house unmatched at $\sigma^{s-1}[\succ]$ over the brokered house, then submitting \succ_i gives agent i his top choice house among houses remaining in round s , and he cannot be better off submitting \succ'_i . Thus, assume that i is a broker in round s under \succ_i and that the brokered house, denoted e , is his top choice remaining house. Under \succ_i , he leaves with his second-best remaining house. Thus, there is an agent j who is matched with e in round s . By R1, in this round, j is an owner of some house h_j and e is his top choice house. By Lemma 3, if i submits \succ'_i , then he is still the broker of e and j is still the owner of h_j in round s . By way of contradiction, suppose that i is matched with e under \succ'_i . Because j 's top choice is still e and because j remains an owner of h_j until matched (by R1), this agent stays unmatched as long as e is unmatched. Agent i is not the broker of e in round s' because he leaves with e in this round and a broker cannot leave with the brokered house while owner j wants it. There is thus a round $s'' \in \{s + 1, \dots, s'\}$ at which

agent i stops being the broker of e . Because j is an owner of h_j at both $s'' - 1$ and s'' , he would have inherited e at s'' by R5. But e is j 's top choice in round s and hence s'' , and hence j would have left with e at s'' ; a contradiction. \square

APPENDIX B: PROOF OF NON-BOSSINESS OF TC

As before, $\psi^{c,b}$ is 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 $\psi^{c,b}[\succ'](i_*) = \psi^{c,b}[\succ](i_*)$. Let h_* be the house i_* obtains, and let s be the round he leaves submitting \succ_{i_*} and let s' be the time he leaves submitting \succ'_{i_*} . We refer to cycles formed under \succ as \succ -cycles, and to cycles formed under \succ' as \succ' -cycles, and we use the $\sigma^{s-1}[\succ]$ notation introduced in Appendix A.⁴² By symmetry, it is enough to consider the case $s \leq s'$. So as to show that $\psi^{c,b}[\succ](i) = \psi^{c,b}[\succ'](i)$ for all $i \in I$, we prove the following stronger statement.

Inductive hypothesis: If a cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$ of length $n \in \{1, 2, \dots\}$ forms and is removed at round r under preference profile \succ , then under preference profile \succ' at least one of the following three possibilities obtains:

- (i) The same cycle $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$ forms.
- (ii) We have $n = 2$ and two cycles form:
 - We have $h^1 \rightarrow i^2 \rightarrow h^1$, or else $g \rightarrow i^2 \rightarrow h^1 \rightarrow i \rightarrow g$ for some agent $i \neq i^2$ and some house $g \neq h^1$, and
 - We have $h^2 \rightarrow i^1 \rightarrow h^2$, or else $h \rightarrow i^1 \rightarrow h^2 \rightarrow j \rightarrow h$ for some agent $j \neq i^1$ and some house $h \neq h^2$.
- (iii) We have $n = 1$ and the cycle $h \rightarrow i^1 \rightarrow h^1 \rightarrow j \rightarrow h$ forms for some agent $j \neq i^1$ and house $h \neq h^1$.

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

Initial step. Consider $r = s$. Under \succ , some house h_*^1 points to agent $i_* = i_*^1$ who points to house $h_* = h_*^2$ that points to some agent i_*^2 that points to \dots that points to some agent i_*^n that points to house h_*^1 , and the cycle $h_*^1 \rightarrow i_*^1 \rightarrow h_*^2 \rightarrow \dots \rightarrow h_*^n \rightarrow i_*^n \rightarrow h_*^1$ is removed at round s . By Lemma 3 all agents from $I_{\sigma^s[\succ]} - \{i_*^1, \dots, i_*^n\}$ are matched by $\sigma^s[\succ']$ in the same way as in $\sigma^s[\succ]$. If $n = 1$, then i_* is the owner of h_* under \succ and, by Lemma 3, under \succ' ; under both profiles he leaves with it in a single-agent cycle. If $n \geq 2$, then Lemma 3 implies that the chain $h_* = h_*^2 \rightarrow \dots \rightarrow h_*^n \rightarrow i_*^n \rightarrow h_*^1 \rightarrow i_*^1 = i_*$ forms at round s under \succ' , and Lemma 4 yields the inductive hypothesis.

Inductive step. Take any round $r > s$ such that $\sigma^r[\succ] - \sigma^{r-1}[\succ]$ is nonempty, and assume that the inductive hypothesis is true for all rounds up to $r - 1$. Take any cycle that forms in round r under \succ , say $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots \rightarrow h^n \rightarrow i^n \rightarrow h^1$, where $n \geq 1$,

⁴²Recall that in cycles of length n the superscripts on houses and agents are modulo n , that is, $i^{n+1} = i^1$ and $h^{n+1} = h^1$.

$i^\ell \in I$, and $h^\ell \in H$ for $\ell = 1, \dots, n$. Since i_* is matched before round r , agents i^1, \dots, i^n have the same preferences in both profiles \succ and \succ' .

CLAIM 1. (i) 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 . (ii) 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 $\sigma \cup \sigma' \cup \{(j', h)\}$.

PROOF. Part (i) follows from R4. To prove (ii), notice that R4 implies that j brokers h at σ' and, hence, $j \neq j'$. R4 also 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 $\sigma \cup \sigma'$, R5 implies that j' owns h at $\sigma \cup \sigma'$, and R6 implies that j owns h' at $\sigma \cup \sigma' \cup \{(j', h)\}$. \square

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. Consider the run of the algorithm under \succ' . If i^ℓ is a $\sigma^{r-1}[\succ]$ -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 a $\sigma^{r-1}[\succ]$ -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. \square

Let t be the earliest round in which one of the houses h^1, \dots, h^n is matched under \succ' ; let h^1 be a house matched in round t under \succ' , and let $\nu = \sigma^{t-1}[\succ] \cup \sigma^{t-1}[\succ']$. By Claim 2, all agents i^1, \dots, i^n are unmatched at $\sigma^{t-1}[\succ']$ because all houses h^1, \dots, h^n are. Thus, all these agents and houses are unmatched at ν .

CLAIM 3. Agent i^1 cannot own h^1 at one of the submatchings $\sigma^{r-1}[\succ]$ or $\sigma^{t-1}[\succ']$ and broker this house at the other.

PROOF. Suppose i^1 owns h^1 at $\sigma^{r-1}[\succ]$ and brokers it at $\sigma^1[\succ']$. Then there exist some agent j and house h such that $h \rightarrow j \rightarrow h^1 \rightarrow i^1$ is part of the cycle occurring in round t under \succ' . Because h^1 is unmatched under $\sigma^{r-1}[\succ]$, so is j by the inductive assumption. Thus, the inductive assumption also implies that h is unmatched under $\sigma^{r-1}[\succ]$ and, by R1, j owns h at all submatchings between $\sigma^{t-1}[\succ']$ and ν . Since i^1 loses brokerage of h^1 between $\sigma^{t-1}[\succ']$ and ν , R5 and R4 imply that j owns h^1 at ν ; a contradiction. The proof that i^1 cannot own h^1 at $\sigma^{t-1}[\succ']$ and broker it at $\sigma^{r-1}[\succ]$ is analogous.⁴³ \square

Let j^1 be the agent controlling house h^1 at $\sigma^{t-1}[\succ']$. We may assume that $j^1 \neq i^1$. Indeed, suppose $j^1 = i^1$, that is, i^1 controls h^1 at $\sigma^{t-1}[\succ']$, and consider three cases:

⁴³Except that we do not need to work to show that one of the owners in the cycle of h^1 under \succ is unmatched at $\sigma^{t-1}[\succ']$.

- If $n = 1$, then i^1 owns h^1 at $\sigma^{r-1}[\succ]$ and, by R4, at ν as well. By Claim 3, i^1 owns h^1 at $\sigma^{t-1}[\succ']$ and, by Claim 2, he points to h^1 and is matched with it. Part (i) of the inductive hypothesis is thus satisfied.
- If $n \geq 2$ and i^1 prefers h^2 over h^1 , then, by Claim 2, he points to it at $\sigma^{t-1}[\succ']$ (he cannot broker it since he controls h^1) and house h^1 is matched in a cycle that contains $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots$.
- If $n \geq 2$ and i^1 prefers h^1 over h^2 , then i^1 is the broker of h^1 at $\sigma^{r-1}[\succ]$. By Claim 3, i^1 brokers h^1 at $\sigma^{t-1}[\succ']$. By Claim 2, i^1 points to h^2 (as he cannot point to h^1) and house h^1 is matched in a cycle that contains $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots$.

Thus if $j^1 = i^1$, then either part (i) of the inductive hypothesis is satisfied or the cycle of h^1 at \succ' contains $h^1 \rightarrow i^1 \rightarrow h^2 \rightarrow \dots$. In the latter case, we define j^2 to be the agent controlling h^2 at $\sigma^{t-1}[\succ']$ and repeat the above procedure for h^2 . In this way, either we prove that part (i) of the inductive hypothesis is satisfied or we find k such that $i^k \neq j^k$. In the latter case, given the symmetry, we can shift the names so that $k = 1$ and, hence, $i^1 \neq j^1$ as required. To conclude the proof, consider two cases.

