Incentive Compatible Allocation and Exchange of Discrete Resources*

Marek Pycia† M. Utku Ünver‡
UCLA *Boston College*

This Draft: February 2014

Abstract

Allocation and exchange of discrete resources such as kidneys, school seats, and many other resources for which agents have single-unit demand is conducted via direct mechanisms without monetary transfers. Incentive compatibility and efficiency are primary concerns in designing such mechanisms. We show that a mechanism is individually strategy-proof and always selects the efficient outcome with respect to some Arrovian social welfare function if and only if the mechanism is group strategy-proof and Pareto efficient. We construct the full class of these 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 strictly Lorenz dominate all previously known mechanisms.

Keywords: Individual strategy-proofness, group strategy-proofness, Pareto efficiency, Arrovian preference aggregation, matching, no-transfer allocation and exchange, single-unit demand.

JEL classification: C78, D78

---

*We thank numerous seminar and conference participants as well as Andrew Atkeson, Sophie Bade, Salvador Barbera, Haluk Ergin, Manolis Galenianos, Ed Green, Matthew Jackson, Philippe Jehiel, Onur Kesten, Fuhito Kojima, Sang-Mok Lee, Vikram Manjunath, Szilvia Pápai, Martine Quinzii, Al Roth, Andrzej Skrzypacz, Tayfun Sönmez, William Thomson, Özgür Yılmaz, William Zame, and three anonymous referees for comments. Kenny Mirkin, Kyle Woodward, and Simpson Zhang provided excellent research assistance. Ünver gratefully acknowledges the research support of the Microsoft Research Lab in New England.

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

‡Boston College, Department of Economics, 140 Commonwealth Ave., Chestnut Hill, MA 02467.

1 Introduction

Microeconomic theory has informed the design of many markets and other institutions. Recently, many new mechanisms have been proposed to allocate resources in environments in which agents have single-unit demands and transfers are not used, or are prohibited. These environments include: allocation and exchange of transplant organs, such as kidneys (cf. Roth, Sönmez, and Ünver, 2004); allocation of school seats in Boston, New York City, Chicago, and San Francisco (cf. Abdulkadiroğlu and Sönmez, 2003); and allocation of dormitory rooms at US colleges (cf. Abdulkadiroğlu and Sönmez, 1999).

The central concerns in the development of allocation mechanisms are incentives and efficiency.[1] We study two incentive compatibility requirements, individual strategy-proofness and group strategy-proofness, and two efficiency requirements, Pareto efficiency and efficiency with respect to an Arrovian social welfare function (Arrovian efficiency). We show that three among four possible combinations of these requirements are equivalent: a direct mechanism is individually strategy-proof and Arrovian efficient if and only if it it is group strategy-proof and Pareto efficient, and also, if and only if it is group strategy-proof and Arrovian efficient. We construct the full class of these mechanisms and analyze the implications of our characterization, as well as the usefulness of the newly constructed mechanisms. The restriction to direct mechanisms is justified by the revelation principle.[2]

Before describing our results and their implications, let us highlight the common features of the standard model we are studying and of the above-mentioned market design problems. There is a finite group of agents, each of whom would like to consume a single indivisible object to which we sometimes refer to as a "house," using the terminology coined by Shapley and Scarf (1974). We allow objects that are the agents' common endowment as well as objects that are privately owned. Agents have strict preferences over the objects, and are indifferent about what objects are allocated to other agents. The outcome of the problem is a matching of agents and objects.

We study two incentive-compatibility concepts, and two efficiency criteria. A mechanism is individually strategy-proof if no agent can benefit by reporting a non-truthful preference ranking. A mechanism is group strategy-proof if no group of agents can jointly manipulate

[1] Incentives and efficiency are also central to the theory of allocation mechanisms. For instance, Bogolomania and Moulin (2004) discuss "a recent flurry of papers on the deterministic assignment of indivisible goods" and state that "the central question of that literature is to characterize the set of efficient and incentive compatible (strategy-proof) assignment mechanisms." The 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 characterization of strategy-proofness and efficiency does not rely on additional assumptions.

[2] See Appendix A for a discussion of the revelation principle in our context.

their reports so that all of them weakly benefit from this manipulation, while at least one in the group strictly benefits. Importantly, strategy-proof mechanisms are immune to manipulation regardless of the information the agents' possess. As importantly, in our setting group strategy-proofness is equivalent to the lack of manipulation opportunities for groups of two agents. This makes group strategy-proofness a desirable property of mechanisms in our setting, particularly because in applications we see attempts at strategic coordination. Coordinated reporting to a mechanism has been, for instance, documented in kidney allocation and exchange (cf. Sönmez and Ünver, 2010; Ashlagi and Roth, 2011) and in school choice (Pathak and Sönmez, 2008).[3] Furthermore, coordinated reporting is effectively the only way a group of agents can manipulate allocations in many of these environments; for instance, without an approval from the school district, two parents cannot trade school admission decisions ex post.[4]

The first efficiency criterion we study is Pareto efficiency. A mechanism is Pareto efficient if, for all preference profiles, the resulting matching is not Pareto dominated by any other matching; a matching Pareto dominates another if all agents weakly prefer the former to the latter, and some agent's preference is strict. The second efficiency criterion we study requires the efficient matching to be the maximum of a social ranking of matchings, in line with Bergson (1938), Samuelson (1947), and Arrow's (1963) reformulation of welfare economics.[5] To formulate this more demanding efficiency criterion, we define a social welfare function (SWF) to be a mapping from profiles of agents' preferences over matchings to partial strict orderings of matchings. We allow partial orderings—such as Pareto dominance—and derive results for complete orderings as corollaries.[6] We require that each SWF satisfies the Pareto

[3]In kidney exchange, transplant centers occasionally try to first conduct kidney exchanges using their internal patient-donor pool, and list their patients and donors in outside exchange programs only if they fail to find a suitable match, thus hindering the efficiency of regional exchange systems (cf. Sönmez and Ünver, 2010; Ashlagi and Roth, 2011). Also, a doctor acting on behalf of several patients can coordinate their reports if it benefits his or her patients. There are known cases of doctors gaming medical systems for the benefit of their patients. For instance, in 2003 two Chicago hospitals settled a Federal lawsuit alleging that some patients had been fraudulently certified as sicker than they were to move them up on the liver transplant queue (Warmbir, 2003). In school choice, Pathak and Sönmez (2008) describe strategic cooperation among parents, e.g. among the members of the West Side Parent Group in Boston.

[4]Non-manipulability is not the only benefit of using strategy-proof mechanisms. Such mechanisms also 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 (cf. Vickrey, 1961; Dasgupta, Hammond, and Maskin, 1979; Pathak and Sönmez, 2008).

[5]Pareto efficiency is, on one hand, the baseline efficiency requirement, and on the other hand, it does not indicate which of the possibly many Pareto-efficient matchings to choose. For instance, Arrow (1963), pp. 36-37, discusses the partial ordering of outcomes given by Pareto dominance, and observes: "But though the study of maximal alternatives is possibly a useful preliminary to the analysis of particular social welfare functions, it is hard to see how any policy recommendations can be based merely on a knowledge of maximal alternatives. There is no way of deciding which maximal alternative to decide on."

[6]See e.g. Sen (1970,1999) for analysis of welfare with partial orderings.

and independence-of-irrelevant-alternatives postulates (Arrow, 1963): (i) a SWF is Pareto if it ranks any matching strictly below any other matching that Pareto dominates it, and (ii) a SWF satisfies the independence of irrelevant alternatives if, given any two profiles of preferences and any two matchings that are socially comparable under both profiles, if all agents rank the two matchings in the same way under both profiles, then the social ranking of the two matchings is the same under both profiles. We call a mechanism Arrovian efficient with respect to a SWF if, for all preference profiles, the resulting matching is the unique maximum of the SWF.[7] For shortness we say that a mechanism is Arrovian efficient if it is Arrovian efficient with respect to some SWF.

Our first main result (Theorem 1) establishes that a mechanism is individually strategy-proof and Arrovian efficient if and only if the mechanism is group strategy-proof and Pareto efficient. Both directions of the equivalence are noteworthy. First, the equivalence tells us that requiring individual strategy-proofness and Arrovian efficiency guarantees group strategy-proofness. As discussed above, group strategy-proofness is a very desirable property of an allocation mechanism. Second, mechanisms such as Serial Dictatorships or Top Trading Cycles (defined below) were known to be group strategy-proof and Pareto efficient; our equivalence allows us to conclude that they are also Arrovian efficient—there are Arrovian SWF that rationalize them.

As far as we know, the present paper is the first to connect the literature on allocation and exchange and the literature on Arrovian preference aggregation. In particular, we seem to be the first to recognize the equivalence of Theorem 1. However, stronger equivalence results—which do not hold true in our setting—are familiar from studies of voting. In voting—unlike in our problem—all agents have strict preferences among all outcomes. In the class of Pareto efficient mechanisms, individual strategy-proofness is then equivalent to group strategy-proofness (Gibbard, 1973, and Satterthwaite, 1975).[8] This stronger equivalence fails in our setting as it admits individually strategy-proof and Pareto-efficient mechanisms that fail group strategy-proofness.

[7] There is a rich social choice literature on the correspondence between choice and the maximum of the SWF ranking in the context of social choice (see below). This literature is interested in rationalizing social choice rather than efficiency of allocation mechanisms, and hence it says that a mechanism, or social choice, is "rationalized by a SWF" rather than "efficient with respect to a SWF."

[8] The equivalence of Theorem 1 also has counterparts in the social choice literature on restricted preference domains—such as single-peaked preferences—in which there are non-dictatorial strategy-proof and Arrow efficient rules. For instance, Moulin (1988) extends a result by Blair and Muller (1983) and shows that in environments such as single-peaked voting, if an Arrovian SWF is monotonic, then the mechanism picking its unique maximal element is group strategy-proof. In particular, this implies that in single-peaked voting individual strategy-proofness and group strategy-proofness are equivalent with no need to restrict attention to efficient mechanisms. In contrast, in allocation environments the equivalence results from the conjunction of incentive and efficiency assumptions, and the equivalence of incentive assumptions alone is not true.

The equivalence of Theorem 1 leads to a question: what mechanisms are individually strategy-proof and Arrovian efficient?

Our second main result, Theorem 2, answers this question and constructs the full class of individually strategy-proof and Arrovian efficient mechanisms, or, equivalently, the full class of group strategy-proof and Pareto efficient mechanisms.

This new class of mechanisms—which we call *trading-cycles* mechanisms, or trading cycles for shortness—is closely related to David Gale's top-trading-cycle mechanism (reported by Shapley and Scarf, 1974), and especially its generalization by Pápai (2000) (known as hierarchical exchange, or simply top trading cycles). Let us describe trading cycles in the special case of our environment in which there are as many objects as agents and each agent initially controls an object. First consider Gale's top trading cycles. The top-trading-cycle algorithm resembles decentralized trading and 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. Since there are a finite 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. The pointing cycles might be short (agent 1 points to object A, which points back to agent 1) or might involve many agents. 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.

Gale's top-trading-cycle mechanism is a special case of trading cycles; Roth (1982) showed that it is group strategy-proof and Pareto efficient. Other examples of trading cycles obtain when we take one of the agents—let us 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 except that we make the broker point to his most preferred unmatched object that is different from the brokered object. Surprisingly, we prove that this modification of top trading cycles remains group strategy-proof and Pareto efficient (and hence, by Theorem 1, also Arrovian efficient).[9] Even more surprisingly, this slight modification of top trading cycles gives us the full class of group strategy-proof and efficient mechanisms.

While we described top trading cycles and trading cycles for a particular environment, the same algorithms can be used in more general environments, for instance when all objects

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

are socially endowed. In such environments, to run top trading cycles we need to specify for each round and each object which agent controls it (see Abdulkadiroğlu and Sönmez, 1999; Pápai, 2000); to run trading cycles we additionally need to specify for every round who, if anyone, is the broker. Provided we are careful how the control rights change from round to round, the resulting mechanisms are group strategy-proof and efficient, and no other mechanisms are.[10]

The main insight brought by our characterization is that every individually strategy-proof and Arrovian efficient mechanism can be obtained by specifying agents' control rights, and allowing them to swap objects. In this sense, our result can be seen as a variant of the Second Fundamental Welfare theorem for the setting without transfers and with single-unit demands.

Knowing the full class of individually strategy-proof and Arrovian efficient mechanisms, allows us to derive some further properties shared by all such mechanisms. In particular, we show that in any such mechanism, for any preference profile, there is a group of agents—the decisive group—all of whom can get one of their two top choices, and all but at most one of whom can get their top choice, irrespective of preferences submitted by agents not in the group. In the trading-cycle algorithm, the decisive group consists of agents who trade in the first round.[11] We further show that all strategy-proof and efficient mechanisms have a recursive structure: the members of the decisive group determine their allocation; given their preferences there is another group of agents who obtain one of top two choices among remaining objects, and who can determine their allocation irrespective of the preferences of others, etc. For instance, in a sequential dictatorship (Satterthwaite and Sonnenschein, 1981; Svensson, 1994, 1999; Ergin, 2000), which is a special case of trading cycles, the first dictator chooses his most preferred object, then a second dictator chooses his most preferred object among the objects which were not chosen by prior dictators, and so forth (we refer to the mechanism as serial dictatorship if the sequence of dictators is exogenously given).

Furthermore, knowing that all individually strategy-proof and Arrovian efficient mechanisms may be represented as trading cycles allows one to determine what can and cannot

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

[11]A similar point was made by Sen (1970) and Gibbard (1969) in the context of voting: every SWF whose ranking of outcomes is a quasi-ordering is determined by the preferences of a group of agents they call oligarchs. Notice that in the context of allocation, the result is more subtle in that who belongs to the decisive group can depend on the profile of preferences, and we might have one member of the decisive group whose preference ranking co-determines the allocation and who obtains one of his top two choices but not necessarily his top choice.

be achieved in a strategy-proof way. The characterization radically simplifies analysis of such questions because it allows us to restrict our attention to trading cycles without loss of generality (we provide examples of such a radical simplification below). In this sense, the role trading-cycles mechanisms play in the single-unit demand no-transfers 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 (cf. Green and Laffont, 1977, and Holmstrom, 1979). Other characterizations of efficient and strategy-proof mechanisms that are non-dictatorial have been obtained by Barberà, Jackson, and Neme (1997) for sharing a perfectly divisible good among agents with single-peaked preferences over their shares; and by Barberà, Gül, and Stacchetti (1993) for voting problems with single-peaked preferences.

To illustrate how Theorems 1 and 2 simplify the analysis of many otherwise difficult questions, we use them to obtain new insights into allocation and exchange, as well as to show that some of the deepest prior results on allocation in environments with single-unit demands and no transfers are their immediate corollaries.

First, we apply our results to the problem of exchange of goods without transfers and with single-unit demands. For example, in kidney exchange, patients (agents) come with a paired-donor kidneys (objects) and have to be matched with at least their paired-donor kidney. Another example is the allocation of dormitory rooms at universities that give some students, such as sophomores, the right to stay in the room they lived in the preceding year. Such exogenous control rights are straightforwardly accommodated by our mechanism class. When some objects are private endowments of agents it is natural to require that the participation in the mechanism is individually rational so that each agent likes the mechanism's outcome at least as much as the best object from his endowment. We show that the class of individually strategy-proof, Arrovian efficient, and individually rational mechanisms equals the class of individually rational trading-cycles mechanisms. A trading-cycles mechanism is individually rational if and only if (i) it may be represented by a consistent 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, top-trading-cycles mechanisms are the unique mechanisms that are individually strategy-proof, Arrovian efficient, and individually rational. In 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 Theorems 1 and 2 is implied by an earlier result of Ma (1994).[12]

[12]Ma shows that Top Trading Cycles is the unique strategy-proof, Pareto-efficient, and individually-rational mechanism in the discrete exchange economy with single-unit demand and single-unit endowment introduced by Shapley and Scarf (1974). There exists a unique core allocation in such an economy that can be reached by Gale's TTC algorithm (cf. Shapley and Scarf, 1974 and Roth and Postlewaite, 1977). Konishi, Quint,

Second, we show that sequential dictatorships, defined above, are the only mechanisms that are individually strategy-proof and Arrovian efficient with respect to a SWF that always generates complete orderings.[13] Dictatorships are the benchmark strategy-proof and efficient mechanisms in many areas of economics. For instance, Gibbard (1973) and Satterthwaite (1975) have shown that all strategy-proof and unanimous voting rules are dictatorial.[14] Still, we find it surprising that this corollary of Theorems 1 and 2 holds true in our environment because—in contrast to the environments where this question was previously studied—ours allows many individually strategy-proof (and even group strategy-proof) and Pareto efficient mechanisms that are not dictatorial.