Case i^1 owns h^1 at $\sigma^{r-1}[\succ]$. By R4, he then owns h^1 at ν . Suppose all agents in the \succ' -cycle of h^1 are unmatched at $\sigma^{r-1}[\succ]$; the inductive assumption implies that all houses in this cycle are also unmatched at $\sigma^{r-1}[\succ]$. 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, there is another agent $j^0 \neq j^1$ and a house h^0 such that $h^0 \rightarrow j^0 \rightarrow h^1 \rightarrow j^1$ is part of the cycle of h^1 under \succ' . By R4 and R5, there can be at most one other agent in this cycle because j^1 loses brokerage rights over h^1 while all agents and houses in this cycle remain unmatched. Thus, the cycle is $h^0 \rightarrow j^0 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$. By R5 and R4 again, j^0 owns h^1 at ν ; hence $j^0 = i^1$. Thus 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.

Thus, at least one agent in the \succ' -cycle of h^1 is matched at $\sigma^{r-1}[\succ]$. By the inductive assumption, $n \in \{1, 2\}$ and the \succ' -cycle of h^1 contains one or two agents. The cycle of h^1 under \succ' cannot be $h^1 \rightarrow j^1 \rightarrow h^1$ as then the inductive assumption would imply that j^1 and h^1 were matched at $\sigma^{r-1}[\succ]$. Thus, there is an agent $j^0 \neq j^1$ and house $h^0 \neq h^1$ such that $h^0 \rightarrow j^0 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$ is the cycle of h^1 under \succ' . By the inductive assumption, agent j^0 cannot be matched at $\sigma^{r-1}[\succ]$; hence agent j^1 and house h^0 are. By R6, j^0 owns h^1 at $\sigma^{t-1}[\succ'] \cup \{(j^1, h^0)\} \subset \nu$. Since j^0 is unmatched at ν , R4 implies he owns h^1 at ν . As i^1 also owns h^1 at ν , we conclude $j^0 = i^1$. Thus, $h^0 \rightarrow i^1 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$ is the cycle of h^1 under \succ' . If $n = 2$, then being a $\sigma^{r-1}[\succ]$ -owner, i^1 prefers h^2 over h^1 . As h^2 is unmatched at $\sigma^{t-1}[\succ']$, we conclude that i^1 brokers h^2 at this submatching. Thus, $h^0 = h^2$, a contradiction as h^0 is matched and h^2 is unmatched at $\sigma^{r-1}[\succ]$. We can conclude that $n = 1$ and part (iii) of the inductive hypothesis is satisfied.

Case i^1 brokers h^1 at $\sigma^{r-1}[\succ]$. In particular, $n \geq 2$, and agent i^2 is the $\sigma^{r-1}[\succ]$ -owner of h^2 .

First, consider how h^1 is matched under \succ' . Let $h^0 \rightarrow j^0 \rightarrow h^1 \rightarrow j^1$ be part of the cycle of h^1 in round t under \succ' ; this is without loss of generality as we allow $h^0 = h^1$ and $j^0 = j^1$. By the inductive assumption, j^0 is unmatched at $\sigma^{r-1}[\succ]$. The following subcases are possible.

- All houses and agents in the cycle of h^1 under \succ' are unmatched at $\sigma^{r-1}[\succ]$.
 First, suppose j^1 brokers h^1 at $\sigma^{t-1}[\succ']$. Then $j^0 \neq j^1$, $h^0 \neq h^1$, and j^0 owns house h^0 at $\sigma^{t-1}[\succ']$. 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 rights, by R5 there are only two agents in the cycle of h^1 under \succ' or \succ , respectively, and 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., $h^0 \rightarrow i^2 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$.
 Second, suppose j^1 owns h^1 at $\sigma^{t-1}[\succ']$. Then R4 implies that j^1 owns h^1 at ν . Thus, i^1 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. By R4 and R5, h^1 is owned by i^2 at ν . Because h^1 is also owned by j^1 at ν , we have $j^1 = i^2$. By Claim 2, h^1 is the best house that i^2 can get under \succ' . Therefore, the cycle of h^1 in round t under \succ' is $h^1 \rightarrow i^2 \rightarrow h^1$.
- At least one agent or house in the cycle of h^1 under \succ' is matched at $\sigma^{r-1}[\succ]$. By the inductive assumption, $n = 2$, j^1 and h^0 are matched at $\sigma^{r-1}[\succ]$ while j^0 and h^1 are not, and the cycle of h^1 under \succ' is $h^0 \rightarrow j^0 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$. R6 implies that j^0 owns h^1 at $\sigma^{t-1}[\succ'] \cup \{(j^1, h^0)\}$ and, hence, by R4, at ν . Then i^1 leaves brokerage of h^1 between $\sigma^{r-1}[\succ]$ and ν . As i^2 owns h^2 at $\sigma^{r-1}[\succ]$, R5 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 $h^0 \rightarrow i^2 \rightarrow h^1 \rightarrow j^1 \rightarrow h^0$.

In all subcases, $n = 2$ and the second part of the inductive hypothesis is satisfied for i^1 and h^2 .

Second, consider how h^2 is matched under \succ' . Because i^1 brokers h^1 and i^2 owns h^2 at $\sigma^{r-1}[\succ]$, R6 implies that i^1 owns h^2 at $\sigma^{r-1}[\succ] \cup \{(i^2, h^1)\}$ and, hence, by R4, at $\nu \cup \{(i^2, h^1)\}$. Let $t^2 \geq t$ be the round in which h^2 is matched under \succ' . Because h^1 is matched with i^2 under \succ' , Claim 2 implies that i^1 is unmatched at t^2 . Suppose $j' \rightarrow h^2 \rightarrow j^2$ is part of the cycle of h^2 under \succ' . Let $\nu^2 = \sigma^{r-1}[\succ] \cup \sigma^{t^2-1}[\succ'] \cup \{(i^2, h^1)\}$. Because h^2 is unmatched at $\sigma^{r-1}[\succ]$ and j' matches with h^2 under \succ' , the inductive assumption implies that j' is unmatched at $\sigma^{r-1}[\succ]$ and, hence, at ν^2 . Since $\nu^2 \supseteq \nu$, R4 implies that i^1 owns h^2 at ν^2 . The following subcases are possible.

- $i^1 = j^2$. If i^1 brokers h^2 at $\sigma^{t^2-1}[\succ']$, then $i^1 \neq j'$ and j' is an owner at $\sigma^{t^2-1}[\succ']$. Since i^1 loses brokerage rights over h^2 between $\sigma^{t^2-1}[\succ']$ and ν^2 (as he owns it at ν^2), R5 and R4 imply that j' owns h^2 at ν^2 , contradicting that i^1 does. We conclude that i^1 owns h^2 at $\sigma^{t^2-1}[\succ']$. Because h^1 is matched with i^2 and not i^1 under \succ' , Claim 2 implies that h^2 is matched in the cycle $h^2 \rightarrow i^1 \rightarrow h^2$ under \succ' .
- $i^1 \neq j^2$ and all houses and agents in the cycle of h^2 under \succ' are unmatched at $\sigma^{r-1}[\succ]$. 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 $j' \neq j^2$. Broker j^2 loses his brokerage right between $\sigma^{t^2-1}[\succ]$ and ν^2 , and, by R5, j' is the only owner in the cycle of h^2 under \succ' . By R5 and R4, j' owns h^2 at ν^2 and, hence, $j' = i^1$. Thus, the cycle of h^2 under \succ' is $h^2 \rightarrow i^1 \rightarrow h^2 \rightarrow j^2 \rightarrow h^2$.

- $i^1 \neq j^2$ and at least one agent or house in the cycle of h^2 under \succ' is matched at $\sigma^{r-1}[\succ]$. By the inductive hypothesis, the cycle of h^2 under \succ' is $h' \rightarrow j' \rightarrow h^2 \rightarrow j^2 \rightarrow h'$, j' is unmatched (because so is h^2), and j^2 is matched with h' at $\sigma^{r-1}[\succ]$. By R6, j' owns h^2 at $\sigma^{t^2-1}[\succ'] \cup \{(j^2, h')\} \subset \nu^2$ and, by R4, he owns it at ν^2 . Because i^1 owns h^2 at ν^2 , we conclude that $i^1 = j'$, and the cycle of h^2 under \succ' is $h' \rightarrow i^1 \rightarrow h^2 \rightarrow j^2 \rightarrow h'$.

In all subcases, the second part of the inductive hypothesis is satisfied for i^1 and h^2 .

APPENDIX C: PROOF THAT GROUP STRATEGY-PROOF AND PARETO-EFFICIENT MECHANISMS ARE TC

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

We use the following notation in our construction. Let $i \in I$, $\sigma \in \overline{\mathcal{M}}$, $n \geq 0$, and $h^1, h^2, \dots, h^n \in \overline{H}_\sigma$. $\mathbf{P}_i[\sigma, h^1, \dots, h^n]$ is the set of preferences \succ_i of agent i such that the following statements hold:

- If $i \in I_\sigma$, then $\sigma(i) \succ_i g$ for all $g \in H - \{\sigma(i)\}$.
- If $i \in \overline{I}_\sigma$, then $h^1 \succeq_i h^2 \succeq \dots \succeq_i h^n \succ_i g \succ_i g'$ for all $g \in \overline{H}_\sigma - \{h^1, \dots, h^n\}$ and $g' \in H_\sigma$.