Third, we show that the some of the deep prior insights of the rich literature on allocation with no-transfers are immediately implied by Theorem 2. Pápai (2000)—the prior work closest to our paper—constructed a class of mechanisms referred to as top trading cycles or hierarchical exchange, which use the same algorithm as Gale's top-trading-cycles mechanism with the exception that the mechanism takes as an input a structure of control rights (without brokers) over objects that—for each round of the mechanism and each unmatched object—determines the agent to whom the object points. She then showed that all group strategy-proof and Pareto efficient mechanisms that satisfy an additional technical property (that she refers to as reallocation-proofness) are in her class.[15] This result is implied by Theorem 2 because trading cycles with brokers do not satisfy Papai's reallocation-proofness property.[16]

and Wako (2001) considered an extension when agents have multi-unit demands and endowments. They showed that a core allocation may not exist when agents have additive preferences over multiple objects. Pápai (2007) showed that when we can rule out some of the types of trades that agents are allowed to make in this multi-unit model, then an extension of Ma's characterization can be restored.

[13]We allow outside options in this result; without outside options we show that this subclass of trading cycles is slightly larger than the class of sequential dictatorships.

[14]Dasgupta, Hammond, and Maskin (1979) extended this result to more general social choice models, Satterthwaite and Sonnenschein (1981) extended it to public goods economies with production, Zhou (1991) extended it to pure public goods economies, and Hatfield (2009) to group strategy-proof quota allocations. In exchange economies, Barberà and Jackson (1995) showed that strategy-proof mechanisms are Pareto inefficient.

[15]A 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, both of them weakly gain and one of them strictly gains by swapping their assignments, and (ii) if only one of them misrepresents his preferences, he cannot change his assignment. Pápai also notes that the more natural reallocation-proofness-type property obtained by dropping condition (ii) conflicts with group strategy-proofness and Pareto efficiency as does allowing the swap of assignments among more than two agents. We do not use reallocation-proofness in our results.

[16]All allocation papers cited above, and the literature in general, shares with our paper the assumption that agents have strict preferences. This is the standard modeling assumption because—as Ehlers (2002) shows—"one cannot go much beyond strict preferences if one insists on efficiency and group strategy-proofness." The full preference domain gives rise to an impossibility result, i.e., when agents can be indifferent among objects, there exists no mechanism that is group strategy-proof and Pareto efficient. For this reason, participants are frequently allowed to submit only strict preference orderings to real-life direct mechanisms in various markets, such as dormitory room allocation, school choice, and matching of interns and hospitals.

Svensson (1999) 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. This result follows from Theorem 2 because neutral and group strategy-proof mechanisms are Pareto efficient, and because, to be neutral, a trading-cycle mechanism must be a serial dictatorship.

Finally, one of the auxiliary contributions of our paper is to recognize the role of brokers in allocation and exchange problems with no transfers. In the context of our paper, the main role played by the brokers is to allow us to construct the full class of strategy-proof and efficient mechanisms. The brokers can also be useful in some mechanism design settings, and we close the paper by providing an example of such a setting. In the example, the trading cycle with one broker described above is the most equitable allocation mechanism. In particular, we prove that it is strictly more equitable—in the sense of Lorenz dominance—than any top-trading-cycles mechanism.

2 Model

2.1 Environment

Let I be a set of **agents** and H be a set of **objects** that we often refer to as **houses** following the standard terminology of the literature. We use letters i, j, k to refer to agents and h, g, e to refer to houses. Each agent i has a **strict preference relation** over H, denoted by $\succ_i$.[17] Let $\mathbf{P}_i$ be the set of strict preference relations for agent i, and let $\mathbf{P}_J$ denote the Cartesian product $\times_{i \in J} \mathbf{P}_i$ for any $J \subseteq I$. Any profile from $\succ = (\succ_i)_{i \in I}$ from $\mathbf{P} \equiv \mathbf{P}_I$ is called a **preference profile**. For all $\succ \in \mathbf{P}$ and all $J \subseteq I$, let $\succ_J = (\succ_i)_{i \in J} \in P_J$ be the restriction of $\succ$ to J.

To simplify the exposition, we make two initial assumptions. Both of these assumptions are fully relaxed in subsequent sections. First, we initially restrict attention to house allocation problems. A **house allocation problem** is the triple $\langle I, H, \succ \rangle$ (cf. Hylland and Zeckhauser, 1979). Throughout the paper, we fix I and H, and thus, a problem is identified with its preference profile. In Section 6.1, we generalize the setting and the results to house allocation and exchange by allowing agents to have initial rights over houses. The results on allocation and exchange turn out to be straightforward corollaries of the results on (pure) allocation. Second, we initially follow the tradition adopted by many papers in the literature (cf. Svensson, 1999) and assume that $|H| \geq |I|$ so that each agent is allocated a house. This

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

assumption is satisfied in settings in which each agent is always allocated a house (there are no outside options), as well as in settings in which agents' outside options are tradable, effectively being indistinguishable from houses. In Section 5.2, we allow for non-tradable outside options and show that analogues of our results remain true irrespective of whether $|H| \geq |I|$ or $|H| < |I|$.

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

A **mechanism** is a mapping $\varphi : \mathbf{P} \longrightarrow \mathcal{M}$ that assigns a matching for each preference profile (or, equivalently, for each allocation problem).[18]

2.2 Strategy-Proofness and Efficiency

A mechanism is individually strategy-proof if truthful revelation of preferences is a weakly dominant strategy for any agent: a mechanism φ is **individually strategy-proof** if for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ_i' \in \mathbf{P}_i$ such that

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

A mechanism is group strategy-proof if there is no group of agents that can misstate their preferences in a way such that each one in the group gets a weakly better house, and at least one agent in the group gets a strictly better house, irrespective of the preference ranking of the agents not in the group. Formally, a mechanism φ is **group strategy-proof** if for all

[18] We study direct mechanisms. By the revelation principle, this is without loss of generality. See Appendix A for a discussion.

$\succ \in \mathbf{P}$, there exists no $J \subseteq I$ and $\succ'_J \in \mathbf{P}_J$ such that

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

and

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

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

Pareto efficiency is a weak efficiency requirement.[19] In order to define the stronger concept of Arrovian efficiency with respect to a social welfare function, denote by P^M the set of strict partial orderings over matchings; we refer to elements of P^M as social rankings. A **social welfare function** (**SWF**) $\Phi : \mathbf{P} \to P^M$ maps agents' preference profiles to social rankings. A SWF Φ is **Pareto** (or unanimous) if: for every preference profile $\succ$ and any two matchings $\mu, \nu \in \mathcal{M}$, if $\mu(i) \succeq_i \nu(i)$ for all $i \in I$, with at least one preference strict, then μ is ranked above ν in the social ranking, $\mu\Phi(\succ)\,\nu$. A SWF Φ satisfies the **independence of irrelevant alternatives** if: for all $\succ, \succ' \in \mathbf{P}$ and all $\mu, \nu \in \mathcal{M}$, if all agents rank μ and ν in the same way and both $\Phi(\succ)$ and $\Phi(\succ')$ rank μ and ν then $\mu\Phi(\succ')\nu \iff \mu\Phi(\succ)\nu$. We restrict attention to SWFs that satisfy the Pareto and independence-of-irrelevant-alternatives postulates. Notice that Pareto dominance is a standard example of a SWF.

A matching μ is **Arrovian efficient** with respect to a social ranking $\Phi(\succ)$ if it maximizes the social welfare, that is $\mu\Phi(\succ)\nu$ for all $\nu \in \mathcal{M}\backslash\{\mu\}$. A mechanism ϕ is **Arrovian efficient** with respect to a SWF Φ if for any profile of agents' preferences $\succ$, the matching $\phi(\succ)$ is Arrovian efficient with respect to $\Phi(\succ)$. If ϕ is Arrovian efficient with respect to some SWF, we simply say that it is Arrovian efficient. The next section offers two examples illustrating the concept of Arrovian efficiency.

3 Main Results: Equivalence

In Theorem 1, we establish the equivalence between three of the pairs of our incentive-compatibility and efficiency concepts. In addition, Example 2 below demonstrates that the

[19] In particular, when imposed on group strategy-proof mechanisms, Pareto efficiency is equivalent to assuming that the mechanism maps $\mathbf{P}$ onto the entire set of matchings $\mathcal{M}$. This surjectivity property is known as citizen sovereignty, or full range.

class of individually strategy-proof and Pareto efficient mechanisms is a strict superset of the mechanisms satisfying any of the equivalent conditions of the theorem.

Theorem 1. *A mechanism is individually strategy-proof and Arrovian efficient if and only if it is group strategy-proof and Pareto efficient, and if and only if it is group strategy-proof and Arrovian efficient.*

To illustrate this equivalence and our concepts let us look a the setting with three agents 1, 2, and 3, three objects (houses) h_1, h_2, and h_3, and no outside options. Consider the following two examples of mechanisms.

Example 1. The serial dictatorship in which 1 chooses first, and 2 chooses second is well-known to be group strategy-proof and Pareto efficient. It is straightforward to see that this serial dictatorship is Arrovian efficient with respect to the following SWF: μ is ranked strictly above ν if and only if (a) 1 strictly prefers μ to ν, or (b) 1 is indifferent and 2 strictly prefers μ to ν.

Example 2. Let us now modify the serial dictatorship of the previous example and consider mechanisms ψ in which 1 chooses first; then 2 chooses second if 1 prefers h_2 over h_3, else 3 chooses second. This mechanism is an example of a sequential dictatorship, and is also individually strategy-proof and Pareto efficient. However, mechanism ψ is neither Arrovian efficient nor group strategy-proof. To see the latter point let us look at the following two preference profiles:

$$1: h_1 \succ h_2 \succ h_3, \;\; 2: h_1 \succ h_2 \succ h_3, \;\; 3: h_1 \succ h_2 \succ h_3,$$

$$1: h_1 \succ' h_3 \succ' h_2, \;\; 2: h_1 \succ' h_2 \succ' h_3, \;\; 3: h_1 \succ' h_2 \succ' h_3.$$

Notice that

$$\begin{aligned} \psi(\succ) &= \{(1, h_1), (2, h_2), (3, h_3)\}, \\ \psi(\succ') &= \{(1, h_1), (2, h_3), (3, h_2)\}. \end{aligned}$$

The mechanism ψ fails group strategy-proofness. For instance, when the true preference profile is $\succ$, then agents 1 and 3 have a profitable manipulation $\{\succ'\}_{\{1,3\}}$. The mechanism ψ also fails Arrovian efficiency. Indeed, by way of contradiction assume that ψ is Arrovian efficient with respect to some SWF Ψ. Then, $\Psi(\succ)$ ranks allocation $\psi(\succ)$ strictly above $\psi(\succ')$, and $\Psi(\succ')$ ranks $\psi(\succ')$ strictly above $\psi(\succ)$. But, this violates the independence of the irrelevant alternatives, a contradiction that shows that ψ is not Arrovian efficient.

The proof of Theorem 1 builds on Example 2. As a preparation for the proof, let us notice three properties of group strategy-proofness. First, in the environment we study group strategy-proofness is equivalent to the conjunction of two non-cooperative properties: individual strategy-proofness and non-bossiness.[20] Non-bossiness (Satterthwaite and Sonnenschein, 1981) means that no agent can misreport his preferences in such a way that his allocation is not changed but the allocation of some other agent is changed: a mechanism φ is **non-bossy** if for all $\succ \in \mathbf{P}$, there is no $i \in I$ and $\succ_i' \in \mathbf{P}_i$ such that

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

The following lemma 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 non-bossy.*

Second, in the environment we study group strategy-proofness is equivalent to Maskin monotonicity (Maskin, 1999). A mechanism φ is **Maskin monotonic** if $\varphi[\succ'] = \varphi[\succ]$ whenever $\succ' \in \mathbf{P}$ is a φ-monotonic transformation of $\succ \in \mathbf{P}$. A preference profile $\succ' \in \mathbf{P}$ is a **φ-monotonic transformation** of $\succ \in \mathbf{P}$ if

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

Thus, for each agent, the set of houses better than the base-profile allocation weakly shrinks when we go from the base profile to its monotonic transformation. The following lemma was proven by Takamiya (2001) for a subset of the problems we study; his proof can be extended to our more general setting.

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

Finally, let us notice the following

Lemma 3. *If a mechanism ϕ is group strategy-proof then no agent can change the outcome of ϕ by changing the ranking of objects worse than the object he obtains, that is if $\succ'$ differs from $\succ$ only in how some agent i ranks objects below $\phi(\succ)(i)$ then $\phi(\succ') = \phi(\succ')$.*

We skip the straightforward proof of this last lemma since we later prove, without reliance on this lemma or Theorem 1, a substantially stronger result, Theorem 2.

[20]Both of these properties are non-cooperative in the sense that they relate a mechanism's outcomes under two scenarios when a single agent makes unilateral preference revelation deviations.

Proof of Theorem 1. Notice that it is sufficient to show that individual strategy-proofness and Arrovian efficiency are equivalent to group strategy-proofness and Pareto efficiency as the third equivalence then follows.

First, consider an individually strategy-proof mechanism ϕ that is Arrovian efficient with respect to some SWF Φ. In light of Lemma 1, to establish the first implication it is enough to show that ϕ is Pareto efficient and non-bossy.

To show that ϕ is Pareto efficient, suppose that for some $\succ \in \mathbf{P}$, $\phi[\succ]$ is not Pareto efficient. Then, there exists some $\mu \in \mathcal{M} \setminus \{\phi[\succ]\}$ such that $\mu(i) \succeq_i \phi[\succ](i)$ for all i, with a strict preference for at least one agent. Since Φ satisfies the Pareto postulate, $\mu \, \Phi(\succ) \, \phi[\succ]$, which contradicts the assumption that ϕ is Arrovian efficient with respect to Φ.

To show that ϕ is non-bossy, let $\succ \in \mathbf{P}$ and $\succ_i' \in \mathbf{P}_i$ be such that

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

Denote $\succ' = (\succ_i', \succ_{-i})$. Since ϕ is Arrovian efficient with respect to Φ, the matching $\phi[\succ]$ is ranked as the unique first by $\Phi(\succ)$ and the matching $\phi[\succ']$ is ranked as the unique first by $\Phi(\succ')$. Thus, $\phi[\succ]$ and $\phi[\succ']$ are comparable under both $\Phi(\succ)$ and $\Phi(\succ')$, and independence of irrelevant alternatives implies that $\phi[\succ]$ and $\phi[\succ']$ are ranked in the same way by $\Phi(\succ)$ and $\Phi(\succ')$. We can thus conclude that $\phi[\succ] = \phi[\succ']$. This establishes that ϕ is non-bossy.

Second, consider a group strategy-proof and Pareto efficient mechanism ϕ. We define the SWF Φ as follows: for any profile of preferences $\succ$ and any matchings μ and $\mu' \neq \mu$, matching μ is ranked by $\Phi(\succ)$ above μ' iff either (i) we have $\mu = \phi(\succ)$ or (ii) for all agents i, we have $\mu(i) \succsim_i \mu'(i)$. Note that Pareto efficiency of ϕ implies that conditions (i) and (ii) are consistent with each other, and hence, that the SWF Φ is well-defined.

By definition, Φ satisfies the Pareto postulate. Furthermore, Φ is transitive: if $\Phi(\succ)$ ranks μ^1 above μ^2 and it ranks μ^2 above μ^3 then it ranks μ^1 above μ^3. Indeed, if one of the μ^i (for $i = 1, 2, 3$) equals $\phi(\succ)$, then it must be that $\mu^1 = \phi(\succ)$ and the claim is proved. If none of the μ^i equals $\phi(\succ)$, then agents unanimously rank μ^1 above μ^2 and unanimously rank μ^2 above μ^3; we can conclude that the agents unanimously rank μ^1 above μ^3 and thus $\Phi(\succ)$ ranks μ^1 above μ^3.

It remains to check that Φ satisfies the independence of irrelevant alternatives. Take two preference profiles $\succ^1$ and $\succ^2$ such that each agent ranks two matchings, say μ and μ', in the same way under the two preference profiles. If the two matchings are comparable under both $\Phi(\succ^1)$ and $\Phi(\succ^2)$, then one of the following cases obtains:

Case 1: one of the matchings is unanimously preferred to the other under $\succ^1$; then the same unanimous preference obtains under $\succ^2$ and the claim is true.

Case 2: there is no unanimous ranking of the two matchings under $\succ^1$; then unanimity cannot obtain under $\succ^2$ either. Since, the matchings are ranked, it must be that $\phi(\succ^1)$ and $\phi(\succ^2)$ take value in $\{\mu, \mu'\}$. Say, $\phi(\succ^1) = \mu$; then we need to check that $\phi(\succ^2) = \mu$ as well. By Lemma 2, we can assume that each agent i ranks $\mu(i)$ and $\mu'(i)$ at the top of his ranking under both $\succ^1$ and $\succ^2$. Furthermore, by Lemma 3 only rankings of objects above agents' allocations (and including their allocations) affect the outcome of a group strategy-proof mechanism; we can thus conclude that $\phi(\succ^1) = \phi(\succ^2)$. QED

4 The Construction of Trading-Cycles Mechanisms

We have established that individual strategy-proofness and Arrovian efficiency are equivalent to group strategy-proofness and Pareto efficiency. We now determine which mechanisms satisfy these properties. Starting with some examples, we construct the full class of such mechanisms; we call them trading cycles (TC).

4.1 Example: Top Trading Cycles

To set the stage for our trading-cycles (TC) mechanism, let us look at the well-known top-trading-cycles (TTC) algorithm adapted by Pápai (2000) to house allocation problems.[21] The class of mechanisms presented in this subsection is identical to Pápai's "hierarchical exchange" class. Our presentation, however, is novel and aims to simultaneously simplify the earlier constructions of Pápai's class, and to introduce some of the terminology we will later use to introduce our class of all group strategy-proof and efficient mechanisms (TC).

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

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

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

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

As we see, the outcome of the TTC algorithm is determined by two types of inputs: agents' preferences and agents' rights of ownership over houses. The preferences are, of course, submitted by the agents. The ownership rights are defined exogenously as part of the mechanism.[22] We formalize this aspect of the mechanism via the following concept.

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

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

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

Each consistent structure of ownership rights $\{c_\sigma\}_{\sigma\in\overline{\mathcal{M}}}$ determines a *hierarchical exchange mechanism* of Pápai (2000). This class consists of mechanisms whose outcomes are found by running the TTC algorithm with consistent structures of ownership rights. Because of this, we will also refer to hierarchical exchange as **TTC mechanisms**. Pápai showed that all TTC mechanisms are group strategy-proof and Pareto efficient, extending an earlier insight of Roth (1982).[23]

4.2 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 3. Consider three agents $i_1, ..., i_3$ and three houses $h_1, ..., h_3$ and an ownership

[22] Recall that we are studying an allocation problem in which objects are a collective endowment. In Section 6.1 we will enlarge the analysis to include exchange problems among agents with private endowments. In exchange problems, some of the mechanism's ownership rights are determined by individual rationality constraints. Notice that ownership rights are related to priorities in school choice in that agents with higher priority at an object correspond to agents who own the object at smaller submatchings, while agents with lower priority at this object correspond to agents who own the object at larger submatchings at which previous owners are already matched (see, for instance, Ergin (2002), Abdulkadiroğlu and Sönmez (2003), and Abdulkadiroğlu and Che (2010) for a discussion of TTC and priorities).

[23] To 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.

structure that allocates ownership of houses according to the following table:

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

Given this structure, let us run TTC with one modification: agent i_1 is not allowed to point to house h_1 as long as there are other unmatched agents. In rounds with other unmatched agents (and hence other unmatched houses), agent i_1 will point to his most preferred house among unmatched houses other than h_1.[24]

For instance, if each agent i has the preference $h_1 \succ_i h_2 \succ_i h_3$ then in the first round agents i_2 and i_3 will point to h_1, but agent i_1 will point to his second-choice house, h_2. We will then have an exchange cycle in which i_1 is matched with h_2 and i_2 is matched with h_1. In the second round, the algorithm matches agent i_3 and house h_3, and terminates.

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

To establish group strategy-proofness we may use an argument similar to the standard proof why TTC are group strategy-proof, see Roth (1982). Instead of replicating this argument, let us consider the following concept. For every agent and every round of the algorithm, let us say that the set of objects *obtainable* by the agent consists of all the objects the agent could obtain in this round by either submitting his true preference ranking or by changing the ranking of not-yet matched objects. One can check that, as long as an agent is unmatched, the agent's set of obtainable objects stays the same or becomes larger with each round. This monotonicity is what drives the strategy-proofness properties of the mechanism.

The above mechanism is indeed different from all TTC mechanisms. To see this, first observe that the mechanism matches house h_1 with agent i_2 under the illustrative preference profile analyzed above, whereas it would match h_1 with another agent, i_3, if agent i_1 submitted preferences $h_1 \succ_{i_1} h_3 \succ_{i_1} h_2$ (and other agents $i \neq i_1$ continued to have preferences

[24] Pápai (2000) gives an example of a non-TTC mechanism. Her construction is different from ours though the resultant mechanisms are identical. As we will show in the next section, the advantage of our construction lies in its generalizability to cover the whole class of group strategy-proof and efficient mechanisms.

$h_1 \succ_i h_2 \succ_i h_3$). However, any TTC mechanism would match h_1 with the same agent in these two preference profiles. Indeed, any TTC ownership structure uniquely determines which agent owns h_1 at the empty submatching, and this agent would be matched with h_1 in the first round of the algorithm under any preference profile in which all agents rank h_1 as their first choice.

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

4.3 Trading Cycles

We turn now to our new algorithm, trading cycles (TC), an example of which we saw in the previous section. Like TTC, the TC is a recursive algorithm that matches agents and houses in exchange cycles over a sequence of rounds. TC is more flexible, however, as it allows two types of intra-round control rights over houses that agents bring to the exchange cycles: ownership and brokerage.

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

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

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

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

The consistency requirement on TC control rights structures consists of three constraints on brokerage at any given submatching (the *within-round* requirements) and three constraints on how the control rights are related across different submatchings (the *across-rounds* requirements, which we will introduce after our new algorithm).[25]

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

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

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

(**R3**) If agent i brokers a house at σ, then i does not own any houses at σ.

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

Requirements R1-R2 are what we need for the TC algorithm to be well defined; R3 is necessary for Pareto efficiency and individual strategy-proofness (see Appendix B). With these requirements in place, we are ready to describe the TC algorithm.

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

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

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

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

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

[25] Properties R1-R6 are defined over submatchings instead of rounds of the TC algorithm. This simplifies the definition of the consistency conditions as it saves us the trouble of keeping track of which submatchings are relevant in the TC algorithm, i.e., which of them can be encountered at the end of a round of a TC algorithm.

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

The algorithm terminates when all agents or all houses are matched.

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, the matching in Step 3 is well defined. Finally, since we match at least one agent-house pair in every round, and since there are finitely many agents and houses, the algorithm terminates after finitely many rounds.

Our algorithm builds on Gale's top-trading-cycles idea, described in Section 4.1, but allows more general trading cycles than top cycles. In TC, brokers do not necessarily point to their top-choice houses.[26] The terminology of owners and brokers is motivated by an imperfect analogy: At any submatching (but not globally through the algorithm), we can think of the broker of house h as representing a latent agent who owns h but prefers any other house over it.

The TC algorithm with a control rights structure satisfying R1-R3 determines a mechanism that maps profiles from $\mathbf{P}$ to Pareto efficient matchings in $\mathcal{M}$.[27] To guarantee that the resulting mechanism is group strategy-proof (and hence also Arrovian efficient) we need to impose the following across-round consistency requirements on the control rights structure.

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

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

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

[26] Looking back at the example of the previous section, we see that the mechanism constructed there was a TC in which agent i_1 brokered house h_1 while other agents owned houses. The difference between TTC and TC is encapsulated in Step 1; the other steps are standard. In contrast to TC, all previous developments of Gale's idea (Shapley and Scarf, 1974)—e.g. the top trading cycles with newcomers (Abdulkadiroğlu and Sönmez, 1999), hierarchical exchange (Pápai, 2000), top trading cycles for school choice (Abdulkadiroğlu and Sönmez, 2003), and top trading cycles and chains (Roth, Sönmez, and Ünver, 2004)—allowed only top trading cycles and had all agents point to their top choice among unmatched houses. All these previous developments may be viewed as using a subclass of TC in which all control rights are ownership rights and there are no brokers (in Section 6.1 we show that TC can easily handle private endowments).

[27] The recursive argument for the efficiency of the non-TTC mechanism from Section 4.2 applies; see Appendix C for details.

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

A good way to understand conditions R4–R6 is to look at the concept of obtainable objects introduced in Section 4.2. The set of obtainable objects is empty for agents who control no objects in a particular round. The set is always non-empty for owners. For brokers, the set of obtainable objects can be either empty or non-empty. Conditions R4–R6 are formulated to ensure that for each agent the set of obtainable objects stays the same or becomes larger with each round.

Requirements R4 and R5 postulate that control rights persist: agents retain control rights as they move from smaller to larger submatchings, or through the rounds of the algorithm. Persistence of ownership is permanent by R4: if an agent becomes an owner, he remains an owner through the execution of the algorithm until he is matched (i.e., once an owner, always an owner). On the other hand, persistence of brokerage is limited in R5: an agent can potentially leave brokerage through the rounds of the algorithm under certain conditions. R4 (*persistence of ownership*) is identical to the consistency assumption we imposed on TTC. The first example of Section 4.1 illustrated why we need such a persistence assumption for the resultant mechanism to be individually strategy-proof. A similar example might convince us that individual strategy-proofness relies also on requirement R5 (*persistence of brokerage*); see Appendix B.

Requirement R6 has two parts.[28] The first part (*consolation for lost control rights)* postulates that when an agent i is matched with a house controlled by i', then i' owns the houses previously controlled by i.[29] R6's second part (*brokered-to-owned house transition*) postulates who obtains the control right over a house when a broker loses his brokerage right.

[28]Sections 6, 6.1 and 6.3 show that brokers and condition R6 are quite easy to work with. However, to sidestep the complication of condition R6, the reader is invited to keep in mind a smaller class of control rights structures in which both of these requirements are replaced by the following strong form of brokerage persistence: "If $|\sigma'| < |I| - 1$ and agent i brokers house h at σ then i brokers h at σ'." We think that by restricting attention to this smaller class of control rights structures, one is not missing much of the flexibility of the TC class of mechanism. We hasten to stress, however, that the complication is there for a reason: there are group strategy-proof and Pareto efficient mechanisms that cannot be replicated by TC control rights structures satisfying the above strengthening of R5-R6. Let us also stress that, a priori, we could expect the class of group strategy-proof and efficient mechanisms to be much more complex than it turned out to be.

[29]It is sufficient to restrict R6 to the case when i' is a broker of h' at σ. Nevertheless, we impose the first part of R6 on both brokers and owners because this gives a smaller class of control rights structures. A key step in seeing why the restriction of R6 to brokers is sufficient is to recognize that, if i' controls h' at σ, a round of the TC algorithm generates σ, and a later round generates a submatching that keeps i' unmatched but contains $\sigma \cup \{(i, h')\}$, then the control right of i' over h' at σ must be brokerage. We would like to thank the referees for drawing our attention to this point.

By R5, the broker can only lose the brokerage right between σ and σ' when no more than one agent is a σ-owner and σ'-owner; this is the agent who obtains the control right at σ' over the house brokered at σ. Requirement R6 is used to guarantee both non-bossiness and individual strategy-proofness of the mechanism; see Appendix B.

A key implication of R6 and R4 is the transfer of ownership rights to ex-brokers: if i' brokers h' at σ but not at σ', and $i \in \overline{I_{\sigma'}}$ owns $h \in \overline{H_{\sigma'}}$ at σ, then R6's first part implies that i' owns h at $\sigma \cup \{(i, h')\}$, and R4 further implies that i' owns h at $\sigma' \cup \{(i, h')\}$. We refer to this consequence of R6 and R4 as *broker-to-heir transition*.

We are now ready to define our mechanism class.

Definition 3. A control rights structure is **consistent** if it satisfies requirements R1-R6. The class of **TC mechanisms** (trading cycles) consists of mechanisms whose outcomes are determined by running the TC algorithm with consistent control rights structures.

The TTC mechanisms of Section 4.1 and the non-TTC mechanism of Section 4.2 are examples of TC. We will denote by $\psi^{c,b}$ the TC mechanism obtained from a consistent control rights structure $\{(c_\sigma, b_\sigma)\}_{\sigma \in \overline{\mathcal{M}}}$. In Section 5.2 we enlarge it to allow for agents' outside options (in particular, we then allow agents to ranks some objects as unacceptable), and in Section 7 we adapt this class of mechanisms to exchange problems.

5 Main Results: Characterization

Our main characterization result tells us that the class of Trading Cycles mechanisms coincides with the class of group strategy-proof and Pareto-efficient mechanisms, and hence with the class of individually strategy-proof and Arrovian efficient mechanisms (by Theorem 1). We first state and prove this result for the model of allocation in which all objects are acceptable, and then relax this simplifying assumption.

Theorem 2. *A mechanism is group strategy-proof and Pareto efficient if and only if it is a Trading Cycles mechanism.*

The argument 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 4.2; see Appendix C for details.

The proof that Trading Cycles are group strategy-proof builds on the following simple observation.

Lemma 4. *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 $[\succ'_i, \succ_{-i}]$.*

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

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

To further show that Trading-Cycles mechanisms are group strategy-proof, recall the lemma of Pápai (2000) (see Lemma 1 above) that implies that to prove group strategy-proofness of an individually strategy-proof mechanism it is enough to show that the mechanism is non-bossy. Proving non-bossiness turns out to be subtle; we provide this part of the proof in Appendix E. To get a sense for this part of the proof, consider a TC mechanism without brokers, and an agent i who gets the same object whether he submits preferences $\succ_i$ or $\succ'_i$. An inductive argument then shows that the algorithm will go though the same cycles under $\succ = (\succ_i, \succ_{-i})$ and $\succ' = (\succ'_i, \succ_{-i})$ even if the rounds at which these cycles are formed may differ. If brokers were strongly persistent, the same argument would apply.

The subtlety in proving non-bossiness is when a broker loses his brokerage right. Condition R5 ensures that cycles of three agents or more are the same under both $\succ$ and $\succ'$, but that cycles of one or two agents can be different. For instance, in the setting of Example 8, consider a preference profile in which agents i_1 and i_3 rank houses $h_1 \succ_{i_1,i_3} h_4 \succ_{i_1,i_3} h_2 \succ_{i_1,i_3} h_3$ and agents i_2 and i_4 rank houses $h_4 \succ_{i_2,i_4} h_2 \succ_{i_2,i_4} h_3 \succ_{i_2,i_4} h_1$. Under this preference profile, $\succ_{\{i_1,i_2,i_3,i_4\}}$, in the first round, broker i_4 obtains object h_2 in a cycle $i_4 \to h_2 \to i_2 \to h_4 \to i_4$. However if i_2 submitted instead preference ranking $\succ'_{i_2}$ identical to $\succ_{i_1,i_3}$, then i_4 and i_2 would not swap houses in first round. They would both still be unmatched in round 2, and i_4 would have lost his brokerage right; house h_4 would now be owned by i_2 (notice that not only it is so in the example but in fact condition R6 requires that i_2 owns h_4 when i_1 becomes matched and i_4 looses the brokerage right). Agent i_2 would then match with h_4 in round 2. In round 3, agent i_4 would become owner of h_2 (again this is so in the example, and, importantly it is guaranteed by the broker-to-heir transition property of R6). Thus, in round 3 agent i_4 would match with h_2. While the cycles are different, the allocations are the same. Looking at requirement R6 (and its broker-to-heir corollary) can give us a sense why—even if one or two agent cycles are different at $\succ$ and $\left(\succ'_{i_2}, \succ_{-i_2}\right)$—such or a similar scenario is bound to happen.

The most interesting part of the proof of Theorem 2 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 natural steps: we first construct the candidate control rights structure (c, b), then show that it satisfies conditions R1-R6, and finally show that the resultant TC mechanism $\psi^{c,b}$ equals φ.

We define a candidate control rights structure in terms of how φ allocates objects for preferences from some special preference classes. To see how this is done, consider the empty submatching and a house h. If φ were a TC and h was owned by an agent then at all preference profiles in which all agents rank h as their most preferred house, φ would allocate h to the same agent – the owner of h at the empty submatching. We thus check whether φ allocates h to the same agent at all above profiles, and if it does, we call this agent the candidate owner of h (in the proof, for brevity, we refer to the candidate owner as the owner*). If φ does not allocate h to the same agent at all above profiles, h is a candidate brokered house. Notice that if φ were a TC and h was brokered by an agent, then at every profile at which every agent ranks h as his most preferred house and some other house h' as his second-most preferred house, φ would allocate h' to the same agent – the broker of h at the empty submatching. We thus check whether there is an agent who always gets his second-most preferred house at the above profiles, and if there is such an agent we call this agent the candidate broker of h (broker* for short). Finally, we prove the key result that every house h either has a candidate owner or a candidate broker. This key result is technically the hardest one to obtain in the paper: we establish it by first showing that if all agents rank some house A as their first choice, and some house B as their second choice, then who gets A does not depend on the rest of the preference profile (Lemma 9); we then slowly refine this insight.