If i is not matched at σ , then $\mathbf{P}_i[\sigma, h^1, \dots, h^n]$ is the set of preferences that rank h^1, \dots, h^n in this order and above other houses unmatched at σ , and that rank unmatched houses above houses matched at σ . If i matched at σ , then $\mathbf{P}_i[\sigma, h^1, \dots, h^n]$ is the set of preferences that rank agent i 's σ -match above all other houses. A generic preference profile in $\mathbf{P}[\sigma, h^1, \dots, h^n]$ is denoted $\succ^{\sigma; h^1, \dots, h^n}$; when σ is fixed, we may suppress it in this notation and write \succ^{h^1, \dots, h^n} . Also, when σ is fixed, we may write $\langle h^1, \dots, h^n, \dots \rangle$ instead of $\mathbf{P}_i[\sigma, h^1, \dots, h^n]$. The profile domain $\mathbf{P}[\sigma, h^1, \dots, h^n] \subseteq \mathbf{P}$ is the Cartesian product of $\mathbf{P}_i[\sigma, h^1, \dots, h^n]$ over all $i \in I$, and $\mathbf{P}^*[\sigma, h] = \bigcup_{h' \in \overline{H}_\sigma - \{h\}} \mathbf{P}[\sigma, h, h']$ is the subset of $\mathbf{P}[\sigma, h]$ in which unmatched agents rank the same house as their second choice.

Using this notation, we state formally the following key concepts introduced in the main body of the paper. A house $h \in \overline{H}_\sigma$ is an owned* house at $\sigma \in \overline{\mathcal{M}}$ if $\varphi[\succ]^{-1}(h) = i$ for fixed $i \in \overline{I}_\sigma$ and all $\succ \in \mathbf{P}[\sigma, h]$; 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 house $\varphi[\succ](k)$ is the second choice of k in \succ_k for all $\succ \in \mathbf{P}^*[\sigma, e]$. Observe that a house cannot be both owned* and brokered* at the same submatching.⁴⁴

Pápai (2000) studied the following natural concept of envy: agent j envies i at \succ if $\varphi[\succ](i) \succ_j \varphi[\succ](j)$. She proved the following lemma.

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

⁴⁴It may appear from the definitions that there is a third option for an unmatched house besides being owned* and brokered* at a submatching. However, Theorem 2, proven below, shows that these are the only two options.

Pápai also proved that for all $i, j \in I$, $\succ \in \mathbf{P}$, and $\succ_j^* \in \mathbf{P}_j$, 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 $\varphi[\succ_i^*, \succ_j^*, \succ_{-\{i,j\}}](i) = \varphi[\succ](j)$. We use this insight to obtain the following lemma.

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

PROOF. If $\varphi[\succ_j^*, \succ_{-j}](i) = \varphi[\succ](i)$, then Pareto efficiency of $\varphi(\succ)$ implies that i does not envy j , and the lemma follows. If $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$, then the lemma follows from the insight of Pápai stated above and the strategy-proofness of φ . \square

C.1 Proof of Theorem 2: The starred control-rights structure is well defined

Theorem 2 tells us that any house is either owned* or brokered* and, hence, the starred control rights are defined for each house. We prove this result through several lemmas.

LEMMA 7. *If $\sigma \in \overline{\mathcal{M}}$, $i \in I_\sigma$, $h \in \overline{H_\sigma}$, and $\succ \in \mathbf{P}[\sigma, h]$, then, $\varphi[\succ](i) = \sigma(i)$.*

PROOF. 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 $\overline{I_\sigma}$ who 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 $\sigma(j)$. Because 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. \square

LEMMA 8. *If $\sigma \in \overline{\mathcal{M}}$, $e, h \in \overline{H_\sigma}$, and $h \neq e$, 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. By way of contradiction, suppose that $\succ, \succ' \in \mathbf{P}[\sigma, e, h]$ are such that $\varphi[\succ](i) = e \neq \varphi[\succ'](i)$ for some agent $i \in \overline{I_\sigma}$. Without loss of generality, we assume that \succ and \succ' differ only in the 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 $\succ_i \in \langle e, h, g, \dots \rangle \subseteq \mathbf{P}_i[\sigma, e, h]$, because 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 nonbossiness, $g' \neq g$, and by strategy-proofness, $g' \neq e, h$. Maskin monotonicity for j allows us also to assume that $\succ_j \in \langle e, h, g, g', \dots \rangle$ and $\succ'_j \in \langle e, h, g', g, \dots \rangle$.

Let $i' \in \overline{I_\sigma}$ be the agent who gets e at \succ' and let $k \in \overline{I_\sigma}$ 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 5 implies that neither i nor i' gets h at \succ . Thus $k \neq i, i', j$.

CLAIM 1. *(i) Under \succ , $\varphi[\succ](i) = e$, $\varphi[\succ](j) = g$, and $\varphi[\succ](k) = h$.*

(ii) Under \succ' , $\varphi[\succ'](i') = e$, $\varphi[\succ'](j) = g'$, $\varphi[\succ'](i) = g$, and $\varphi[\succ'](k) = h$.

PROOF. The first five equalities were proved or assumed above and are listed for convenience only. The last two equalities require an argument. First, we show that $\varphi[\succ'](i) = g$. Because agent j envies i at \succ and $\varphi[\succ](j) = g$, Lemma 6 implies that i gets at least g at $\succ' = (\succ_{-j}, \succ'_j)$. Hence, $\varphi[\succ'](i) \in \{h, g\}$. Furthermore, Lemma 5 tells us that j cannot envy i at \succ' . Hence, $\varphi[\succ'](i) = g$.

Second, we show that $\varphi[\succ'](k) = h$. Consider an auxiliary preference ranking $\tilde{\succ}_k \in \langle e, h, g, \dots \rangle$ that agrees with \succ_k except possibly for the relative ranking of g . Maskin monotonicity implies that $\varphi[\tilde{\succ}_k, \succ_{-k}] = \varphi[\succ]$. Thus, agent j envies k at $[\tilde{\succ}_k, \succ_{-k}]$ and $\varphi[\tilde{\succ}_k, \succ_{-k}](j) = g$, and 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 $[\tilde{\succ}_k, \succ_{-\{k,j\}}, \succ'_j]$. To prove that k gets h it is thus enough to show that i gets g at $[\tilde{\succ}_k, \succ_{-\{k,j\}}, \succ'_j]$. The proof is analogous to the proof of the equality $\varphi[\succ'](i) = g$ above: i gets at least g at $[\tilde{\succ}_k, \succ_{-\{k,j\}}, \succ'_j]$, and, because j cannot envy i at $[\tilde{\succ}_k, \succ_{-\{k,j\}}, \succ'_j]$ (by Lemma 5), 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, $\varphi[\tilde{\succ}_k, \succ_{-\{k,j\}}, \succ'_j] = \varphi[\succ_k, \succ_{-\{k,j\}}, \succ'_j] = \varphi[\succ']$. Thus, $\varphi[\succ'](k) = h$ and the claim is proved. \triangleleft

The above claim and Maskin monotonicity allow us to assume that $\succ_k \in \langle e, h, g, \dots \rangle$. We fix three additional auxiliary preference rankings: $\succ_i^* \in \langle h, e, g, \dots \rangle$, $\succ_{i'}^* \in \langle h, e, \dots \rangle$, and $\succ_k^* \in \langle e, g, h, \dots \rangle$.

CLAIM 2. (i) We have $\varphi[\succ_i^*, \succ_{-i}](i) = h$ and $\varphi[\succ_{i'}^*, \succ'_{-i'}](i') = h$.

(ii) We have $\varphi[\succ_i^*, \succ_{-i}](j) = g$.

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

(ii) By Maskin monotonicity for i , $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ']$.⁴⁵ 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, 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 nonbossiness 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 $\succ' = [\succ'_j, \succ_{-j}]$, and, by Lemma 5, i must get the same house at these two profiles. This contradiction proves Claim 2. \triangleleft

CLAIM 3. (i) We have $\varphi[\succ_k^*, \succ_{-k}](k) = g$.

(ii) We have $\varphi[\succ_k^*, \succ'_{-k}] = \varphi[\succ_k^*, \succ_{-k}]$.

⁴⁵In more detail, $\varphi[\succ'](i) \succeq_i^* a$ implies $\varphi[\succ'](i) \succeq_i^* a$ for all houses a ; hence, $[\succ_i^*, \succ'_{-i}]$ is a φ -monotonic transformation of \succ' , and the Maskin monotonicity of φ gives $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ']$.

PROOF. (i) 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{T}_σ ranks h over g . Thus, by Lemma 7 and Pareto efficiency, agent k should get g at $[\succ_k^*, \succ_{-k}]$.

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

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

PROOF. 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(ii), $\varphi[\succ_k^*, \succ_{-k}](i') \in \{e, h\}$. 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(i) 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(i), we must thus have $\varphi[\succ_i^*, \succ_{-i}](k) = g$. But this contradicts Claim 2(ii). \triangleleft

CLAIM 5. (i) We have $\varphi[\succ_k^*, \succ_{-k}] = \varphi[\succ_k^*, \succ_{i'}^*, \succ_{-\{k,i'\}}] = \varphi[\succ_k^*, \succ_{i'}^*, \succ'_{-i}, \succ_{-\{k,i',j\}}]$.

(ii) We have $\varphi[\succ_k^*, \succ_{i'}^* \succ'_{-\{k,i'\}}](k) = g$.

PROOF. The first equality of part (i) follows from Maskin monotonicity for i' and Claim 4. To prove the second equality of part (i), 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(i). Thus the second equality follows from Maskin monotonicity for j . Now, part (ii) of the claim follows from part (i) and Claim 3(i). \triangleleft

CLAIM 6. We have $\varphi[\succ_{i'}^*, \succ'_{-i'}](i) = e$.

PROOF. Strategy-proofness for k and Claim 5(ii) imply that agent k gets at least g at $(\succ_{i'}^*, \succ'_{-i'})$ but does not get e . By Claim 2(i), k gets g at $(\succ_{i'}^*, \succ'_{-i'})$. By nonbossiness for k and part (ii) of Claim 5,

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

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

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

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

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

PROOF. We 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.

Because \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 $[\succ_{\{i,i'\}}^*, \succ'_{-\{i,i'\}}] = [\succ_{\{i,i',j\}}^*, \succ_{-\{i,i',j\}}', \succ'_j]$. This, however, leads to a contradiction with Lemma 5, because Claim 2 and the second above-displayed equality imply that $\varphi[\succ_{\{i,i'\}}^*, \succ_{-\{i,i',j\}}, \succ_j](i) = h$. Thus, we have shown that $\varphi[\succ_{i'}^*, \succ'_{-i'}](i) \neq e$. \triangleleft

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

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

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

PROOF. 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, $\varphi[\succ^*] = \varphi[\succ']$ and $\varphi[\succ^*] = \varphi[\succ]$. Also $\succ^*, \succ'^* \in \mathbf{P}[\sigma \cup \{(i, e)\}, h, h']$. Thus, by Lemma 8, $\varphi[\succ^*]^{-1}(h) = \varphi[\succ'^*]^{-1}(h)$. Hence, $\varphi[\succ']^{-1}(h) = \varphi[\succ'^*]^{-1}(h) = \varphi[\succ^*]^{-1}(h) = \varphi[\succ]^{-1}(h)$. \triangleleft

CLAIM 2. If $h, h' \in \overline{H_\sigma} - \{e\}$, $h' \neq h$, $\succ \in \mathbf{P}[\sigma, e, h]$, $\succ' \in \mathbf{P}[\sigma, e, h', h]$, and $\varphi[\succ']^{-1}(e) \neq \varphi[\succ]^{-1}(e)$, then $\varphi[\succ'](k') = h'$ for agent $k' = \varphi^{-1}[\succ](h)$.

PROOF. 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' . By Claim 1, $\varphi[\succ^*]^{-1}(h') = \varphi[\succ']^{-1}(h') = k'$ and Maskin monotonicity gives

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

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

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

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

CLAIM 3. If $h, h' \in \overline{H_\sigma} - \{e\}$, $h \neq h'$, $\succ \in \mathbf{P}[\sigma, e, h]$, and $\succ' \in \mathbf{P}[\sigma, e, h', h]$, then $\varphi[\succ'](k') = h'$ for agent $k' = \varphi^{-1}[\succ](h)$.

PROOF. 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 available outcome $h'' \neq e$ at σ such that for some $\succ'' \in \mathbf{P}[\sigma, e, h'']$, we have $\varphi[\succ'']^{-1}(e) \neq \varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e)$. By Lemma 8, $h'' \neq h$, and we may assume that $\succ'' \in \mathbf{P}[\sigma, e, h'', h]$. By Claim 2, for $k' = \varphi[\succ]^{-1}(h)$ we have $\varphi[\succ''](k') = h''$. Claim 2 and $\varphi[\succ''](k') = h''$ hence give us $\varphi[\succ'](k') = h'$. \triangleleft

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. By Lemma 8, $\varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e)$. Because e is brokered* at σ , there exists some available outcome $h'' \neq e$ at σ 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, for $k' = \varphi[\succ]^{-1}(h)$, $\varphi[\succ''](k') = h''$, and as $\varphi[\succ''](k') = h''$ we have $\varphi[\succ'](k') = h$ by Claim 3 again. \triangleleft

To complete the proof of the lemma, notice that e being brokered* implies that there are at least two houses in $\overline{H_\sigma} - \{e\}$. Let $h, h' \in \overline{H_\sigma} - \{e\}$, $h' \neq h$, be any two such houses. Take $\succ^* \in \mathbf{P}[\sigma, e, h, h']$ and let $k' = \varphi[\succ^*]^{-1}(h)$. Claim 3 implies that $\varphi[\succ'](k') = h'$ for all $\succ' \in \mathbf{P}[\sigma, e, h']$, and Claim 4 implies that $k' = \varphi[\succ]^{-1}(h)$ for all $\succ \in \mathbf{P}[\sigma, e, h]$. Thus, k' obtains his second choice at any profile in $\mathbf{P}^*[\sigma, e]$ and, hence, this agent is the unique broker* of e at σ . \square

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. We 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 $\succ' \in \mathbf{P}[\sigma, h, g]$. Then $\varphi[\succ_i^*, \succ_{-i}](i) = g$ for all $\succ_i^* \in \langle g, \dots \rangle$ and all $\succ_{-i} \in \mathbf{P}_{-i}[\sigma, h]$.

PROOF. Let $\succ_{-i} \in \mathbf{P}_{-i}[\sigma, h]$. Take any $\succ_i \in \langle h, g, \dots \rangle$. If $\varphi[\succ](i) = h$, then Pareto efficiency and strategy-proofness imply that $\varphi[\succ_i^*, \succ_{-i}](i) = g$ for all $\succ_i^* \in \langle g, h, \dots \rangle$ and, furthermore, by strategy-proofness, for all $\succ_i^* \in \langle g, \dots \rangle$. 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 except for how the unmatched agents rank 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 $\succ_i^* \in \langle g, h, \dots \rangle$. By strategy-proofness, agent i gets at least h at $[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}]$; thus, by Pareto efficiency, agent i gets g . 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 that $\varphi[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}] = \varphi[\succ_i^*, \succ_{-i}]$. In particular, $\varphi[\succ_i^*, \succ_{-i}](i) = g$. \triangleleft

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 $\succ' \in \mathbf{P}[\sigma, h, g]$. If $\succ \in \mathbf{P}[\sigma, h]$ and there is some $\succ' \in \mathbf{P}[\sigma, h, g]$ such that $\succ_k \in \langle h, g, \dots \rangle$ for $k = \varphi[\succ']^{-1}(g)$, then $\varphi[\succ](i) = h$.

PROOF. By way of contradiction, suppose $\succ \in \mathbf{P}[\sigma, h]$ and there is some $\succ' \in \mathbf{P}[\sigma, h, g]$ such that $\succ_k \in \langle h, g, \dots \rangle$ for $k = \varphi[\succ']^{-1}(g)$ and yet $\varphi[\succ](i) \neq h$. By strategy-proofness, we can choose $\succ_i \in \langle h, g, \dots \rangle$. Furthermore, we can choose \succ and \succ' so that they only differ 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, $\succ_k^* \in \langle h, g, \dots \rangle$, and Maskin monotonicity implies that $\varphi[\succ^*](i) = h$.

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. Let h' be the house that j gets at \succ . Since $j \neq k$, it follows that $h' \neq g$. By strategy-proofness, $h' \neq h$. Thus, Maskin monotonicity allows us to assume that $\succ''_j \in \langle h, g, h', \dots \rangle$.

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 of 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 nonbossiness,

$\varphi[\succ''_j, \succ_{-j}](j) = g$. 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 $[\succ''_j, \succ_{-j}] = [\succ''_j, \succ'_{-j}]$. \triangleleft

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 $\varphi[\succ]^{-1}(h) \neq i$. For all unmatched houses $g \neq h$ at σ , define $\succ^{h,g}$ 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 who gets g_1 at \succ^{h,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 who gets g_2 at \succ^{h,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 show that i gets h at any profile $\succ' \in \mathbf{P}[\sigma, h]$ at which k_2 ranks g_1 second. Indeed, suppose $\succ'_{k_2} \in \langle h, g_1, \dots \rangle$ and i does not get h at \succ' . Let $\succ''_i \in \langle h, g_1, \dots \rangle$. 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 5. 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 who gets g_3 at \succ^{h,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 houses is finite, by repeating the procedure we arrive at an agent k_n who ranks one of the houses g_1, \dots, g_n second at \succ . That means that i gets h at \succ , a contradiction that concludes the proof. \square

The proof of Theorem 2 is a direct corollary of Lemmas 9 and 10 and the definitions of owned* and brokered* houses.⁴⁶

Lemmas 9 and 10 also imply the following theorem.