The construction of candidate control rights at non-empty submatchings is similar. The only modification is that instead of looking at preferences at which all agents agree on their most preferred house (or two most preferred houses), we impose this commonality only on unmatched agents, and at the same time we assume the matched agents rank the houses they are matched with at the top, while all other agents rank matched houses at the bottom. Thanks to the simplifying assumption that $|H| \geq |I|$, the Pareto efficiency of TC mechanisms implies that the above procedure would work well if φ was a TC, and we prove that indeed it works well whenever φ is group strategy-proof and efficient.[30]

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

5.1 Properties of Strategy-Proof and Efficient Mechanisms

Knowing that all individually strategy-proof and Arrovian-efficient mechanisms—equivalently all group strategy-proof and Pareto-efficient mechanisms—are trading-cycles mechanisms allows us to derive properties common to all such mechanisms.

Let us start by noticing that in any trading cycle mechanism, and for any preference profile, there is a group of agents—the decisive group—each of whom can get one of their two top choices, and all but at most one of them can get their top choice, irrespective of preferences submitted by agents not in the group.

Corollary 1. *(Decisive Group) Fix an individually strategy-proof and Arrovian-efficient mechanism ϕ. For any preference profile $\succ$, there is a group of agents $I_1 \subseteq I$ such that: (i) all agents from I_1 get one of their two top choices, and all but at most one of them get their top choices, and (ii) the allocation of agents from I_1 does not depend on preferences of agents not in I_1, that is for all $\succ'$ we have*

$$\phi\left(\succ\right)|_{I_1} = \phi\left(\succ_{I_1}, \succ'_{I-I_1}\right)|_{I_1}.$$

Notice that this corollary implies that if all agents rank house A first and house B second then who gets A does not depend on how the agents rank houses below B.

We further observe that all strategy-proof and efficient mechanisms have a recursive structure: the agents in the decisive group determine their allocation; given their preferences there is another group of agents who obtain one of their top two choices and who can determine their allocation irrespective of the preferences of others, etc. For instance, in a serial dictatorship (Satterthwaite and Sonnenschein, 1981; Svensson, 1994, 1999; Ergin, 2000), which is a special case of trading cycles, the first dictator chooses his most preferred object, then a second dictator chooses his most preferred object among the objects which were not chosen by the prior dictators, and so on until all agents have objects.

[30]This point in the construction requires more care in the case of $|H| < |I|$; see Section 5.2.

Corollary 2. *(Recursive Structure) Fix an individually strategy-proof and Arrovian-efficient mechanism ϕ. For every preference profile $\succ$, there is a partition $I_1, ..., I_k$ of the set of agents such that: (i) all agents from I_ℓ get one of their two top choices among objects unmatched at $\phi(\succ)(I_1 \cup ... \cup I_{\ell-1})$, and all but at most one of them gets their top choice among objects unmatched at $\phi(\succ)(I_1 \cup ... \cup I_{\ell-1})$, and (ii) the allocation of agents from I_ℓ does not depend on preferences of agents not in $I_1 \cup ... \cup I_{\ell-1} \cup I_\ell$.*

5.2 Outside Options

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

We generalize the concept of submatching as follows: For $J \subseteq I$, a submatching is a one-to-one function $\sigma : J \to H \cup Y$ such that each agent is matched with a house or his outside option.[31] A terminological warning is in order. A natural interpretation of the outside option is remaining unmatched. We will not refer to the outside option in this way, however, in order to avoid confusion with our submatching terminology. As in the main body of the paper, whenever we say that an agent is unmatched at σ, we refer to agents from $\overline{I_\sigma} = I - I_\sigma$. An agent is considered matched even if he is matched to his outside option.

The control rights structures (c, b) and their consistency R1-R6 are defined as before though the meaning of some terms such as submatching has changed, as explained above. In particular, (i) only houses are owned or brokered, the outside options are not; and (ii) control rights are defined for all submatchings, including submatchings in which some agents are matched with their outside options.[32]

We adjust the definition of round r of the TC algorithm to accommodate the outside

[31] As before, $\mathcal{S}$ denotes the set of submatchings, I_σ denotes the set of agents matched by σ, $H_\sigma \subseteq H$ denotes the set of houses matched by σ, and we use the standard function notation so that $\sigma(i)$ is the assignment of agent $i \in I_\sigma$, $\sigma^{-1}(h)$ is the agent that got house $h \in \sigma(I_\sigma)$, and $\sigma^{-1}(Y)$ is the set of agents matched to their outside options. A matching is a maximal submatching, that is, $\mu \in \mathcal{S}$ is a matching if $I_\mu = I$. As before, $\mathcal{M} \subset \mathcal{S}$ is the set of matchings. A mechanism is a mapping $\varphi : \mathbf{P} \longrightarrow \mathcal{M}$ that assigns a matching for each preference profile. Mechanisms, efficiency, and group strategy-proofness are defined as before.

[32] Notice that if a control rights structure is consistent on the domain with outside options, and $|H| \geq |I|$, then the restriction of the control rights structure to submatchings in which all agents are matched with houses is a consistent control rights structure in the sense of Section 4.

options as follows:

- In Step 1 we add the provisions that: (i) if an agent prefers his outside option to all unmatched houses, the agent points to the outside option; (ii) if there is a broker for whom the brokered house is the only acceptable house, such a broker also points to his outside option; and (iii) the outside option of each agent points to the agent.

- In Step 2 we allow exchange cycles in which agents points to their outside options.

- In Step 3 we match agents in each exchange cycle that does not contain the broker; we match agents in the exchange cycle of the broker if and only if there is at least one owner who points to the brokered house. We also re-define σ^r as the union of σ^{r-1} and the set of agent-house pairs and agent-outside option pairs matched in Step 3.

We refer to this modified algorithm as outside-options TC, and when there is no risk of confusion, simply as TC. We refer to the mechanism $\psi^{c,b}$ resulting from running the outside-options TC on consistent control rights structures as **outside-options TC**, or **TC**.

The outside-options TC mechanism described above can be used to allocate houses in the setting without outside options, and in this setting it is identical to the simpler TC mechanism described previously. In particular, the Step 3 provision that we match the broker if and only if there is at least one owner who points to the brokered house is only binding if the broker points to the outside option.[33]

In the presence of outside options, the TC class of mechanisms again coincides with the class of Pareto-efficient and group strategy-proof direct mechanisms, and hence with the class of individually strategy-proof and Arrovian efficient mechanisms.

Theorem 3. *In the environment with outside options, the following statements are equivalent:*

- *a mechanism is individually strategy-proof and Arrovian efficient,*
- *the mechanism is group strategy-proof and Pareto efficient,*
- *the mechanism is group strategy-proof and Arrovian efficient,*
- *the mechanism is an outside-options Trading-Cycles mechanism.*

The proof resembles the proofs of Theorems 1 and 2; the required modifications are discussed in Appendix G.

[33]In the presence of the outside options, this provision is important to assure the Pareto efficiency of the outside-options TC mechanisms.

6 Applications

6.1 Individually Rational House Allocation and Exchange

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

6.1.1 Model of House Allocation and Exchange

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

Pareto efficiency and group strategy-proofness are defined in the same way as in Section 2. In particular, the equivalence between group strategy-proofness and the conjunction of individual strategy-proofness and non-bossiness 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. Formally, a matching μ is individually rational if

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

For agents with empty endowments, $H_i = \emptyset$, this condition is tautologically true. While we formulate this definition for the setting without outside options all the results of this section, and their proofs, remain true if we allow outside options: we then define a matching μ to be individually rational if $\mu(i)$ is at least as good as the outside option and $\mu(i) \succeq_i h \qquad \forall i \in I, \forall h \in H_i$.

6.1.2 Results on Individually Rational Allocation and Exchange

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

Theorem 4. *In house allocation and exchange problems, the following three properties are equivalent:*

a mechanism is individually rational, group strategy-proof, and Pareto efficient;

a mechanism is individually rational, individually strategy-proof, and Arrovian efficient; and

a mechanism is an individually rational TC mechanism.

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

Proposition 1. *In house allocation and exchange problems, a TC mechanism is individually rational if and only if it may be represented by a consistent control rights structure in which each agent is given the initial ownership rights of all houses from his endowment.*[34]

Proof of Proposition 1. To prove individual rationality of the above subclass of TC mechanisms, consider an agent i and assume that at the empty submatching i owns a house h from his endowment. Then R4 ensures that i owns h throughout the execution of the TC algorithm. Thus, the TC mechanism will allocate to i house h or a house that i prefers to h. Now, let ψ be an individually rational TC mechanism. Recall that ownership* was defined in the proof of Theorem 2. For any agent i and house h from i's endowment, i is owner* of h because individual rationality implies that $\psi[\succ](i) = h$ for any $\succ \in \mathbf{P}[\varnothing, h]$, which is the set of preference profiles that rank h first for all agents. The construction from the proof of Theorem 2 thus represents ψ and yields a control rights structure that assigns to each agent the initial ownership rights over the houses from his endowment. **QED**

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

Theorem 5. *In house allocation and exchange problems where each agent has a nonempty endowment, the following three properties are equivalent:*

[34] Notice that when one agent is endowed with all houses, there are individually rational mechanisms that might be represented both by a control rights structure that assigns this agent initial ownership rights over all houses, and by an alternative control rights structure that assigns this agent ownership rights over all houses but one, 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.

a mechanism is individually rational, group strategy-proof, and Pareto efficient;

a mechanism is individually rational, individually strategy-proof, and Arrovian efficient; and

a mechanism is a TTC mechanism (aka hierarchical exchange) that assigns all agents the initial ownership rights of houses from their endowment.

Proof of Theorem 5. By Theorem 4, a mechanism φ is individually rational, Pareto efficient and group strategy-proof if and only if there exists an individually rational and consistent control rights structure (c, b) such that $\varphi = \psi^{c,b}$. By Proposition 1 we may assume that each agent has initial ownership rights over the houses from their endowment. By condition R4 of consistency all unmatched agents own a house throughout the mechanism, and hence R3 implies that no agent is a broker. $\psi^{c,b}$ is thus a TTC mechanism. **QED**

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

6.1.3 Market Design Environments

The assumptions of Theorem 4 are satisfied by the *house allocation problem with existing tenants* of Abdulkadiroğlu and Sönmez (1999). Theirs is the subclass of house allocation and exchange problems in which each agent is endowed with one or zero houses. In the former case, the agent is referred to as an *existing tenant*. The house allocation problem with existing tenants is modeled after dormitory assignment problems in US college campuses. In each such college, at the beginning of the academic year, there are new senior, junior, and sophomore students, each of whom already occupies a room from the last academic year. There are vacated rooms by the graduating class and there are new freshmen who would like to obtain a room, though they do not currently occupy any.

The assumptions of Theorem 5 are satisfied by the *kidney exchange* with strict preferences (Roth, Sönmez, and Ünver, 2004), and the *kidney exchange problem with good Samaritan donors* (Sönmez and Ünver, 2006). Kidney transplant patients are the agents and live kidney donors are the houses. Each agent is endowed with a live donor who would like to donate a kidney if his paired-donor receives a transplant in return. Thus, all agents have nonempty endowments. The model also allows for unattached donors known as good Samaritan donors who would like to donate a kidney to any patient. In the US, good Samaritan donors have been the driving force behind kidney exchange since 2006. Many regional programs such as the Alliance for Paired Donation (centered in Toledo, Ohio) and the New England Program for Kidney Exchange (centered in Newton, Massachusetts) have used good Samaritan donors in many of kidney exchanges conducted since 2006 (cf. Rees, Kopke, Pelletier, Segev, Rutter, Fabrega, Rogers, Pankewycz, Hiller, Roth, Sandholm, Ünver, and Montgomery, 2009).

The kidney exchange context underscores the importance of group strategy-proofness. The doctors of patients are the ones who have the information about patients' preferences over kidneys and it is known that doctors (or transplant centers) themselves at times manipulate the system to benefit their patients.[35] An individually strategy-proof mechanism that is not group strategy-proof could thus by manipulated by doctors. Group strategy-proofness guarantees that no doctor is able to manipulate the mechanism on behalf of his or her patients without harming at least one of them.

6.2 Complete Social Welfare Rankings

Our main results show that the class of individually strategy-proof and Arrovian efficient mechanisms is exactly the class of group strategy-proof and Pareto efficient mechanisms, which is exactly the class of Trading-Cycles mechanisms. In these results we allowed welfare functions to incompletely rank social outcomes. As an application of the main results, we now show that sequential dictatorships are exactly the mechanisms that are strategy-proof and Arrovian efficient with respect to complete SWF, that is SWF that always rank all outcomes.[36]

Let us start with an example showing that not all Top Trading Cycles are efficient with respect to a complete Arrovian SWF.

Example 4. Consider allocating two objects to two agents. Let ϕ be a top-trading-cycles mechanism in which agent 1 owns house h_1 and agent 2 owns house h_2 at the empty submatching. We will show that there is no complete SWF such that ϕ is efficient.

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

[36] We study the setting with outside options. Without outside options there are non-dictatorial mechanisms that are individually strategy-proof and Arrovian efficient with respect to a complete SWF. Each such mechanisms allocates objects as a sequential dictatorship as long as there are three or more objects left. With only two objects left the mechanism can proceed as a sequential dictatorship, or it can assign the two objects to two different agents, and run a round of TTC. This is similar to a situation from voting theory, wherein there are non-dictatorial mechanisms to vote on two outcomes. The proof of this characterization is based on similar ideas as the proof of Theorem 6 below, and we omit it.

Consider the preference profile $\succ$ such that

$$h_2 \succ_1 h_1 \succ_1 \emptyset \quad \text{and} \quad h_1 \succ_2 h_2 \succ_2 \emptyset,$$

where $\emptyset$ denotes the outside option. Consider also the following four additional preference profiles

$$\succ^1 = \begin{array}{|c|c|} 1 & 2 \\ \hline h_2 & \emptyset \\ \emptyset & \end{array}, \quad \succ^2 = \begin{array}{|c|c|} 1 & 2 \\ \hline h_2 & h_2 \\ \emptyset & \emptyset \end{array}, \quad \succ^3 = \begin{array}{|c|c|} 1 & 2 \\ \hline \emptyset & h_1 \\ & h_2 \end{array}, \quad \succ^4 = \begin{array}{|c|c|} 1 & 2 \\ \hline h_1 & h_1 \\ \emptyset & \emptyset \end{array},$$

where non-listed objects are worse than the outside option $\emptyset$. Denote

$$\begin{aligned} \mu_1 &= \phi\left(\succ^1\right) = \{(1, h_2), (2, \emptyset)\}; \\ \mu_2 &= \phi\left(\succ^2\right) = \{(1, \emptyset), (2, h_2)\}; \\ \mu_3 &= \phi\left(\succ^3\right) = \{(1, \emptyset), (2, h_1)\}; \\ \mu_4 &= \phi\left(\succ^4\right) = \{(1, h_1), (2, \emptyset)\}. \end{aligned}$$

Now, if there is a complete SWF Φ such that ϕ is efficient, then $\Phi(\succ^1)$ ranks $\phi(\succ^1)$ strictly above $\phi(\succ^4)$, and, by the independence of irrelevant alternatives, this implies that $\Phi(\succ)$ ranks $\phi(\succ^1)$ strictly above $\phi(\succ^4)$. Similarly, $\Phi(\succ^2)$ ranks $\phi(\succ^2)$ strictly above $\phi(\succ^1)$, and, by the independence of irrelevant alternatives, this implies that $\Phi(\succ)$ ranks $\phi(\succ^2)$ strictly above $\phi(\succ^1)$. Further, and again similarly, $\Phi(\succ^3)$ ranks $\phi(\succ^3)$ strictly above $\phi(\succ^2)$, and, by the independence of irrelevant alternatives, this implies that $\Phi(\succ)$ ranks $\phi(\succ^3)$ strictly above $\phi(\succ^2)$. Finally, $\Phi(\succ^4)$ ranks $\phi(\succ^4)$ strictly above $\phi(\succ^3)$, and, by the independence of irrelevant alternatives, this implies that $\Phi(\succ)$ ranks $\phi(\succ^4)$ strictly above $\phi(\succ^3)$. But, then $\Phi(\succ)$ fails transitivity, showing that there does not exist a complete SWF with respect to which ϕ is efficient.

We will use this example to prove the following

Theorem 6. *A mechanism is individually strategy-proof and Arrovian efficient with respect to a complete SWF if and only if it is a sequential dictatorship.*

Proof. ($\Longrightarrow$) Consider a mechanism ϕ that is individually strategy-proof and efficient with respect to a complete Arrovian welfare function. By Theorems 1 and 2, ϕ is a Trading Cycles mechanism $\psi^{c,b}$. Fix an arbitrary preference profile $\succ_I$.