THEOREM 6. *If there is exactly one agent who owns* houses at σ , then there are no brokered* houses at σ .*

While there are no brokers* at submatchings with a single owner*, in R1–R6 we allow such a possibility. For any control rights obeying R1–R6, there is another one in which there are brokers only at submatchings with two or more owners. To construct such an equivalent control-rights structure, simply set the control rights to be equal to the original ones at all submatchings, except possibly submatchings in which there is exactly one owner; in the latter submatchings make this unique owner the owner of all houses. R6 guarantees that if there is a broker in a submatching with a single owner in the original control-rights structure, then in the submatching that matches this single owner with the originally brokered house, both the original and the equivalent control rights give the original broker the ownership right over the remaining houses.

⁴⁶The proof so far did not rely on the assumption that $|H| > |I|$.

C.2 The starred control-rights structure satisfies R1–R6

Before proving R1–R6, we state and prove one more auxiliary result.

LEMMA 11 (Relationship between brokerage* and ownership*). *Let $\sigma \in \overline{\mathcal{M}}$, agent k broker* house e at σ , $h \in \overline{H_\sigma} - \{e\}$, and $\succ \in \mathbf{P}[\sigma, e, h]$. Then there is an agent $i \in \overline{I_\sigma} - \{k\}$ such that $\varphi[\succ](i) = e$, and either (i) i is the owner* of h at σ or (ii) $|\overline{H_\sigma}| = 3$, and there exists a third agent in $\overline{I_\sigma}$ who brokers* h at σ and i brokers* the remaining third house at σ .⁴⁷*

PROOF. By Pareto efficiency, for any $\succ \in \mathbf{P}[\sigma, e, h]$ there is an agent $i \in \overline{I_\sigma} - \{k\}$ such that $i = \varphi[\succ]^{-1}(e)$.

CLAIM 1. *For all $\succ^{e,h} \in \mathbf{P}[\sigma, e, h]$, $\varphi[\succ^{e,h}](i) = e$ and $\varphi[\succ^{e,h}](k) = h$.*

The first claim follows from Lemma 8, and the second from Lemma 9.

Because k brokers* e at σ , Claim 1 implies that there exist agents $i, j \in \overline{I_\sigma}$ and houses $g, h \in \overline{H_\sigma}$ such that $j \neq i$ and $g \neq h$, as well as $\varphi[\succ^{e,h}](i) = e$ for all $\succ^{e,h} \in \mathbf{P}[\sigma, e, h]$ and $\varphi[\succ^{e,g}](j) = e$ for all $\succ^{e,g} \in \mathbf{P}[\sigma, e, g]$. Furthermore, $\varphi[\succ^{e,h}](k) = h$ and $\varphi[\succ^{e,g}](k) = g$. The pairs (i, h) and (j, g) play symmetric roles.

CLAIM 2. *For all $\succ_i^{h,e}$ and $\succ_{-i}^{e,h}$, $\varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}](i) = h$.*

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

CLAIM 3. *For all $\succ_k^{h,e}$ and $\succ_{-k}^{e,h}$, $\varphi[\succ_k^{h,e}, \succ_{-k}^{e,h}] = \varphi[\succ^{e,h}]$.*

PROOF. Because $\varphi[\succ^{e,h}](k) = h$, profile $[\succ_k^{h,e}, \succ_{-k}^{e,h}]$ is a monotonic transformation of $\succ^{e,h}$; thus, by Maskin monotonicity of φ , we have the desired result. ◁

CLAIM 4. *For all $\succ_{\{i,k\}}^{h,e}$ and $\succ_{-\{i,k\}}^{e,h}$, if $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) = h$, then $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](k) \neq e$.*

PROOF. For an indirect argument, suppose that $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) = h$ and $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](k) = e$ for some $\succ_{\{i,k\}}^{h,e}$ and $\succ_{-\{i,k\}}^{e,h}$. Then $\varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}](k) = e$ by the strategy-proofness of φ . Because e is a brokered* house at σ , there exists some available outcome $g \notin \{e, h\}$ at σ and some preference profile $\succ^{e,g}$ such that $\varphi[\succ^{e,g}](j) = e$ for some available agent $j \notin \{i, k\}$ at σ . By Lemma 8, we may assume that each agent $i' \in \overline{I_\sigma}$ ranks outcomes other than g and h in the same way at $\succ_{i'}^{e,g}$ and $\succ_{i'}^{e,h}$ and that $\succ^{e,g} = \succ^{e,g,h}$ for some profile $\succ^{e,g,h}$. Because k is the broker* of e at σ , we have $\varphi[\succ^{e,g,h}](k) = g$. By Maskin monotonicity, $\varphi[\succ^{e,g,h}] = \varphi[\succ_{\{i,k\}}^{e,g,h}, \succ_{-\{i,k\}}^{e,h}]$. Now i gets a house weakly worse than h at $[\succ_{\{i,k\}}^{e,g,h}, \succ_{-\{i,k\}}^{e,h}]$. However, if i and k manipulated and submitted $\succ_{\{i,k\}}^{h,e}$ instead

⁴⁷Case (ii) in Lemma 11 cannot happen when $|H| > |I|$. Case (ii) is relevant when we relax this assumption in the Supplement.

of $\succ_{\{i,k\}}^{e,g,h}$, they would get h and e , respectively, at $[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}]$. Both agents weakly improve, while k strictly improves. This contradicts the fact that φ is group strategy-proof. \triangleleft

CLAIM 5. For all $\succ_{\{i,k\}}^{h,e}$ and $\succ_{-\{i,k\}}^{e,h}$, $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) = h$.

PROOF. By way of contradiction, suppose $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) \neq h$. By Claim 3, $\varphi[\succ_k^{h,e}, \succ_{-k}^{e,h}](i) = \varphi[\succ^{e,h}](i) = e$, and by strategy-proofness of φ , $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) = e$. The nonbossiness of φ and Claim 3 imply that $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}] = \varphi[\succ_k^{h,e}, \succ_{-k}^{e,h}] = \varphi[\succ^{e,h}]$. By Claim 1, $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](k) = h$, and by Maskin monotonicity of φ and Claim 1, $\varphi[\succ^{h,e}](k) = h$. Hence, h is either owned* by k or brokered*.

Suppose h is owned* by k . As k also brokers* e , there exists some available outcome $g \neq h, e$ at σ and agent $j \in \overline{I_\sigma} - \{i, k\}$ such that for all $\succ^{e,g,h}$, $\varphi[\succ^{e,g,h}](j) = e$ and $\varphi[\succ^{e,g,h}](k) = g$. By Maskin monotonicity, $\varphi[\succ_k^{e,g,h}, \succ_i^{h,e}, \succ_{-\{i,k\}}^{e,h}] = \varphi[\succ^{e,g,h}]$. Because k owns* h , Maskin monotonicity and strategy-proofness for k give $\varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}](k) \succeq_k^{e,h} h$. By Claim 2, $\varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}](i) = h$ and, thus, $\varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}](k) = e$. By Maskin monotonicity, $\varphi[\succ_k^{e,g,h}, \succ_i^{h,e}, \succ_{-\{i,k\}}^{e,h}] = \varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}]$. Thus, $\varphi[\succ^{e,g,h}] = \varphi[\succ_i^{h,e}, \succ_{-i}^{e,h}]$, but the argument also shows that k gets different houses under these two profiles; a contradiction.

Suppose h is brokered* and consider the three possible subcases:

- Agent i brokers* h . Then there exists a house $g' \neq h, e$ and an agent $j' \neq k, i$ such that for all $\succ^{h,g',e}$, we have $\varphi[\succ^{h,g',e}](i) = g'$ and $\varphi[\succ^{h,g',e}](j') = h$. Thus, $e \succeq_k^{h,g',e} \varphi[\succ^{h,g',e}](k)$. However, both i and k improve, and at least one of them strictly improves if they report $\succ_{\{i,k\}}^{h,e}$, since $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{h,g',e}](i) = h$ and $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{h,g',e}](k) = e$. This is a contradiction to group strategy-proofness of φ .
- Agent k brokers* h . Then $\varphi[\succ^{h,e}](k) = e$, contradicting $\varphi[\succ^{h,e}](k) = h$.
- Some $j' \notin \{i, k\}$ brokers* h . Then $\varphi[\succ^{h,e}](j') = e$ for all $\succ^{h,e}$. But $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](i) = e$ and $\varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}](k) = h$ (see the beginning of the proof of Claim 5). By Maskin monotonicity of φ , $\varphi[\succ^{h,e}] = \varphi[\succ_{\{i,k\}}^{h,e}, \succ_{-\{i,k\}}^{e,h}]$, contradicting $\varphi[\succ^{h,e}](j') = e$. \triangleleft

CLAIM 6. For all $\succ_{\{i,k\}}^{h,e,g}$ and $\succ_{-\{i,k\}}^{e,h,g}$, $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](i) = h$, $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](k) = g$, and $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](j) = e$.