We claim that at any round r of the algorithm $\psi^{c,b}$, there is exactly one agent who controls all objects. We prove it in two steps.

First, let us show that there cannot be two (or more) agents who each owns an object. By way of contradiction, suppose that some agent 1 controls object h_1 and some other agent 2 controls object h_2 in round r. Let σ be the submatching created by the algorithm $\psi^{c,b}$ before round r. Consider four auxiliary preference profiles $\succ^{\ell}$ that all share the following properties: (i) all agents matched under σ rank objects under $\succ^{\ell}$, $\ell = 1, ..., 4$, in the same way they rank them under $\succ$, (ii) all agents unmatched at σ and different from agents 1 and 2 find all objects to be unacceptable, and (iii) agents 1 and 2 rank all objects matched at σ above h_1 and h_2 and they find all objects other than h_1, h_2 and objects matched at σ to be unacceptable. The four profiles differ only in how agents 1 and 2 rank objects h_1 and h_2, and which of these objects are acceptable to them: the ranking of these two objects is the same as in the four preference profiles of the above example. Notice that

$$\psi^{c,b}\left(\succ^{\ell}\right) = \sigma \cup \mu_{\ell},$$

where μ_{ℓ} are defined as in the above example. Furthermore, the same argument we used in the example shows that there can be no welfare functional that ranks all four μ_{ℓ}, is transitive, and satisfies the independence of irrelevant alternatives. Hence, there is no complete Arrovian welfare function that makes $\psi^{c,b}$ efficient, a contradiction that implies that there cannot be two agents who own objects in a round of the algorithm.

As $\psi^{c,b}$ never allows two owners in a round of the algorithm, Proposition 3 in Appendix F.1 allows us to assume that there are no brokers in any round, either. Hence, in each round of the algorithm there is a single agent who controls (and owns) all houses. That means that $\psi^{c,b}$ is a sequential dictatorship.

($\Longleftarrow$) Consider a sequential dictatorship $\psi^{c,b}$. We construct a complete Arrovian welfare functional Φ such that $\psi^{c,b}$ is efficient with respect to Φ. Under Φ any two matchings are ranked according to preferences of the first-round dictator; if he is indifferent then the matchings are ranked according to the preferences of the second-round dictator, etc. Formally, for any $\succ \in \mathbf{P}$ and any two distinct $\mu, \nu \in \mathcal{M}$, let $\mu \, \Phi(\succ) \, \nu$ if and only if there exists $k \in \{1, ..., |I|\}$ such that $\mu(i_1) = \nu(i_1)$, ... and $\mu(i_{k-1}) = \nu(i_{k-1})$, and agent i_k strictly prefers $\mu(i_k)$ over $\nu(i_k)$, where agents $i_1, ..., i_k$ are defined recursively: $i_1 = c(\emptyset)$, and in general $i_{\ell} = c\left(\left(\left(i_1, \mu(i_1)\right), ..., \left(i_{\ell-1}, \mu(i_{\ell-1})\right)\right)\right)$ for $\ell = 1, ..., k$. It is straightforward to verify that Φ is a complete Arrovian welfare functional and that $\psi^{c,b}$ is efficient with respect to Φ. QED

6.3 Further Illustrative Applications

In Section 6, we have so far illustrated how our main results allow one to easily obtain new insights into allocation and exchange problems.[37] We now show that some of the deepest prior insights easily follow from our Theorem 2.

6.3.1 Serial Dictatorships

Neutrality and group strategy-proofness were characterized through serial dictatorships by Svensson (1999) when there are no outside options. In a serial dictatorship agents are ordered, the first agent in the ordering gets his most preferred house, the second agent in the ordering gets her most preferred among houses unassigned to agents higher in the ordering, etc. Whether we allow outside options or not, Svensson's result is a corollary of Theorem 2 as illustrated below.

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

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

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

It remains to show that any neutral trading cycles mechanism $\psi^{c,b}$ is equivalent to a serial dictatorship. Let $\sigma \in \overline{\mathcal{M}}$. By R1 and R3 there is an agent $i \in \overline{I}_\sigma$ 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 $j \in I_\sigma$.

[37] Our main results have many additional applications. As an example, consider a mechanism ϕ that is invariant, that is for any agent $i \in I$ and any object $h \in H$ if for all $g \succsim_i h$ we have $\phi(\succ_i, \succ_{-i})(i) = g \Longleftrightarrow \phi(\succ_i', \succ_{-i})(i) = g$, then for all $g \succsim_i h$ and all $j \in I$ we have $\phi(\succ_i, \succ_{-i})(j) = g \Longleftrightarrow \phi(\succ_i', \succ_{-i})(j) = g$. We can then show that ϕ is individually strategy-proof and Pareto efficient if and only if it is a hierarchical exchange mechanism of Pápai (2000). For the proof, notice that invariance implies non-bossiness, and hence, together with individual strategy-proofness it implies that ϕ is group strategy-proof. The rest of the proof resembles the proof of Corollary 2 below.

Now, $\succ^{\pi} \in \mathbf{P}[\sigma'; g]$ and by neutrality $\psi^{c,b}[\succ^{\pi}](i) = \pi(h) = g$. Maskin monotonicity implies that i is allocated the best unmatched house at σ' as long as $I_{\sigma'} = I_{\sigma}$. The mechanism $\psi^{c,b}$ is thus equivalent to a serial dictatorship. **QED**

6.3.2 Reallocation-proofness

Pápai (2000) showed that group strategy-proof and Pareto efficient mechanisms that satisfy an additional condition can be implemented as hierarchical exchange (i.e., TTC) mechanisms. The condition she relies on is as follows: a mechanism φ is **reallocation-proof** if there exists no pair of agents $i, j \in I$ such that for some $\succ \in \mathbf{P}$, $\succ'_i \in \mathbf{P}_i$, and $\succ'_j \in \mathbf{P}_j$ with $\varphi[\succ'_i, \succ_{-i}] = \varphi[\succ'_j, \succ_{-j}] = \varphi[\succ]$, we have $\varphi[\succ'_{\{i,j\}}, \succ_{-\{i,j\}}](j) \succ_i \varphi[\succ](i)$ and $\varphi[\succ'_{\{i,j\}}, \succ_{-\{i,j\}}](i) \succsim_j \varphi[\succ](j)$. We can derive the key insight of Pápai (2000) as 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 of Corollary 2. Let φ be a group strategy-proof, efficient, and reallocation-proof mechanism. By Theorem 2, it is equivalent to a reallocation-proof TC mechanism $\psi^{c,b}$. It remains to show that the control rights structure (c, b) can be chosen in such a way that there are no brokers. Take any submatching $\sigma \in \overline{\mathcal{M}}$. First notice that if there are two owners, j and k at σ, then no house is σ-brokered. Indeed, by way of contradiction assume that some house h is σ-brokered by an agent i, and let h_j be a house owned by j and h_k be a house owned by k. Consider a preference profile $\succ \in \mathbf{P}[\sigma]$ and such that $\succ_i \in \mathbf{P}_i[\sigma; h, h_k]$, $\succ_j \in \mathbf{P}_j[\sigma; h_j]$, and $\succ_k \in \mathbf{P}_k[\sigma; h]$. Then, the deviation to $\succ'_i \in \mathbf{P}_i[\sigma; h, h_j, h_k]$ and $\succ'_j \in \mathbf{P}_j[\sigma; h, h_j]$ violates the reallocation-proofness condition. Hence, (c, b) can allow brokers only at submatchings with a unique owner. But then $\psi^{c,b}$ is equivalent to $\psi^{c',b'}$ such that (c', b') is identical to (c, b) except that at any submatching σ at which (c, b) gives brokerage right over a house h to an agent i, the primed control rights structure (c', b') gives ownership of h to the unique (c, b) owner j at σ, and it gives i the ownership of all unmatched houses at $\sigma \cup \{(j, h)\}$. **QED**

7 Conclusion

We study allocation and exchange in environments without transfers and with single-unit demands. Addressing central concerns of both practical allocation problems and the relevant theoretical literature, we (i) show that a mechanism is individually strategy-proof and always selects the efficient outcome with respect to a social ranking satisfying Arrovian postulates if and only if the mechanism is group strategy-proof and Pareto efficient, and (ii) construct the full class of these mechanisms.

Our construction relies on the introduction of brokers to allocation and exchange problems. While in the context of our paper, the main role played by the brokers is to allow us to construct the full class of strategy-proof and efficient mechanisms, let us conclude with an example showing how brokers can be useful in some mechanism design settings. Before describing the example, let us stress that we have already seen that there are social welfare criteria (functionals) that can only be satisfied by brokered TC mechanisms. The example below shows that in asymmetric settings we may want to use brokered TC over TTC because of equity considerations.

Consider a manager who assigns n tasks $h_1, ..., h_n$ to n employees $i_1, ..., i_n$ 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 not as proficient as the other employees in conducting this task). That is, if there is a Pareto-efficient matching that avoids assigning h_n to i_n, she would like to choose such a matching. Because she does not know employees' preferences, she wants to use a group strategy-proof mechanism.

The manager can use one of the Top Trading Cycles mechanisms of Pápai (2000); to do so the manager would need to initially endow employees $i_1, ..., i_{n-1}$ with all the tasks, and can allow i_n to inherit some task only after either (i) all other employees are already matched, or (ii) the task is h_n is matched. The manager can also use 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.

Because the manager's problem is asymmetric, all mechanisms treat agents in an asymmetric way. We show, however, that the above brokered trading-cycles mechanism is strictly more equitable than any top-trading-cycles mechanism. We compare the equitability of the mechanisms using the equity ranking introduced by Lorenz (1905).[38] The outcomes of employees depend of course on the profile of their preferences, and in particular the comparison of the equitability of the mechanisms depends on the preference profile. To get a Lorenz comparison we thus look at many problems that differ only in employees' preference profiles. We can interpret this as either taking an ex ante view, before the employees drew their rankings from a distribution, or as looking at a population of managers and their employees. We further restrict attention to the canonical case in which all employees rank the tasks in the same way and each ranking is equally likely (or, in the population interpretation, represented by an identical number of managers' problems).[39]

[38] In line with the rest of the paper, we define Lorenz dominance in our ordinal setting. Let us say however that the cardinal analogue of our result in this section is, and its proof is effectively the same. In particular, the best TC mechanism dominates all TTC mechanisms in terms of the Gini coefficient.

[39] Our result would also be true if we looked at other natural distributions of preference profiles, for instance

To formally define Lorenz dominance let us first introduce notation capturing the welfare of agents: given a distribution $\hat{\succ}$ over preference profiles, we denote by $\rho^{\varphi[\hat{\succ}]}(i,r)$ the probability that i receives at least his r'th choice object under a mechanism φ. We say that a mechanism φ **Lorenz dominates** mechanism ψ for a random preference profile $\hat{\succ}$ if for all $r=1,...,|H|$ there is an ordering of agents $i_1,...,i_{|I|}$ such that for any $k=1,...,|I|$ we have

$$\min_{J\subseteq I,|J|=k}\sum_{j\in J}\rho^{\varphi[\hat{\succ}]}(j,r) \ \geq \ \sum_{\ell=1}^{k}\rho^{\psi[\hat{\succ}]}(i_\ell,r).$$

The dominance is strict if there are r and k such that the inequality is strict.

With these definitions in place, we can state our concluding result

Proposition 2. *In the task assignment problem, the TC mechanisms such that i_n brokers h_n and each agent i_ℓ for $\ell=1,...,n-1$ owns h_ℓ strictly Lorenz dominates any TTC mechanism satisfying the managers constraints.*

A Appendix: Revelation Principle for Group Strategy-Proofness

The classical dominant-strategy revelation principle has a natural extension to group dominant strategies. A group dominant-strategy equilibrium of an indirect mechanism is a pure strategy profile such that for any realization of preferences, when any coalition of agents together report any other messages than the ones prescribed by the original strategy profile, then for any reports of agents outside of the coalition, either at least one of coalition members receives a lower payoff than he would get or all of the coalition members receive the same payoffs as they would get if the coalition played according to their prescribed strategies. The group dominant-strategy revelation principle says that for any indirect mechanism with a group dominant-strategy equilibrium, there exists a direct group strategy-proof mechanism under which truth-telling achieves the same payoffs for all agents as in this equilibrium of the original mechanism. Its proof is analogous to that of the classical dominant-strategy revelation principle.

B Appendix: Comments on Consistency Requirements

This appendix explains the consistency requirements R1-R6.

if we restricted attention to profiles in which all employees of a manager agreed on a ranking of the tasks.

Requirements R1 and R2 are needed to ensure that in Step 1 of the TC algorithm there is always an owned house for the broker to point to.

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

We discussed the role of requirement R4 in Section 4.1.

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

Example 5. ***Why do we need R5 to prevent individual manipulation?*** Consider four agents $i_1, ..., i_4$. Assume that at the empty submatching agent i_2 brokers a house and other agents own one house each. Denote by h_k the house controlled by agent i_k. Let us maintain R1-R3, R4, and R6, and violate R5 by assuming that h_2 is owned by i_4 at submatching $\{(i_1, h_1)\}$. Now, there are two previous owners unmatched at $\{(i_1, h_1)\}$, i_3 and i_4. Moreover, i_2 is no longer a broker. Consider now a preference profile such that h_1 is i_1's and i_2's mutual first-choice house, h_2 is the first choice of the other agents, and h_3 is the second choice of i_2 and i_3. Under this preference profile and control rights structure, in the first round of the TC algorithm, i_1 and i_2 point to h_1, while i_3 and i_4 point to h_2. House h_1 points to its owner i_1, and h_2 points to its broker i_2. There is a unique cycle: $h_1 \to i_1 \to h_1$, and we obtain the submatching $\{(i_1, h_1)\}$. In the second round, all remaining agents point to h_2 which is owned by i_4. Hence the unique cycle is $h_2 \to i_4 \to h_2$, and i_4 is matched with h_2. Thus, agent i_2 is neither matched with his first or second choice. This agent would benefit by misrepresenting his preferences and declaring h_3 to be his first choice: then in the first round of TC, he would point to h_3 completing the cycle $h_3 \to i_3 \to h_2 \to i_2 \to h_3$ and ending up matched to h_3, his true second choice house.

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

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

Example 6. ***Why do we need R6 to prevent individual manipulation?*** Consider four agents $i_1, ..., i_4$. Assume that at the empty submatching agent i_2 brokers h_2, i_1 owns h_1, h_4, and i_3 owns h_3. At submatching $\{(i_1, h_1)\}$, assume that i_3 owns h_2 as well, and i_2 loses his brokerage right. Now, i_4 inherits h_4 as an owner. We assume R1-R5, and violate R6. R5 is not violated, as there is a single previous owner unmatched at $\{(i_1, h_1)\}$, and he is i_3. However, R6 is violated, as at the submatching $\{(i_1, h_1), (i_3, h_2\}$, i_2 is not the heir to i_3. That is, i_2 does not own the ex-owned house h_3 of i_3, but i_4 does. Consider the preference profile at which agents i_1 and i_2 have house h_1, i_3 has h_2 and i_4 has h_3 as their first choices; and agent i_2's second choice is h_3. Then, i_2 would benefit by ranking h_3 first.

Example 7. ***Why do we need R6 to prevent bossiness?*** Consider the same control rights structure as in Example 6. Consider the preference profile at which i_1 and i_3's first choices are h_1, and i_2 and i_4's first choices are h_3, and i_3's second choice is h_2. Now, whether agent i_3 reports the above ranking or ranks h_2 first, he receives house h_2. However, in the first case, i_2 receives h_4, while in the latter, he receives h_3. Thus, without R6, the mechanism could be bossy.

Example 8. ***Can we replace R5-R6 by a simpler (and stronger) persistence of brokerage property?*** Consider the following property "if $|\sigma'| < |I| - 1$ and agent i brokers house h at σ and is unmatched at $\sigma' \supset \sigma$, then i brokers h at σ' " (an analogue of R4 for brokers). The following example shows that we cannot replace R5-R6 with this. Consider an environment with four agents, i_1, i_2, i_3, i_4, four houses, h_1, h_2, h_3, h_4, and a TC mechanism $\psi^{c,b}$ whose control rights structure (c, b) is explained below and illustrated by the table in Figure 1.