PROOF. Claim 5 implies that $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](i) = h$. By way of contradiction, suppose the other parts of Claim 6 are not true. Claim 4 then implies $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](k) \neq e$. Thus, $g \succeq_k^{h,e,g} \varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](k)$ and $e \succeq_j^{e,h,g} \varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](j)$, where at least one preference is strict. Thus, k and j can jointly misreport $\succ_{\{k,j\}}^{e,g}$ and secure $\varphi[\succ_i^{h,e,g}, \succ_{\{k,j\}}^{e,g}, \succ_{-\{i,j,k\}}^{e,h,g}](j) = e$ and $\varphi[\succ_i^{h,e,g}, \succ_{\{k,j\}}^{e,g}, \succ_{-\{i,j,k\}}^{e,h,g}](k) = g$. This is a contradiction to the group strategy-proofness of φ . \triangleleft

CLAIM 7. For all $\succ_{\{i,j\}}^{h,e}$ and $\succ' \in \mathbf{P}[\sigma]$, $\varphi[\succ_{\{i,j\}}^{h,e}, \succ'_{-\{i,j\}}](i) = h$ and $\varphi[\succ_{\{i,j\}}^{h,e}, \succ'_{-\{i,j\}}](j) = e$.

PROOF. The strategy-proofness for j and Claim 6 (i.e., $\varphi[\succ_{\{i,k\}}^{h,e,g}, \succ_{-\{i,k\}}^{e,h,g}](j) = e$) together with Maskin monotonicity imply that

$$\varphi[\succ_{\{i,j,k\}}^{h,e,g}, \succ'_{-\{i,j,k\}}](j) \in \{e, h\}.$$

Also $\varphi[\succ_i^{h,e,g}, \succ_{-i}^{e,h,g}](i) \in \{h, e\}$ because i can get e by reporting $\succ_i^{e,h,g}$. If now $\varphi[\succ_i^{h,e,g}, \succ_{-i}^{e,h,g}](i) = e$, then by reporting $\succ_i^{e,h,g}$, agent i does not change his assignment, and by nonbossiness, everybody else's assignment is unchanged and, therefore, $\varphi[\succ_i^{h,e,g}, \succ_{-i}^{e,h,g}](k) = h$. But this would contradict the above-displayed statement. Hence, $\varphi[\succ_i^{h,e,g}, \succ_{-i}^{e,h,g}](i) = h$ and the above-displayed statement implies that $\varphi[\succ_{\{i,j,k\}}^{h,e,g}, \succ'_{-\{i,j,k\}}](j) = e$. By Maskin monotonicity, the claim follows. \square

Claim 7 implies that i and j are the only agents who can get e when agents report some profile $\succ^{e,g'}$. Indeed, if not, then there would be a j' for whom Claim 7 would also be true (in place of j), which is not possible. Claim 7 also implies that either i owns* h or j brokers* h at σ .

If i owns* h at σ , then the first case of the lemma obtains. Suppose thus that j brokers* h at σ . By symmetry, all claims hold for j and g instead of i and h . In particular, Claim 7 implies that j owns* g or i brokers* g . To conclude the proof let us first show that j (who brokers* h) cannot own* g . By way of contradiction, suppose that j does. Then by Maskin monotonicity, j gets g at $[\succ_j^g, \succ_{-j}]$ for all $\succ \in \mathbf{P}[\sigma]$. Let $j^* \neq j$ be such that he gets h at $\succ^{h,g}$ and let $j'' \neq j, j^*$ be such that he gets h at $\succ^{h,g'}$ for some house $g' \neq g$; because h is brokered* by j , Claim 1 (applied to j and h instead of k and e) implies that such agents j^* and j'' exist. Now, Claim 7 applied to agents j, j^*, j'' and houses h, g instead of agents k, i, j and houses e, h (respectively) implies that $\varphi[\succ_{\{j^*,j''\}}^{g,h}, \succ_{-\{j^*,j''\}}](j^*) = g$ for all $\succ \in \mathbf{P}[\sigma]$. This is a contradiction as $j \neq j^*$ gets g at all profiles $[\succ_j^g, \succ_{-j}]$ and this set of profiles contains all $[\succ_{\{j^*,j''\}}^{g,h}, \succ_{-\{j^*,j''\}}]$. Therefore, i brokers* g .

Since k brokers* e and, in the case now considered, j brokers* h and i brokers* g , it is enough to show that $|\overline{H}_\sigma| = 3$ to prove that the second part of the lemma obtains. By way of contradiction, suppose that there exists some other house $g' \in \overline{H}_\sigma \setminus \{e, g, h\}$. Because $j' \in \{i, j\}$ is the agent who gets e when everybody reports $\succ^{e,g'}$, either j' owns* g' or $j'' \in \{i, j\} \setminus \{j'\}$ brokers* g' . We just showed that both i and j are brokers*. By the above argument for why j cannot own* while simultaneously brokering*, i cannot own* any houses either. We next show that neither of i or j can broker* any additional house. There are two agents $i' \neq i''$ different from i such that $\varphi[\succ^{g,g'}](i') = g$ and $\varphi[\succ^{g,g''}](i'') = g$ for a house $g'' \in \overline{H}_\sigma - \{g, g'\}$ (as i brokers* g at σ). Then one of i' or i'' owns* or brokers* g' . Because each house has a unique owner* or broker*, i cannot broker* g' . A similar argument, using h instead of g , shows that j cannot broker* g' . This is a contradiction to the existence of such a house g' . Hence, $|\overline{H}_\sigma| = 3$ and the second case of the lemma obtains. \square

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

LEMMA 12 (R1, uniqueness of a brokered* house). *Let $\sigma \in \overline{\mathcal{M}}$. There is at most one brokered* house at σ .*

PROOF. Suppose e is a brokered* house at σ . By Lemma 9, house e has a broker* at σ . Because $|H| > |I|$, Lemma 11 implies that each house $h \in \overline{I_\sigma} - \{e\}$ is owned*. \square

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

PROOF. 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 . \square

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. 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, by Lemma 11, i gets h at all $\succ \in \mathbf{P}^*[\sigma, h]$, a contradiction as $k \neq i$ owns* h . \square

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

PROOF. 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 $j \in I_{\sigma'} - I_\sigma - \{i\}$ gets h at \succ' . Take $\succ \in \mathbf{P}[\sigma, h]$ such that each agent $k \in I_{\sigma'} - I_\sigma$ ranks $\sigma'(k)$ as his second choice, and $\succ_k = \succ'_k$ for agents $k \notin I_{\sigma'} - I_\sigma$. Each $k \in I_{\sigma'} - I_\sigma$ gets $\sigma'(k)$ under \succ' by Lemma 7, and he gets the same house under \succ by the Pareto efficiency of φ and the fact that $\varphi[\succ](i) = h$. Thus, these agents are indifferent between \succ and \succ' . This contradicts group strategy-proofness of φ because $\succ' = \succ$ for agents not in $I_{\sigma'} - I_\sigma$ and because j is strictly better off under \succ' . \square

LEMMA 16 (R5, part 1, limited persistence of brokerage*). *Let $\sigma, \sigma' \in \overline{\mathcal{M}}$ be such that $\sigma' \supseteq \sigma$. 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. By Lemma 11, agent i gets e at all $\succ \in \mathbf{P}[\sigma, e, h]$, agent i' gets e at all $\succ \in \mathbf{P}[\sigma, e, h']$, and agent k gets h and h' , respectively. 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^{ih} 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 $\succ^{ih'}$ be obtained analogously from $\succ^{h'}$. By Maskin monotonicity, $\varphi[\succ^{ih}]^{-1}(e) = i \neq i' = \varphi[\succ^{ih'}]^{-1}(e)$. Because $\succ^{ih}, \succ^{ih'} \in \mathbf{P}^*[\sigma', e]$, house e is a brokered* house at σ' .

We complete the proof by way of contradiction. 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. Notice that

$$\varphi[\succ'](k') = h,$$

because $k' \neq k$ brokers* e at σ' . Let $\succ \in \mathbf{P}[\sigma, e, h]$ 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 in \succ and \succ' . Because 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 by Maskin monotonicity, $\varphi[\succ'] = \varphi[\succ]$. Thus, $\varphi[\succ'](k) = h$, which contradicts the above-displayed equation. \square

LEMMA 17 (R5, part 2, 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. By Lemmas 10 and 11 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle e, \dots \rangle$ and $\succ_k \in \langle e, h, \dots \rangle$, we have $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$. Thus, Theorem 2 implies that either i owns* e at σ' or k brokers* e at σ' . By assumption, the latter is not true; hence i owns* e at σ' . \square

LEMMA 18 (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 owns* h , and j controls* g at σ . Then j owns* h at $\sigma' = \sigma \cup \{(i, g)\}$.*

PROOF. Take any $\succ \in \mathbf{P}[\sigma', h]$. We need to show that $\varphi[\succ](j) = h$. By Lemma 7 and Maskin monotonicity, we may assume that $\succ_i \in \langle g, h, \dots \rangle$. There are two cases.

Case j owns* g at σ . Let \succ' be such that j ranks g second and otherwise all agents rank the houses as in \succ . Because at any $\succ'' \in \mathbf{P}[\sigma]$ at which i ranks h at the top and j ranks g at the top, they receive these houses, the group strategy-proofness and Pareto efficiency of φ imply that $\varphi[\succ'](j) = h$, and, by Maskin monotonicity, $\varphi[\succ](j) = h$ as required.