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

House h_2 is owned by i_2. When i_2 is matched but h_2 is not then h_2 is inherited by one of the unmatched agents; who inherits h_2 depends on how the matched agents are matched

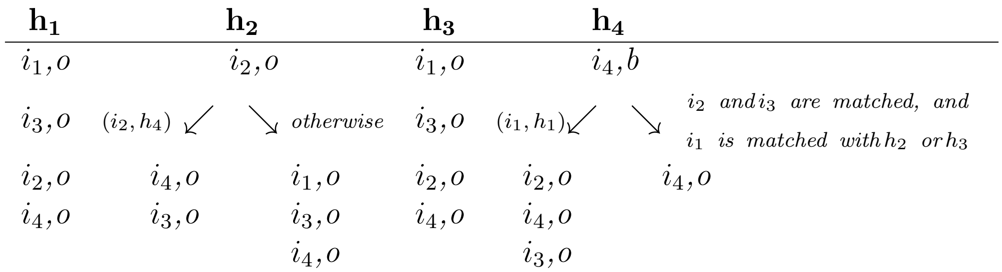

Figure 1: A control rights structure with broker-to-heir transition

(the submatching). If i_1 is matched with h_1 and i_2 is matched with h_4, then the next owners of h_2 are i_4 and i_3, in this order. In all other cases, the order of the next owners of h_2 is i_1, i_3, and i_4.

House h_4 is initially brokered by agent i_4 (denoted by "b" next to i_4 in the figure). Agent i_4 continues to broker h_4 as long as he is unmatched with two exceptions: (i) if i_1 is matched with h_1 then i_4 loses the brokerage right, and h_4 becomes an owned house with the order of owners i_2, i_4, and i_3; and (ii) if i_4 is the only remaining agent, then he owns h_4. The second exception is dictated by R2. We explain in detail how the first exception occurs and why it is consistent with our conditions. We use the notation of the statement of R6 and denote $\sigma = \emptyset$, $\sigma' = \{(i_1, h_1)\}$, $i = i_2$, $h = h_2$, $i' = i_4$, $h' = h_4$. Now, i_4 brokers h_4 at $\emptyset$ and i_2 owns h_2 at $\emptyset$: at $\sigma' = \{(i_1, h_1)\}$, as it is allowed by R5 that only one $\emptyset$-owner, i_2, is left unmatched, i_4 loses his brokerage right of h_4. By the second part of R6, at $\sigma' = \{(i_1, h_1)\}$, i_2, the only remaining $\emptyset$-owner, owns h_4. On the other hand, at submatching $\{(i_2, h_4)\}$, by the first part of R6, i_4 owns h_2. However, this will only be relevant, when i_4 loses his brokerage right. This happens at σ'. Hence, by R4, i_4 owns h_2 at $\sigma' \cup \{(i_2, h_4)\} = \{(i_1, h_1), (i_2, h_4)\}$. This is what we refer to as the broker-to-heir transition.

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

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

for all $\succ \in \mathbf{P}$ such that all agents rank h_4 first and any $g \in \{h_1, h_3, h_2\}$ second.

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

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

However,

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

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

C Appendix: Proof of the Pareto Efficiency of TC Mechanisms (Part 1 of the Proof of Theorem 2)

We prove the proposition by a simple recursion. Consider the TC algorithm. Each agent matched in the first round of the algorithm gets his first or second choice house and is matched with a house controlled by an agent matched in the same round. Moreover, there is at most one agent who gets his second choice in the first round, as there is at most one broker. Therefore, if an agent matched in the first round gets his second choice, then getting his first choice would harm another agent matched in that round.

In general, each agent matched in the r-th round of the algorithm gets his first or second choice among the remaining houses and is matched with a house controlled by an agent matched in the same round. Moreover, there is at most one agent who gets his second choice in this round, as there is at most one broker. Therefore, if an agent matched in the r-th round were given a better house, this would harm some other agent matched at the same or an earlier round. **QED**

D Appendix: Proof of the Individual Strategy-Proofness of TC Mechanisms (Part 2 of the Proof of Theorem 2)

Let $\psi^{c,b}$ be a TC mechanism. Let $\succ$ be a preference profile. We fix an agent $i \in I$. We will show that i cannot benefit by submitting $\succ'_i \neq \succ_i$ when the other agents submit $\succ_{-i}$. Let s

be the round i leaves (with house h) under $\succ_i$ and s' be the time i leaves (with h') under $\succ'_i$ in the algorithm. We will consider two cases.

Case 1. $s \leq s'$**:** At round s, the same houses and agents are in the market under both $\succ_i$ and $\succ'_i$ by Lemma 4. If i is not a broker at time s under $\succ_i$, then, by submitting $\succ_i$, agent i gets the top-choice house among the remaining ones in round s, implying that he cannot be better off by submitting $\succ'_i$.

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

It remains for us to consider the situation in which e is broker i's top-choice remaining house, and to show that i cannot get e by submitting the profile $\succ'_i$. For an argument by contradiction, assume that under $\succ'_i$ agent i leaves at round s' with house e. Because agent i is a broker when he leaves at $\succ_i$, there is an agent j who is matched with house e at time s. At this time, j is an owner of some owned house h_j, and e is his top-choice house. By Lemma 4, the control rights structure at round s is the same under both $\succ_i$ and $\succ'_i$. Hence, i is also a broker at time s after submitting $\succ'_i$, and j is an owner of h_j. Moreover, j's top choice is still house e. That means that under $\succ'_i$ agent j will stay unmatched until $s' + 1$. Since agent i leaves with e at s', he cannot be the broker of e at this round, because a broker cannot leave with the brokered house, while another owner j is unmatched. Thus, there is a round $s'' \in \{s+1, ..., s'\}$ at which agent i stops being the broker of e. Since e is still unmatched at this round, there is a broker-to-heir transition between $s'' - 1$ and s'' (by R6). Because j is an owner of h_j at both $s'' - 1$ and s'', he would have inherited e at s'' (by R6). Thus, j would have left with e at s'', as e is j's top choice among houses left at s (and hence those left at s''). A contradiction.

Case 2. $s > s'$**:** At round s', the same houses and agents are in the market under both $\succ_i$ and $\succ'_i$ by Lemma 4. Consider round s' under both $\succ_i$ and $\succ'_i$. Under $\succ'_i$, agent i points to house $h' = h^1$ that points to agent i^1 that points to ... that points to object h^n that points to agent $i = i^n$ (and this cycle leaves at round s'). If the cycle is trivial ($n = 1$) and h' points back to i, then i owns h'. Since ownership persists by R4, i will own h' at $s > s'$, and thus at round s, agent i would leave with a house at least as good as h'.

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

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

If i^ℓ brokers h^ℓ for some $\ell \in \{1, ..., n\}$, then the chain $h' = h^1 \to i^1 \to h^2 \to ... \to h^n \to i$ will stay in the system as long as i^ℓ continues brokering h^ℓ (since there are no other brokerages and ownerships persist by R4). If i^ℓ brokers h^ℓ at round s under $\succ_i$, then we are done, since the same cycle would have formed. Thus suppose that at a round $s'' \in \{s'+1, ..., s\}$ broker i^ℓ loses his broker status. Because $n \geq 2$, agent $i^{\ell+1}$ is an owner at both rounds $s''-1$ and s''. Hence, the loss of brokerage status means that i^ℓ enters a broker-to-heir transition. We must then have $n = 2$ (since by R5, only one previous owner can remain unmatched during the broker-to-heir transition). There are two cases, as by R6's second part, the unique previous owner owns both houses, his previously owned house and the ex-brokered house: either i^1 owns $h^1 = h'$ and h^2 (and $i^2 = i^\ell$ is the heir) or $i^2 = i$ owns $h^1 = h'$ and h^2. In the former case, i^1 who wants h^2, will leave with it in round s'' under $\succ_i$, and $i^2 = i$ will inherit $h^1 = h'$ at $s''+1$ by R6, more precisely, by the broker-to-heir transition property. In the latter case, $i^1 = i$ owns $h^1 = h'$ already in round s''. In both cases, in round $s \geq s''$ agent i can only leave with a house at least as good as h' under $\succ_i$. **QED**

E Appendix: Proof of the Non-Bossiness of TC Mechanisms (Part 3 of the Proof of Theorem 2)

Let $\psi^{c,b}$ be a TC mechanism. Fix an agent $i_* \in I$ and two preference profiles $\succ = [\succ_{i_*}, \succ_{-i_*}]$ and $\succ' = [\succ'_{i_*}, \succ_{-i_*}]$ such that

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

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

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

we will prove the following stronger statement:

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

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

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

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

   or

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

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

By Lemma 4, the above hypothesis is true for all $r < s$. The proof for $r \geq s$ proceeds by induction over the round r.

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

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

is removed in round s. Lemma 4 implies that the same houses and agents are in the market at time s under both $\succ$ and $\succ'$ and that all agents from $I_{\sigma^s[\succ]} - \{i_*^1, ..., i_*^n\}$ are matched by $\sigma^s[\succ']$ in the same way as in $\sigma^s[\succ]$. Lemma 4 also implies that the chain $h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1 \to i_*^1$ forms at round s under preferences $\succ'$.

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

If $n > 1$ and i_*^ℓ brokers h_*^ℓ for some $\ell \in \{2, ..., n\}$, then the chain $h_*^2 \to ... \to h_*^n \to i_*^n \to h_*^1 \to i_*^1$ will stay in the system as long as i_*^ℓ continues to broker h_*^ℓ. If i_*^ℓ continues to broker h_*^ℓ until round s' under $\succsim'$, then the initial step is proved. Otherwise, there is a round $s'' \in \{s+1, ..., s'\}$ such that agent i_*^ℓ has the brokerage right over h_*^ℓ at rounds $s, ..., s''-1$

but not at round s''. By R6's broker-to-heir transition property, $n = 2$ and $i_*^{\ell+1}$ owns h_*^{ℓ} at $\sigma^{s''}[\succ']$ because he owns $h_*^{\ell+1}$at both $\sigma^{s''-1}[\succ']$ and $\sigma^{s''}[\succ']$. As $i_*^{\ell+1}$'s top preference is then h_*^{ℓ}, he leaves with it at s''. Applying again the broker-to-heir transition property, we see that agent i_*^{ℓ} inherits $h_*^{\ell+1}$ at $s''+1$ and will be matched with it. This case ends the proof of the inductive hypothesis for $r = s$.

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

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

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

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

Claim 2. Under $\succ'$:

- all houses i^{ℓ} prefers over $h^{\ell+1}$, except possibly h^{ℓ}, are matched with agents other than i^{ℓ};
- if i^{ℓ} is a $\sigma^{r-1}(\succ)$-owner then all houses i^{ℓ} prefers over $h^{\ell+1}$ are matched with agents other than i^{ℓ}.

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

is $\sigma^{r-1}[\succ]$-broker then all houses i^ℓ prefers over $h^{\ell+1}$, except possibly h^ℓ, are matched before round r under $\succ$, and the inductive assumption yields the claim. QED

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- $j^1 \notin I_{\sigma^{r-1}[\succ]}$. Then, both agents j^1 and i^1 are unmatched at the submatching ν, and Claim 1 implies that j^1 is a broker of h^1 at $\sigma^{t-1}[\succ']$. Let j' be an agent who obtains h^1 at time t under $\succ'$. As j^1 is a broker of h^1 at $\sigma^{t-1}[\succ']$, we have $j' \neq j^1$. Thus j' is an owner of a house h' at $\sigma^{t-1}[\succ']$, and $h' \neq h^1$.

  Agent j' is unmatched at $\sigma^{r-1}[\succ]$ because if he were matched then the inductive assumption would imply that h^1 is matched at $\sigma^{r-1}[\succ]$, which is false. Now, if h' was matched at $\sigma^{r-1}[\succ]$ then the inductive assumption would imply that either all agents in the $\succ'$-cycle of h' are unmatched at $\sigma^{r-1}[\succ]$ (which is false), or h' is matched under $\succ'$ in a cycle of one or two agents. Since the $\succ'$-cycle of h' contains at least two different agents j^1 and j', we would conclude that j^1 obtains h' under $\succ'$, and hence by the inductive assumption j^1 would be matched with h' in $\sigma^{r-1}[\succ]$. Since we consider the situation $j^1 \notin I_{\sigma^{r-1}[\succ]}$, we may conclude that h' is unmatched at $\sigma^{r-1}[\succ]$. Thus j'

and h' are unmatched at $\sigma^{r-1}[\succ]$, and hence at ν. Then, R4, R5, and R6 imply that j' owns h^1 at ν, and thus $j' = i^1$. Since i^1 points to h^1, they are matched in the cycle $h' \to i^1 \to h^1 \to j \to h'$, and the inductive hypothesis (part 3) is true for i^1.

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

Step 1. First, consider how h^1 is matched under $\succ'$. Let $h^0 \to j^0 \to h^1 \to j^1$ 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 $\sigma^{r-1}[\succ]$, and one of the two subcases obtains:

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

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

We handle these two subcases separately:

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

  - ⋆ Assume j^1 brokers h^1 at $\sigma^{t-1}[\succ']$. Then, $j^0 \neq j^1$, $h^0 \neq h^1$, 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 the brokerage right, R5 implies that there are only two agents in the cycle of h^1 under $\succ'$ or $\succ$, respectively; moreover, by R6, j^0 or i^n (respectively) owns h^1 at ν. However, then neither j^1 nor i^1 can broker h^1 at ν, implying that both lose brokerage rights, and hence, h^1 is owned by both j^0 and i^n at ν. Thus, $j^0 = i^n$, and $n = 2$. We conclude that i^2 is matched with h^1 under $\succ'$ and the cycle he gets matched in has two agents, i.e., $h^0 \to i^2 \to h^1 \to j^1 \to h^0$.
  - ⋆ Assume j^1 owns h^1 at $\sigma^{t-1}[\succ']$. Then, R4 implies that j^1 owns h^1 at ν. 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 R6, h^1 is owned by i^2 at ν. Since 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 $h^1 \to i^2 \to h^1$.

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

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

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

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

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

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

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

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

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

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

R6 and R4 j' owns h^2 at ν^2. Thus, $j' = i^1$. Hence, the cycle of h^2 under $\succ'$ is $h' \to i^1 \to h^2 \to j^2 \to h'$.

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

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

Case 3. *$n > 1$ and i^1 owns h^1 at $\sigma^{r-1}[\succ]$* (in particular, R4 implies that i^1 owns h^1 at ν). We will show that this case cannot happen. By the inductive assumption either

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

(b) $n = 2$ and the cycle $h^0 \to j^0 \to h^1 \to j^1 \to h^0$ occurs under $\succ'$ for some house h^0, and j^1 is matched with h^0 in $\sigma^{r-1}[\succ]$, that is $\{(j^1, h^0)\} \subseteq \sigma^{r-1}[\succ]$.

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

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

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

F Appendix: Proof of the Implementation Part of Theorem 2 (Part 4 of the Proof of Theorem 2)

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

Let us start by introducing some useful terms and notation. Let $\sigma \in \overline{\mathcal{M}}$, $n \geq 0$, $h^1, h^2, ..., h^n \in \overline{H_\sigma}$, and $i \in I$. Denote by $\mathbf{P}_i[\sigma, h^1, ..., h^n]$ the set of preferences $\succ_i$ of agent i such that

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

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

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

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

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

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

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

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

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

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

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

A house $e \in \overline{H_\sigma}$ is a **brokered* house at** $\sigma \in \overline{\mathcal{M}}$ if there exist some $\succ$ and $\succ' \in \mathbf{P}^*[\sigma, e]$ such that $\varphi[\succ]^{-1}(e) \neq \varphi[\succ']^{-1}(e)$. Agent k is the **broker* of** e **at** σ if e is a brokered* house at σ and for all $\succ \in \mathbf{P}^*[\sigma, e]$ house $\varphi[\succ](k)$ is the second choice of k in $\succ_k$. Observe that a house cannot be both owned* and brokered* at the same submatching. [40]

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

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

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

Lemma 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*

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

Lemma 6. *(Pápai 2000) For all $i, j \in I$, all $\succ \in \mathbf{P}$, and all $\succ_j^* \in \mathbf{P}_j$, if j envies i at $\succ$ and $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$, then there exists $\succ_i^* \in \mathbf{P}_i$ such that*

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

This last lemma allows us to prove

Lemma 7. *For all $i, j \in I$, all $\succ \in \mathbf{P}$, and all $\succ_j^* \in \mathbf{P}_j$, if j envies i at $\succ$, then*

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

Proof of Lemma 7. If $\varphi[\succ_j^*, \succ_{-j}](i) \neq \varphi[\succ](i)$ then the lemma follows from Lemma 6 and strategy-proofness of φ. If $\varphi[\succ_j^*, \succ_{-j}](i) = \varphi[\succ](i)$ then Pareto efficiency of $\varphi(\succ)$ implies that i cannot envy j when j envies i and hence the claim of the lemma follows. QED

[40] It may appear from the definitions that there is a third option for an unmatched house besides being owned* and brokered* at a submatching. However, Proposition 3 below shows that these are the only two options.

F.1 The Starred Control Rights Structure is Well Defined

The main result in this subsection (Proposition 3) shows that any house is either owned* or brokered*. Thus, the starred (candidate) control rights structure is well defined. We build up toward this result through several lemmas. All lemmas in this section are formulated and proven at a fixed submatching $\sigma \in \overline{\mathcal{M}}$.

Lemma 8. *Let $\sigma \in \overline{\mathcal{M}}$. For all $i \in I_\sigma$ and all $h \in \overline{H_\sigma}$,*

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

Proof of Lemma 8. Suppose that an agent in $i \in I_\sigma$ does not get $\sigma(i)$ at $\varphi[\succ]$. Then we can create a new matching by assigning all agents in $\overline{I_\sigma}$ that get a house in H_σ a house in $\overline{H_\sigma}$ that was assigned to an agent in I_σ, all other agents j in $\overline{I_\sigma}$ the house $\varphi[\succ](j)$, and all agents j in I_σ the house $\sigma(j)$. Since each agent in $\overline{I_\sigma}$ ranks houses in H_σ lower than houses in $\overline{H_\sigma}$ and each agent in I_σ ranks his σ-house as his first choice, this new matching Pareto dominates $\varphi[\succ]$, contradicting the fact that φ is Pareto efficient. **QED**

Lemma 9. *Let $\sigma \in \overline{\mathcal{M}}$ and $e, h \in \overline{H_\sigma}$. Then there exists some agent $i \in \overline{I_\sigma}$ such that $\varphi[\succ](i) = e$ for all $\succ \in \mathbf{P}[\sigma, e, h]$.*

Proof of Lemma 9. 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 preferences of a single agent $j \in \overline{I_\sigma}$. Let $g = \varphi[\succ](j)$. By strategy-proofness for j, we have $j \neq i$ and $g \neq e$. Moreover, by Maskin monotonicity, if it were true that $g = h$, then $\varphi[\succ'] = \varphi[\succ]$ would be true, contradicting that $\varphi[\succ'] \neq \varphi[\succ]$. Thus, $g \neq h$. We may further assume that

$$\succ_i \in \langle e, h, g, ... \rangle \subseteq \mathbf{P}_i[\sigma, e, h],$$

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

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

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

Let $i' \in \overline{I_\sigma}$ be the agent who gets e at $\succ'$, and $k \in \overline{I_\sigma}$ be the agent who gets h at $\succ$. Notice that such agents exist because of Pareto efficiency. Because neither i nor j gets e at

$\succ'$, we have $i' \neq i, j$. Furthermore, we saw above that j does not get h at $\succ$, and Lemma 5 implies that neither i nor i' gets h at $\succ$. Thus $k \neq i, i', j$.

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

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

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

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

Proof of Claim 1. The first five equalities were proved or assumed above and are listed for convenience only. The last two equalities require an argument. First, let us show that $\varphi[\succ'](i) = g$. Since agent j envies i at $\succ$ and $\varphi[\succ](j) = g$, Lemma 7 implies that i gets at least g at $\succ' = \left(\succ_{-j}, \succ'_j\right)$. 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, let us show that $\varphi[\succ'](k) = h$. Consider an auxiliary preference ranking $\tilde{\succ}_k \in \langle e, h, g, ...\rangle$ that agrees with $\succ_k$ except possibly for the relative ranking of g. Maskin monotonicity implies that

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

Thus, agent j envies k at $(\tilde{\succ}_k, \succ_{-k})$ and $\varphi[\tilde{\succ}_k, \succ_{-k}](j) = g$, and thus Lemma 7 implies that $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j](k) \succeq_k g$. Strategy-proofness for k implies that k cannot get e at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$. To prove that k gets h it is thus enough to show that i gets g at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$. The proof is analogous to the proof of the equality $\varphi[\succ'](i) = g$ above: i gets at least g at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$, and because j cannot envy i at $\left(\tilde{\succ}_k, \succ_{-k,j}, \succ'_j\right)$ (by Lemma 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 it must be that $\varphi[\tilde{\succ}_k, \succ_{-k,j}, \succ'_j] = \varphi[\succ_k, \succ_{-k,j}, \succ'_j] = \varphi[\succ']$. Thus, $\varphi[\succ'](k) = h$, and the claim is proved. QED

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

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

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

$$\succ_i^* \in \langle h, e, g, ...\rangle \,,$$
$$\succ_{i'}^* \in \langle h, e, ...\rangle \,, \text{ and}$$
$$\succ_k^* \in \langle e, g, h, ...\rangle \,.$$

We will prove a number of claims.

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

Proof of Claim 2.
(1) By strategy-proofness for i, $\varphi[\succ_i^*, \succ_{-i}](i) \succeq_i^* e$. Everybody else in $\overline{I_\sigma}$ ranks e over h. Thus, by Lemma 8 and Pareto efficiency, i should get h at $[\succ_i^*, \succ_{-i}]$. The symmetric argument yields $\varphi[\succ_{i'}^*, \succ'_{-i'}](i') = h$.

(2) By Maskin monotonicity for i, $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ']$.[41] Thus, j gets g' at $[\succ_i^*, \succ'_{-i}]$. By strategy-proofness for j, agent j gets at least g' and no house better than g at $[\succ_i^*, \succ_{-i}]$ (recall that between $\succ_{-i}$ and $\succ'_{-i}$ only j changes preferences). Thus, in order to prove the claim that j gets g at $[\succ_i^*, \succ_{-i}]$ it is enough to show that he cannot get g' at $[\succ_i^*, \succ_{-i}]$. Assume to the contrary that j gets g' at $[\succ_i^*, \succ_{-i}]$. Then, non-bossiness would imply that i gets h at $[\succ_i^*, \succ'_{-i}]$. By strategy-proofness for i, he gets at least h at $\succ'$. But then j envies i both at $\succ$ and $\succ' = (\succ'_j, \succ_{-j})$ and by Lemma 5, i must get the same house at these two profiles. This contradiction proves Claim 2. QED

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

Proof of Claim 3. (1) Because k gets h at $\succ$, strategy-proofness implies that k cannot get e and gets at least h at $[\succ_k^*, \succ_{-k}]$. Thus, k gets h or g at $[\succ_k^*, \succ_{-k}]$. Everybody else in $\overline{I_\sigma}$ ranks h over g. Thus, by Lemma 8 and Pareto efficiency, agent k should get g at $[\succ_k^*, \succ_{-k}]$.

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

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

[41] We use this short-hand terminology often: When we say "by Maskin monotonicity for an agent i, $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ']$" we mean that as $\varphi[\succ'](i) \succ'_i a \implies \varphi[\succ'](i) \succ_i^* a$ for all houses a, $[\succ_i^*, \succ'_{-i}]$ is a φ-monotonic transformation of $\succ'$, and by Maskin monotonicity of φ, $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ']$.

Proof of Claim 4. Because agent k envies agent i at $\succ$, Lemma 7 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 7 implies that i' gets at least $h = \varphi[\succ'](k)$ at $[\succ_k^*, \succ'_{-k}]$. Hence $\varphi[\succ_k^*, \succ'_{-k}](i') \in \{e, h\}$. By Claim 3(2), $\varphi[\succ_k^*, \succ_{-k}](i') \in \{e, h\}$. Thus,

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

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

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

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

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

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

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

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

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

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

Proof of Claim 7. Let us first prove that $\varphi[\succ_{\{i,i'\}}^*, \succ'_{-\{i,i'\}}](i) \neq h$. Suppose not. Then, Maskin monotonicity for i' gives $\varphi[\succ_i^*, \succ'_{-i}] = \varphi[\succ_{\{i,i'\}}^*, \succ'_{-\{i,i'\}}]$, and in particular, $\varphi[\succ_i^*, \succ'_{-i}$

$](i) = h$. By strategy-proofness for i, $\varphi[\succ'](i) \succeq_i h$, contradicting that $\varphi[\succ'](i') = e$ and $\varphi[\succ'](k) = h$, and proving the required inequality.

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

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

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

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

Contrary to the claim we are proving, suppose that $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) = e$. Then, the first of the above-displayed equalities implies $\varphi[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}](i) = e$ and, hence, j envies i under $[\succ^*_{\{i,i'\}}, \succ'_{-\{i,i'\}}] = [\succ^*_{\{i,i'\}}, \succ_{-\{i,i',j\}}, \succ'_j]$. This, however, leads to a contradiction with Lemma 5, because Claim 2 and the second above-displayed equality implies that $\varphi[\succ^*_{\{i,i'\}}, \succ_{-\{i,i',j\}}, \succ_j](i) = h$. Thus, we have shown that $\varphi[\succ^*_{i'}, \succ'_{-i'}](i) \neq e$. QED

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

Lemma 10. ***(Existence and uniqueness of a broker* for each brokered* house)*** *Let $\sigma \in \overline{\mathcal{M}}$ and e be a brokered* house at σ. Then there exists an agent $k \in \overline{I_\sigma}$ who is the unique broker* of e at σ.*

Proof of Lemma 10. Let $\sigma \in \overline{\mathcal{M}}$ and e be a brokered* house at σ. We start with the following preparatory claim:

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

Proof of Claim 1. By Lemma 9, $\varphi[\succ']^{-1}(e) = \varphi[\succ]^{-1}(e)$. 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 under $\succ^*$, $\succ$ and $\succ'^*$, $\succ'$ are respectively the same. By Maskin monotonicity, $\varphi[\succ^{*\prime}] = \varphi[\succ']$ and $\varphi[\succ^*] = \varphi[\succ]$. Also $\succ^*, \succ'^* \in \mathbf{P}[\sigma \cup \{(i,e)\}, h, h']$. Thus, by Lemma 9, $\varphi[\succ^*]^{-1}(h) = \varphi[\succ^{*\prime}]^{-1}(h)$. Hence, $\varphi[\succ']^{-1}(h) = \varphi[\succ'^*]^{-1}(h) = \varphi[\succ^*]^{-1}(h) = \varphi[\succ]^{-1}(h)$. QED

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

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

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

In particular, $\varphi[\succ^*_{k'}, \succ_{-k'}](k') = h'$. By strategy-proofness for k', we have $\varphi[\succ](k') \in \{h, h'\}$. On the other hand, by Lemma 9,

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

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

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

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

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

By Lemma 9, $h'' \neq h$. By the same lemma, we assume that $\succ'' \in \mathbf{P}[\sigma, e, h'', h]$.

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

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

Proof of Claim 4. By Lemma 9, $\varphi[\succ]^{-1}(e) = \varphi[\succ']^{-1}(e)$. Because e is brokered* at σ, there exists some $h'' \in \overline{H_\sigma} - \{e\}$ such that for some $\succ'' \in \mathbf{P}[\sigma, e, h'']$,

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

By Lemma 9, $h'' \neq h$ and, by the same lemma, we may assume $\succ'' \in \mathbf{P}[\sigma, e, h'', h]$. By Claim 3, $\varphi[\succ'']^{-1}(h'') = \varphi[\succ]^{-1}(h)$ and $\varphi[\succ'']^{-1}(h'') = \varphi[\succ']^{-1}(h)$, implying that $\varphi[\succ]^{-1}(h) = \varphi[\succ']^{-1}(h)$. QED

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

Lemma 11. *Let $\sigma \in \overline{\mathcal{M}}$, $i \in \overline{I_\sigma}$, and $h \in \overline{H_\sigma}$. If $\varphi[\succ](i) = h$ for all $\succ \in \mathbf{P}^*[\sigma, h]$, then i owns* h at σ.*

Proof of Lemma 11. Let us start with two preparatory claims:

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

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

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

Let $\succ_i^* \in \langle g, h, ...\rangle$. By strategy-proofness, agent i gets at least h at $[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}]$; and by Pareto efficiency, agent i gets g. Also recall that $\varphi[\succ](i) \prec_i g$ and $\varphi[\succ](k) = g$. Thus, $\varphi[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}](k) \neq h$ because otherwise agents i and k could jointly improve upon their $\varphi[\succ]$ allocation by submitting $[\succ_i^*, \succ'_k]$ at $\succ$, contradicting group strategy-proofness. Thus, $g \succ'_k \varphi[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}](k)$, and furthermore, Maskin monotonicity implies $\varphi[\succ_i^*, \succ'_k, \succ_{-\{i,k\}}] = \varphi[\succ_i^*, \succ_{-i}]$. In particular, $\varphi[\succ_i^*, \succ_{-i}](i) = g$. QED

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

Proof of Claim 2. Assume that $\varphi[\succ'']^{-1}(h) = i \in \overline{I_\sigma}$ for all $\succ' \in \mathbf{P}[\sigma, h, g]$. By way of contradiction suppose $\succ \in \mathbf{P}[\sigma, h]$ and there is some $\succ' \in \mathbf{P}[\sigma, h, g]$ such that $\succ_k \in \langle h, g, ...\rangle$

for $k = \varphi[\succ']^{-1}(g)$, and yet $\varphi[\succ](i) \neq h$. By strategy-proofness, we can choose $\succ_i \in \langle h, g, ...\rangle$. Furthermore, we can choose $\succ$ such that $\succ$ and $\succ'$ differ only in the preferences of a single agent $j \in \overline{I_\sigma}$ and in how house g is ranked by the agents.

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

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

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

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

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

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

Furthermore, we will show that i gets h at any profile $\succ' \in \mathbf{P}[\sigma, h]$ at which k_2 ranks g_1 second. Indeed, suppose $\succ'_{k_2} \in \langle h, g_1, ...\rangle$ and i does not get h at $\succ'$. Let $\succ''_i \in \langle h, g_1, ...\rangle$. By Claim 1 and strategy-proofness, agent i gets g_1 at $[\succ''_i, \succ'_{-i}]$. By the previous paragraph and strategy-proofness, k_2 does not get h at $[\succ''_i, \succ'_{-i}]$, and thus k_2 envies i at $[\succ''_i, \succ'_{-i}]$. However, by the previous paragraph k_2 can improve the outcome of agent i, contrary to Lemma 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 that gets g_3 at $\succ^{g_3}$. As above, we can show that i gets h at any profile in $\mathbf{P}[\sigma, h]$ at which k_3 ranks g_3 or g_2 or g_1

second.

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

Lemmas 10 and 11 and the definitions of owned* and brokered* houses give us the key result of this subsection:

Proposition 3. ***(Houses are either brokered* or owned*)*** *For any $\sigma \in \overline{\mathcal{M}}$, any house $h \in \overline{H_\sigma}$ is owned* or brokered* at σ, but not both. In particular, if there is only one agent who owns* houses at σ then h is not a brokered* house.*

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

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

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

Lemma 12. ***(Relationship between brokerage* and ownership*).*** *Let $\sigma \in \overline{\mathcal{M}}$, agent k be a broker* of house e at σ, and $\succ'' \in \mathbf{P}^*[\sigma, e]$. Then agent $\varphi[\succ'']^{-1}(e)$ is the owner* of house $\varphi[\succ''](k)$ at σ.*

Proof of Lemma 12. Let $\succ'' \in \mathbf{P}^*[\sigma, e]$ and $h = \varphi[\succ''](k)$. Because k is a broker* at σ, Lemma 10 implies that house h is agent k's second choice. Since $\succ'' \in \mathbf{P}^*[\sigma, e]$, house h is the second choice of all agents in $\overline{I_\sigma}$ at $\succ''$, and thus,

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

There exists an agent $i \in (\overline{I_\sigma}) - \{k\}$ such that $\varphi[\succ'']^{-1}(e) = i$. By Lemma 9, for all $\succ \in \mathbf{P}[\sigma, e, h]$, agent i gets e at $\succ$. We are to show that i is the owner* of h at σ.