Case j brokers* g at σ . Let \succ' be such that i ranks h second, all agents unmatched at σ' put g at the top, and otherwise all agents rank houses as in \succ . Lemma 11 implies that $\varphi[\succ'](i) = g$ and $\varphi[\succ'](j) = h$. By Maskin monotonicity, $\varphi[\succ](j) = h$ as required. \square

C.3 The TC mechanism defined by the starred control-rights structure equals φ

We showed that the starred control-rights structure (c, b) is well defined and satisfies R1–R6. To close the proof of Theorem 1, we now show that the resulting TC mechanism, $\psi^{c,b}$, maps preferences to outcomes in the same way as φ does.

Fix $\succ \in \mathbf{P}$. For each agent i , let h_i be the unique house that points to him and is removed in the same cycle as i under $\psi^{c,b}[\succ]$. The auxiliary preference profile \succ^* is defined as follows.

- 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 agent $j \neq i$ in the cycle of i is assigned a brokered house, then we construct \succ_i^* from \succ_i by moving h_i just after $\psi^{c,b[\succ]}(i)$.⁴⁸
- If some agent $j \neq i$ in the cycle of i is assigned a brokered house e , then we construct \succ_i^* from \succ_i by moving e just after $\psi^{c,b[\succ]}(i)$ and by moving h_i just after e .
- If i was removed as a broker, then we construct \succ_i^* from \succ_i by moving h_j just below $\psi^{c,b[\succ]}(i)$, where j is the agent in the cycle of i who was assigned the brokered house.

By Maskin monotonicity, $\psi^{c,b[\succ]} = \psi^{c,b[\succ^*]}$ and $\varphi[\succ^*] = \varphi[\succ]$. Thus, to show that $\varphi[\succ] = \psi^{c,b[\succ]}$, it is sufficient to prove that $\psi^{c,b[\succ^*]} = \varphi[\succ^*]$. Denoting by I^s the set of agents removed in round s of $\psi^{c,b[\succ^*]}$, we can restate this equality as

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

The claim is empty and true for $r = 0$, and we prove it by induction on rounds r of $\psi^{c,b[\succ^*]}$.

Fix $r \geq 1$ and let σ^{r-1} be the submatching fixed before round r (in particular, $\sigma^0 = \emptyset$). For the inductive step, assume that $\varphi[\succ^*](i) = \psi^{c,b[\succ^*]}(i)$ for all $i \in \bigcup_{s \leq r-1} I^s = I_{\sigma^{r-1}}$.

CLAIM 1. *We have $\varphi[\succ^*](i) \succeq_i^* h_i$ for $i \in I^r$ who own h_i in round r of $\psi^{c,b[\succ^*]}$.*

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

By the inductive assumption, Lemma 7, and Maskin monotonicity,

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

Because h_i is owned by i at σ^{r-1} under $\psi^{c,b}$, the construction of the control-rights structure (c, b) implies that h_i is owned* by i at σ^{r-1} in φ , and, hence,

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

and at \succ' agents from $\overline{I_{\sigma^{r-1}}}$ do not obtain h_i and, by Lemma 7, they do not obtain any houses from $H_{\sigma^{r-1}}$ either. Thus, Maskin monotonicity implies that

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

⁴⁸Note that i might be assigned the brokered house. We do not change the ranking of houses other than h_i .

The first-displayed equation of this proof, the strategy-proofness of φ , and the third- and second-displayed equations give us

$$\begin{aligned} \varphi[>^*](i) &= \varphi[>''_{(I-I_{\sigma^{r-1}})-\{i\}}, >'_{I_{\sigma^{r-1}}}, >^*_i](i) \succeq_i^* \varphi[>''_{(I-I_{\sigma^{r-1}})-\{i\}}, >'_{I_{\sigma^{r-1}} \cup \{i\}}](i) \\ &= \varphi[>'](i) = h_i. \end{aligned} \quad \square$$

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

PROOF. The inductive assumption implies that all houses better than $\psi^{c,b}[>^*](i)$ are already given to other agents; hence, $\psi^{c,b}[>^*](i) \succeq_i^* \varphi[>^*](i)$. By way of contradiction, suppose $\varphi[>^*](i) \neq \psi^{c,b}[>^*](i)$. Then Claim 1 and the construction of $>^*$ imply that $\varphi[>^*](i) = h_i$. Let $h_i \rightarrow i \rightarrow h_{i^2} \rightarrow i^2 \rightarrow \dots \rightarrow h_{i^n} \rightarrow i^n \rightarrow h_i$ be the cycle in which i is removed under $\psi^{c,b}[>^*]$. From $\varphi[>^*](i) = h_i = \psi^{c,b}[>^*](i^n)$, we conclude that $\varphi[>^*](i^n) \neq \psi^{c,b}[>^*](i^n)$, and Claim 1 and the construction of $>^*$ imply that $\varphi[>^*](i^n) = h_{i^n} = \psi^{c,b}[>^*](i^{n-1})$. As we continue iteratively, we obtain $\varphi[>^*](j) = h_j$ for all $j \in \{i, i^n, \dots, i^2\}$. Hence, the matching obtained by assigning $\psi^{c,b}[>^*](j)$ to each agent $j \in \{i, i^2, \dots, i^n\}$ and $\varphi[>^*](j)$ to each agent $j \in I - \{i, i^2, \dots, i^n\}$ Pareto dominates $\varphi[>^*]$ at $>^*$, contradicting Pareto efficiency of $\varphi[>^*]$. \square

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

PROOF. Let e be the brokered house and let $k \equiv i^{n+1} \equiv i^0$ be the broker at σ^{r-1} . Let

$$h_{i^1} \rightarrow i^1 \rightarrow h_{i^2} \rightarrow \dots \rightarrow i^n \rightarrow e \rightarrow k \rightarrow h_{i^1}$$

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

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

We continue in two steps.

Step 1. $\varphi[>^*](i^n) = e$. Suppose not. Then (2) gives us $\varphi[>^*](i^n) = h_{i^n}$. Note that $\varphi[>^*](i^\ell) \neq e$ as otherwise agents i^n and i^{n-1} would be better off by swapping their allocations, contrary to Pareto efficiency. Thus, (2) gives us $\varphi[>^*](i^\ell) = h_{i^\ell}$. Iterating this argument, we show that $\varphi[>^*](i^\ell) = h_{i^\ell}$ for $\ell = n, n-1, \dots, 1$.

Let $>' \in \mathbf{P}[\sigma^{r-1}]$ rank houses from $\overline{H}_{\sigma^{r-1}}$ the same way as $>^*$, except that i^1 ranks h^2 below all other houses from $\overline{H}_{\sigma^{r-1}}$. By our assumptions, k is the broker of e and i^1 is the owner of h^1 at σ^{r-1} . By construction of the control-rights structure (c, b) , agent k brokers* e and agent i^1 owns* h^1 at σ^{r-1} in φ . Because $>'_k \in \mathbf{P}_k[\sigma^{r-1}, h_{i^1}, \dots] \cup \mathbf{P}_k[\sigma^{r-1}, e, h_{i^1}, \dots]$ and $>'_{i^1} \in \mathbf{P}_k[\sigma^{r-1}, e, h_{i^1}, \dots]$, we conclude that $\varphi[>'](i^1) = e$. The inductive assumption and the conclusion of the first paragraph of this step imply that $>'$ is a φ -Maskin-monotone transformation of $>^*$, and, hence, $\varphi[>^*](i^1) = e$. But we have shown above that $\varphi[>^*](i^1) = h^1 \neq e$. This contradiction proves the claim of Step 1.

Step 2. $\varphi[>^*](i^\ell) = h_{i^{\ell+1}} = \psi^{c,b}[>^*](i^\ell)$ for all $\ell = 0, \dots, n-1$. By way of contradiction, suppose there is $m \in \{0, \dots, n-1\}$ such that $\varphi[>^*](i^m) \neq h_{i^{m+1}}$. Then, (2) and Step 1 imply that $\varphi[>^*](i^m) = h_{i^m}$. Thus, $\varphi[>^*](i^{m-1}) \neq h_{i^{(m-1)+1}}$. Iterating this argument, we show that $\varphi[>^*](i^\ell) = h_{i^\ell}$ for $\ell = m, m-1, \dots, 1$.