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

Proof of Claim 1. The first claim follows from Lemma 9, and the second from Lemma 10. QED

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

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

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

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

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

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

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

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

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

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

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

Proof of Claim 6. For an indirect argument, suppose that $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](i) = h$ and $\varphi[\succ^*_{\{i,k\}}, \succ_{-\{i,k\}}](k) = e$. Then, $\varphi[\succ^*_i, \succ_{-i}](k) = e$ by the strategy-proofness of φ. Since e is a brokered* house at σ, there exists some house $g \notin \{e, h\}$ and some preference profile

$\succ' \in \mathbf{P}[\sigma, e, g]$ such that $\varphi[\succ']^{-1}(e) = j$ for some agent $j \notin \{i, k\}$. By Lemma 9, we may assume that each agent $i' \in \overline{I_\sigma}$ ranks houses other than g and h in the same way at $\succ'_{i'}$ and $\succ_{i'}$ and that $\succ'_{i'} \in \langle e, g, h, ...\rangle$. Since k is the broker* of e at σ, we have $\varphi[\succ'](k) = g$. By Maskin monotonicity,

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

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

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

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

Lemma 13. ***(R1; Uniqueness of a brokered* house)***. *Let $\sigma \in \overline{\mathcal{M}}$. If e is a brokered* house at σ, then no other house is a brokered* house at σ (and all other unmatched houses are owned* houses).*

Proof of Lemma 13. Let e be a brokered* house at σ. By Lemma 10, there is a broker* of e at σ; let us denote him as k. Consider a house $h \in \overline{I_\sigma} - \{e\}$. By Lemma 9, there is an agent i who gets e at all profiles in $\mathbf{P}[\sigma, e, h]$. By Lemma 11, i is the owner* of h. Thus h is not a brokered* house at σ. **QED**

Lemma 14. ***(R2; Last unmatched agent is an owner).*** *Let $\sigma \in \overline{\mathcal{M}}$, such that there exists a unique agent i unmatched at σ. Then i owns* all unmatched houses at $\sigma \in \overline{I_\sigma}$.*

Proof of Lemma 14. Let $\succ \in \mathbf{P}[\sigma, h]$ for $h \in \overline{H_\sigma}$. By Pareto efficiency of φ, $\varphi[\succ](i) = h$, implying that i owns* h at σ. **QED**

Lemma 15. ***(R3; Broker* does not own*)***. *Let $\sigma \in \overline{\mathcal{M}}$. If agent k is the broker* of house e at σ, then he cannot own* any houses at σ.*

Proof of Lemma 15. Suppose that k owns* a house $h \neq e$ at σ. By Lemma 9, there exists some agent $i \neq k$ who gets e at all profiles in $\mathbf{P}[\sigma, e, h]$. Thus, i gets h at all $\succ \in \mathbf{P}^*[\sigma, h]$, contradicting the fact that k owns* h. **QED**

Lemma 16. ***(R4; Persistence of ownership*).*** *Let i own* h at some $\sigma \in \overline{\mathcal{M}}$. If $\sigma' \supsetneq \sigma$, and i and h are unmatched at σ', then i owns* h at σ'.*

Proof of Lemma 16. Imagine to the contrary that i gets h at all $\succ \in \mathbf{P}[\sigma, h]$, but there is some $\succ' \in \mathbf{P}[\sigma', h]$ such that some agent $j \in I_{\sigma'} - I_\sigma$, such that $j \neq i$, gets h at $\succ'$. Take $\succ \in \mathbf{P}[\sigma, h]$ such that

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

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

- at $\succ'$ agent k gets $\sigma'(k)$ by Lemma 8,
- at $\succ$ agent k gets $\sigma'(k)$ by the Pareto efficiency of φ and the fact that $\varphi[\succ](i) = h$.

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

Lemma 17. ***(R5; Limited persistence of brokerage*)*** *Let $\sigma, \sigma' \in \overline{\mathcal{M}}$ be such that $\sigma' \supsetneq \sigma$. Suppose that agent k is the broker* of house e at σ, agent i is the owner* of house h at σ, and agent $i' \neq i$ is the owner* of house h' at σ. If k, i, i', e, h, h' are unmatched at σ', then k brokers* e at σ'.*

Proof of Lemma 17. First, notice that Lemma 12 implies that i gets e at all $\succ \in \mathbf{P}[\sigma, e, h]$ and i' gets e at all $\succ \in \mathbf{P}[\sigma, e, h']$, and that k gets h and h', at the respective profiles. Take $\succ^h \in \mathbf{P}[\sigma, e, h]$ and $\succ^{h'} \in \mathbf{P}[\sigma, e, h']$ such that each agent $j \in I_{\sigma'} - I_\sigma$ has $\sigma'(j)$ as his third choice and each agent $j \in I - I_{\sigma'}$ ranks each house unmatched at σ' above all houses matched at σ' at both preference profiles. Let profile $\succ'^h$ be obtained from $\succ^h$ by moving $\sigma'(j)$ for all $j \in I_{\sigma'} - I_\sigma$ up to be the first choice of j. Let $\succ'^{h'}$be obtained analogously from $\succ^{h'}$. By Maskin monotonicity, $\varphi[\succ'^h]^{-1}(e) = i \neq i' = \varphi[\succ'^{h'}]^{-1}(e)$. Since $\succ'^h$ and $\succ'^{h'} \in \mathbf{P}^*[\sigma', e]$, house e is a brokered* house at σ'.

For an indirect argument for the second part of the proof, suppose that k is not the broker* of e at σ'. Then, by Lemma 10 there exists some other agent $k' \neq k$ who brokers* e at σ'.

Let $\succ' \in \mathbf{P}[\sigma', e, h]$ be arbitrary and $\succ \in \mathbf{P}[\sigma, e, h]$ be such that each agent j in $I_{\sigma'} - I_\sigma$ lists $\sigma'(j)$ as his third choice at $\succ$, each agent in $I - I_{\sigma'}$ lists houses in $H_{\sigma'}$ lower than houses in $H_{\sigma'} - H_\sigma$ at $\succ$, and the rest of the relative rankings of the houses are the same between $\succ$ and $\succ'$. Since k brokers* e at σ and i owns* h at σ, by Lemma 12 $\varphi[\succ](k) = g$ and $\varphi[\succ](i) = e$. Then, by Pareto efficiency, $\varphi[\succ'](j) = \sigma'(j)$ for all $j \in I_{\sigma'} - I_\sigma$, and thus, by Maskin monotonicity, $\varphi[\succ'] = \varphi[\succ]$. Now, $\varphi[\succ'](k) = h$, however, this contradicts the fact that agent $k' \neq k$ brokers* e at σ' and thus, $\varphi[\succ'](k') = h$. Therefore, k brokers* e at σ', as well. **QED**

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

Proof of Lemma 18. First consider the case i brokers* h and j owns* g at σ. By Lemmas 11 and 12 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle h, g, ... \rangle$, $\succ_j \in \langle h, ... \rangle$, we have $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Then, by Maskin monotonicity, for any $\succ' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ](i) = g$, i.e., i owns* g at σ'.
Next consider the case i owns* h and j owns* g at σ. For all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle g, h ... \rangle$ and $\succ_j \in \langle h, g, ... \rangle$, strategy-proofness for i implies $\varphi[\succ](i) \succeq_i h$ as $\varphi[\succ'_i, \succ_{-i}](i) = h$ for $\succ'_i \in \langle h, ... \rangle$. Similarly, $\varphi[\succ](j) \succeq_j g$. Pareto efficiency implies $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Hence, by Maskin monotonicity, for all $\succ'' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ''](i) = g$, i.e., i owns* g at σ'.
Finally consider the case where i owns* h and j brokers* g at σ. By Lemmas 11 and 12 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle g, ... \rangle$ and $\succ_j \in \langle g, h ... \rangle$, we have $\varphi[\succ](i) = g$ and $\varphi[\succ](j) = h$. Then, by Maskin monotonicity, for any $\succ' \in \mathbf{P}[\sigma', g]$, $\varphi[\succ](i) = g$, i.e., i owns* g at σ'. **QED**

Lemma 19. ***(R6; Brokered*-to-Owned* House Transition)*** *Let $\sigma \in \overline{\mathcal{M}}$, $k, j, i \in \overline{I_\sigma}$, and $e, g, h \in \overline{H_\sigma}$ be such that $k \neq j$ and $e \neq g$, k brokers* e at σ but not at $\sigma' = \sigma \cup \{(j, g)\}$, and i owns* h at σ. Then i owns* e at σ'.*

Proof of Lemma 19. By Lemmas 11 and 12 and Maskin monotonicity, for all profiles $\succ \in \mathbf{P}[\sigma]$ such that $\succ_i \in \langle e, ... \rangle$ and $\succ_k \in \langle e, h, ... \rangle$, we have $\varphi[\succ](i) = e$ and $\varphi[\succ](k) = h$. Since $\mathbf{P}[\sigma'] \subset \mathbf{P}[\sigma]$, Proposition 3 implies that either i owns* e at σ' or k brokers* e at σ'. The latter is not true, by an assumption made in the lemma; hence i owns* e at σ'. **QED**

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

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

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

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

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

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

We will show that

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

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

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

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

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

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

By Maskin monotonicity,

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

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

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

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

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

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

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

QED

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

Proof of Claim 2. The inductive assumption implies that all houses better than $\psi^{c,b}[\succ^*](i)$ are already given to other agents; hence

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

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

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

Let

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

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

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

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

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

As we continue iteratively, we obtain that

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

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

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

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

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

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

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

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

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

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

Let us construct an auxiliary preference profile $\succ' \in \mathbf{P}[\sigma^{r-1}]$ from $\succ^*$ by pushing up $\sigma^{r-1}(i)$ in preferences of agents $i \in I_{\sigma^{r-1}}$, pushing down houses matched at σ^{r-1} in preferences of agents

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

Step 2. Let us show that

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

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

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

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

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

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

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

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

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

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

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

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

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

In particular,

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

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

G Appendix: Proof of Theorem 3

The equivalence of the three pairs of incentive and efficiency conditions follows the same steps as in Theorem 1. Also, the argument for the Pareto efficiency of TC follows the same steps as in the case without outside options. As before, group strategy-proofness is equivalent to individual strategy-proofness and non-bossiness.

Lemma 20. *In the environment with outside options, a mechanism is group strategy-proof if and only if it is individually strategy-proof and non-bossy.*

The proof is analogous to the proof of Lemma 1 in Pápai (2000).

Our arguments for individual strategy-proofness and non-bossiness go through with two modifications. First, when in the proof of Theorem 2 we assume that an agent is matched with a house, we should now substitute "matched with a house or the agent's outside option." If the agent is matched in a cycle of a length above 1, we can then conclude that the agent is indeed matched with a house. Second, in some steps of the proof we consider separately the case when a broker is matched with his outside option. We handle these cases below. This allows us to assume this case away in the relevant parts of the original proof.

Consider the proof of individual strategy-proofness. In Case 1: $s \leq s'$, let i be a broker of house e and suppose he leaves with his outside option in round s under $\succ_i$. Since the same houses are matched under $\succ_i$ and $\succ_i'$, under $\succ_i'$ the best the broker can do is to leave either with his outside option, or—if he prefers the brokered house e to his outside option—to leave with the brokered house e. We need to prove that the latter cannot happen. By Lemma 4, in round s of TC under $\succ_i'$, agent i is a broker of e and there is an owner j whose first preference is e. For i to be matched with e, he would need to lose the brokerage right, but by R5-R6 if this happens then j becomes the owner of e, and is then matched with it, ending the argument for Case 1. In Case 2: $s > s'$, if i be a broker of house e matched with his outside option under $\succ_i'$, then submitting this preference profile cannot be better than submitting the true profile $\succ_i$, as under any profile agent i is matched with at least his outside option.

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

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

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

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

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

$$h^1 \succeq_i h^2 \succeq ... \succeq_i h^n \succ_i y_i \succ g \text{ for all } g \in H_\sigma.$$

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

H Proof of Proposition 2 (Lorenz Dominance)

Let ψ^{c^*,b^*} be the TC mechanism in which at the empty submatching i_ℓ owns h_ℓ for all $\ell \neq n$ and i_n brokers h_n (the brokerage right is permanent). Let $\psi^{c,b}$ be any TTC mechanism satisfying the manager's constraints, which implies that i_n does not own any objects as long as other agents are present and h_n is unallocated. Let $\overline{\succ}$ be the uniform lottery over profiles of preferences in which all agents rank objects in the same way, and for any mechanism ϕ let $\rho^{\phi[\overline{\succ}]}(i,r)$ be the average of $\rho^{\phi[\succ]}(i,r)$ over the distribution $\overline{\succ}$.

Since all agents share the same ranking under every realization of $\overline{\succ}$, in every Pareto efficient mechanism ϕ we have

$$\rho^{\phi(\overline{\succ})}(i_1,r) + ... + \rho^{\phi(\overline{\succ})}(i_n,r) = r. \tag{3}$$

By symmetry among agents $i_1, ..., i_{n-1}$, the value $\rho^{\psi^{c^*,b^*}[\overline{\succ}]}(i_k,r)$ is the same for all of them and for all $r < n$. Moreover, this common value is larger than $\rho^{\psi^{c^*,b^*}[\overline{\succ}]}(i_n,r)$. Thus,

$$\min_{J\subseteq I,|J|=k} \sum_{j\in J} \rho^{\psi^{c^*,b^*}[\overline{\succ}]}(j,r) = \frac{k-1}{n-1}\left(r - \rho^{\psi^{c^*,b^*}[\overline{\succ}]}(i_n,r)\right) + \rho^{\psi^{c^*,b^*}[\overline{\succ}]}(i_n,r).$$

Let us denote this value by $X^*(k)$.

Now, consider agent i_n and $k-1$ agents $i_1, ..., i_{n-1}$ with the lowest $\rho^{\psi^{c,b}}(\cdot,r)$. By (3), the sum of $\rho^{\psi^{c,b}}(\cdot,r)$ over these k agents is at most

$$\min_{J\subseteq I-\{i_n\},|J|=k-1} \sum_{j\in J} \rho^{\psi^{c^*,b^*}[\overline{\succ}]}(j,r) + \rho^{\psi^{c,b}[\overline{\succ}]}(i_n,r)$$

Let us denote this value by $X(k)$.

To show that $X^*(k) \geq X(k)$ for all k notice that

$$X(k) \leq \frac{k-1}{n-1}\left(r - \rho^{\psi^{c,b}[\overline{\succ}]}(i_n,r)\right) + \rho^{\psi^{c,b}[\overline{\succ}]}(i_n,r).$$

It is thus enough to show that $\rho^{\psi^{c^*,b^*}[\overline{\succ}]}(i_n,r) \geq \rho^{\psi^{c,b}[\overline{\succ}]}(i_n,r)$. This inequality obtains because (i) in a TTC satisfying the manager's constraints agent i_n cannot be allocated an object before another agent takes h_n or before i_n is the only unmatched agent remaining, and (ii) with any common ranking any TTC allocates one object per round, from the most preferred object down everybody's list, and ψ^{c^*,b^*} allocates objects in the same way in rounds preceding the allocation of h_n, and then in the round in which h_n is allocated ψ^{c^*,b^*} allocates to i_n his most preferred object among objects that remain.

To close the proof notice that for at least one k we must have $X^*(k) > X(k)$. Indeed, if $X^*(1) = X(1)$ then i_n has the same $\rho^{\cdot}(i_n, r)$ at both mechanisms. Then, however, $X^*(k) > X(k)$ for all $k = 2, ..., n-1$. Indeed, in $\psi^{c,b}$ there is at least one agent who owns at least two objects at the empty submatching. Such an agent i has $\rho^{\psi^{c,b}[\succ]}(i, r)$ strictly larger than the common value of $\rho^{\psi^{c^*,b^*}[\succ]}(i_\ell, r)$ for $\ell = 1, ..., n-1$. By (3), the sum of the relevant $\rho^{\cdot}(\cdot, r)$ for the remaining $n-2$ agents is thus strictly less under $\psi^{c,b}$ than under ψ^{c^*,b^*}. **QED**

References

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

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

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

Arrow, K. J. (1963): *Social Choice and Individual Values*. New York: Wiley, 2nd edition edn.

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

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

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

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

Bergson, A. (1938): "A Reformulation of Certain Aspects of Welfare Economics," *Quarterly Journal of Economics*, 52(2), 310–334.

Blair, D., and E. Muller (1983): "Essential Aggregation Procedures on Restricted Domains of Preferences," *Journal of Economic Theory*, 30, 34–53.

Bogolomania, A., and H. Moulin (2004): "Random Matching Under Dichotomous Preferences," *Econometrica*, 72, 257–279.

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

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

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

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

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

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

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

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

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

Holmstrom, B. (1979): "Groves' Scheme on Restricted Domains," *Econometrica,*, 47(5), 1137—1144.

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

Konishi, H., T. Quint, and J. Wako (2001): "On the Shapley-Scarf Economy: The Case of Multiple Types of Indivisible Goods," *Journal of Mathematical Economics*, 35, 1–15.

Lorenz, M. O. (1905): "Methods of Measuring the Concentration of Wealth," *Publications of the American Statistical Association*, 9(70), 209—219.

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

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

Moulin, H. (1988): *Axioms of Cooperative Decision Making*. Cambridge: Cambridge University Press.

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

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

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

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

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

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

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

Samuelson, P. A. (1947): *Foundations of Economic Analysis*. MA: Harvard University Press.

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

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

Sen, A. (1970): "Interpersonal Aggregation and Partial Comparability," *Econometrica*, 38(3), 393–409.

——— (1999): "The Possibility of Social Choice," *American Economic Review*, 89(3), 349–378.

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

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

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

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

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

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

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

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

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