Let $>' \in \mathbf{P}[\sigma^{r-1}]$ rank houses from $\overline{H}_{\sigma^{r-1}}$ the same way as $>^*$, except that $k \equiv i^0$ ranks h^1 below all other houses from $\overline{H}_{\sigma^{r-1}}$. By our assumptions, k is the broker of e and i^1 is the owner of h^1 at σ^{r-1} . By construction of the control-rights structure (c, b) , agent k brokers* e and agent i^n owns* h^n at σ^{r-1} in φ . Because $>'_k \in \mathbf{P}_k[\sigma^{r-1}, h_{i^n}, \dots] \cup \mathbf{P}_k[\sigma^{r-1}, e, h_{i^n}, \dots]$ and $>'_{i^1} \in \mathbf{P}_k[\sigma^{r-1}, e, h_{i^n}, \dots]$, we get $\varphi[>'](k) = h_{i^n}$. The conclusion of the first paragraph of the current step implies that $\varphi[>^*](k) \neq h_{i^1}$; given the inductive assumption $>'$ is a φ -Maskin-monotone transformation of $>^*$ and, hence, $\varphi[>^*](k) = h_{i^n}$ as well. As a result, (2) and Step 1 imply that $\varphi[>^*](i^{n-1}) = h_{i^{n-1}}$. Thus, $\varphi[>^*](i^{n-2}) \neq h_{i^{n-1}}$. Iterating this argument, we conclude that $\varphi[>^*](i^\ell) = h_{i^\ell}$ for $\ell = n-1, n-2, \dots, 1$. Together with Step 1, this conclusion implies that $\varphi[>^*]$ is Pareto dominated by the allocation in which each agent i^m , $m = 0, \dots, n-1$, gets house h^{m+1} and all other agents get their $\varphi[>^*]$ houses. This contradiction proves Step 2 and Claim 3. \square

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

REFERENCES

- Abdulkadiroğlu, Atila and Yeon-Koo Che (2010), “The role of priorities in assigning indivisible objects: A characterization of top trading cycles.” Unpublished paper, Duke University and Columbia University. [304]
- Abdulkadiroğlu, Atila and Tayfun Sönmez (1999), “House allocation with existing tenants.” *Journal of Economic Theory*, 88, 233–260. [288, 289, 293, 297, 301]
- Abdulkadiroğlu, Atila and Tayfun Sönmez (2003), “School choice: A mechanism design approach.” *American Economic Review*, 93, 729–747. [297]
- Ashlagi, Itai and Alvin E. Roth (2014), “Free riding and participation in large scale, multi-hospital kidney exchange.” *Theoretical Economics*, 9, 817–863. [288]
- Bade, Sophie (2014), “Pareto optimal, strategyproof and non-bossy matching mechanisms.” Unpublished paper, Royal Holloway University. [301]
- Barberà, Salvador, Faruk Gul, and Ennio Stacchetti (1993), “Generalized median voter schemes and committees.” *Journal of Economic Theory*, 61, 262–289. [290]
- Barberà, Salvador, Matthew O. Jackson, and Alejandro Neme (1997), “Strategy-proof allotment rules.” *Games and Economic Behavior*, 18, 1–21. [290]
- Bogomolnaia, Anna, Rajat Deb, and Lars Ehlers (2005), “Strategy-proof assignment on the full preference domain.” *Journal of Economic Theory*, 123, 161–186. [291]
- Clarke, Edward H. (1971), “Multipart pricing of public goods.” *Public Choice*, 11, 17–33. [290]

- Dur, Umut Mert (2013), "A characterization of the top trading cycles mechanism in the school choice problem." Unpublished paper, North Carolina State University. [304]
- Dur, Umut Mert and M. Utku Ünver (2015), "Two-sided matching via balanced exchange: Tuition and worker exchanges." SSRN Working Paper 2180357. [304]
- Ehlers, Lars (2002), "Coalitional strategy-proof house allocation." *Journal of Economic Theory*, 105, 298–317. [291]
- Ehlers, Lars and Bettina Klaus (2004), "Resource-monotonicity for house allocation problems." *International Journal of Game Theory*, 32, 545–560. [297]
- Ehlers, Lars and Bettina Klaus (2007), "Consistent house allocation." *Economic Theory*, 30, 561–574. [303]
- Ehlers, Lars, Bettina Klaus, and Szilvia Pápai (2002), "Strategy-proofness and population-monotonicity for house allocation problems." *Journal of Mathematical Economics*, 38, 329–339. [297]
- Ergin, Haluk (2000), "Consistency in house allocation problems." *Journal of Mathematical Economics*, 34, 77–97. [303]
- Green, Jerry and Jean-Jacques Laffont (1977), "Characterization of satisfactory mechanisms for revelation of preferences for public goods." *Econometrica*, 45, 427–438. [290]
- Groves, Theodore (1973), "Incentives in teams." *Econometrica*, 41, 617–631. [290]
- Hakimov, Rustamdjan and Onur Kesten (2014), "The equitable top trading cycles mechanism for school choice." WZB Discussion Paper SP II 2014-210. [305]
- Harless, Patrick (2015), "From behind the veil: Evaluating allocation rules by ex-ante properties." Unpublished paper, University of Rochester. [305]
- Harless, Patrick and Vikram Manjunath (2016), "Learning matters: Reappraising object allocation rules when agents strategically investigate." Unpublished paper, University of Glasgow and Texas AM University. [305]
- Holmström, Bengt (1979), "Groves' scheme on restricted domains." *Econometrica*, 47, 1137–1144. [290]
- Lorenz, Max O. (1905), "Methods of measuring the concentration of wealth." *Publications of the American Statistical Association*, 9, 209–219. [305]
- Ma, Jinpeng (1994), "Strategy-proofness and the strict core in a market with indivisibilities." *International Journal of Game Theory*, 23, 75–83. [290]
- Maskin, Eric (1999), "Nash equilibrium and welfare optimality." *Review of Economic Studies*, 66, 23–38. [293]
- Morrill, Thayer (2013), "An alternative characterization of the deferred acceptance algorithm." *International Journal of Game Theory*, 42, 19–28. [304]
- Morrill, Thayer (2015a), "Making just school assignments." *Games and Economic Behavior*, 92, 18–27. [305]

Moulin, Hervé (1980), “On strategy-proofness and single-peakedness.” *Public Choice*, 35, 437–455. [290]

Pápai, Szilvia (2000), “Strategyproof assignment by hierarchical exchange.” *Econometrica*, 68, 1403–1433. [288, 289, 290, 291, 292, 293, 294, 297, 298, 299, 301, 303, 304, 305, 311, 312]

Pápai, Szilvia (2007), “Exchange in a general market with indivisible goods.” *Journal of Economic Theory*, 132, 208–235. [290]

Pathak, Parag A. and Tayfun Sönmez (2008), “Leveling the playing field: Sincere and sophisticated players in the Boston mechanism.” *American Economic Review*, 98, 1636–1652. [288]

Pycia, Marek (2016), “Swaps on networks.” SSRN Working Paper 2735524. [290, 304]

Pycia, Marek and M. Utku Ünver (2014), “Outside options in neutral discrete resource allocation.” Unpublished paper, UCLA and Boston College. [303]

Pycia, Marek and M. Utku Ünver (2011), “Trading cycles for school choice.” SSRN Working Paper 1899344. [297]

Pycia, Marek and M. Utku Ünver (2016), “Arrovian efficiency in allocation of discrete resources.” Unpublished paper, UCLA and Boston College. [305]

Roth, Alvin E. (1982), “Incentive compatibility in a market with indivisible goods.” *Economics Letters*, 9, 127–132. [289, 294, 295]

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver (2004), “Kidney exchange.” *Quarterly Journal of Economics*, 119, 457–488. [288, 297, 302]

Satterthwaite, Mark A. and Hugo Sonnenschein (1981), “Strategy-proof allocation mechanisms at differentiable points.” *Review of Economic Studies*, 48, 587–597. [292, 294]

Shapley, Lloyd S. and Herbert E. Scarf (1974), “On cores and indivisibility.” *Journal of Mathematical Economics*, 1, 23–37. [288]

Sönmez, Tayfun and M. Utku Ünver (2006), “Kidney exchange with good samaritan donors: A characterization.” Unpublished Paper, Boston College and University of Pittsburgh. [288, 302, 303]

Sönmez, Tayfun and M. Utku Ünver (2013), “Market design for kidney exchange.” In *The Handbook of Market Design* (Zvika Neeman, Nir Vulkan, and Alvin E. Roth, eds.), 93–137, Oxford University Press. [288]

Sprumont, Yves (1991), “The division problem with single-peaked preferences: A characterization of the uniform allocation rule.” *Econometrica*, 59, 509–519. [290]

Svensson, Lars-Gunnar (1994), “Queue allocation of indivisible goods.” *Social Choice and Welfare*, 11, 323–330. [294]

Svensson, Lars-Gunnar (1999), “Strategy-proof allocation of indivisible goods.” *Social Choice and Welfare*, 16, 557–567. [290, 303]

Takamiya, Koji (2001), “Coalition strategy-proofness and monotonicity in Shapley–Scarf housing markets.” *Mathematical Social Sciences*, 41, 201–213. [293]

Tang, Qianfeng and Yongchao Zhang (2016), “Hierarchical exchange rules and the core in indivisible objects allocation.” Unpublished paper, Shanghai University of Finance and Economics. [290, 293, 304]

Velez, Rodrigo A. (2014), “Consistent strategy-proof assignment by hierarchical exchange.” *Economic Theory*, 56, 125–156. [303]

Vickrey, William (1961), “Counterspeculation, auctions and competitive sealed tenders.” *Journal of Finance*, 16, 8–37. [290]

Warmbir, Steve (2003), “UIC hospital sued for medicare fraud facility accused of rigging list of patients awaiting new livers.” Chicago Sun-Times. July 29. HighBeam Research. Available at <https://www.highbeam.com/doc/1P2-1497626.html>. [288]

Co-editor George J. Mailath handled this manuscript.

Manuscript received 11 June, 2015; final version accepted 17 February, 2016; available online 23 February, 2016.