

A Simple Family of Top Trading Cycles Mechanisms for Housing Markets with Indifferences

C. Gregory Plaxton *

April 13, 2013

Abstract

Recently, two new families of strategyproof mechanisms have been presented for housing markets with indifferences. These mechanisms produce a Pareto-efficient core allocation, and produce a strict core allocation when the strict core is nonempty. We propose a novel family of mechanisms and prove that this family achieves the same combination of properties. Our family of mechanisms is based on a generalization of the top trading cycles algorithm. We establish a confluence property of our algorithm, and use this property to establish that the associated mechanisms are strategyproof. We also provide a simple $O(n^3)$ -time deterministic implementation of our family of mechanisms, where n denotes the number of agents.

*Department of Computer Science, University of Texas at Austin, 2317 Speedway, Stop D9500, Austin, Texas 78712-1757. Email: plaxton@cs.utexas.edu. This research was supported by NSF Grant CCF-1217980.

1 Introduction

In a seminal paper, Shapley and Scarf [11] study a simple housing market involving n agents. Each agent owns a house, and has a preference order over the n houses. The goal is to determine a suitable allocation of houses to agents, with the understanding that no monetary transfers are allowed. For this setting, Shapley and Scarf [11] present the *top trading cycles (TTC)* algorithm, which they attribute to David Gale. Under weak preferences (i.e., the natural setting in which indifferences are allowed), the TTC algorithm can produce different allocations. Shapley and Scarf show that the set of TTC allocations coincides with the set of competitive allocations, that the set of competitive allocations is contained in the core, that the latter containment can be proper (even when preferences are strict), and that the strict core can be empty.

When preferences are strict, the housing market problem is well understood. Roth and Postlewaite [10] show that the strict core corresponds to the unique competitive allocation. Roth [9] shows that the mechanism defined by the TTC algorithm is strategyproof. Bird [4] shows that the TTC mechanism is group strategyproof. Ma [7] shows that the TTC mechanism is the only individually rational, Pareto-efficient, and strategyproof mechanism.

For weak preferences, the housing market problem is substantially more challenging. Wako [12] shows that the strict core is contained in the set of competitive allocations, and that this containment can be proper. Quint and Wako [8] characterize the class of instances for which the strict core is nonempty, and provide an $O(n^3)$ -time algorithm to obtain a strict core allocation on such instances. Recent independent work of Alcalde-Unzu and Molis [2], and of Jaramillo and Manjunath [6], provides new mechanisms for housing markets with weak preferences. These results are highly relevant to the present paper, and are discussed in greater detail below.

Alcalde-Unzu and Molis [2] present an algorithm called Top Trading Absorbing Sets (TTAS). This algorithm yields a family of strategyproof mechanisms called the TTAS mechanisms. The mechanisms in this family produce a Pareto-efficient core allocation, and produce a strict core allocation when the strict core is nonempty. The main shortcoming of the TTAS algorithm is its high time complexity, which arises because TTAS can trade along many “bad” cycles, i.e., cycles where each of the associated agents is already tentatively assigned to a house in its preferred set (amongst the remaining houses). Alcalde-Unzu and Molis leave open the question of whether the TTAS algorithm runs in polynomial time. Aziz and de Keijzer [3] answer this question in the negative by exhibiting a family of instances on which the TTAS algorithm runs in exponential time.

Jaramillo and Manjunath [6] present an algorithm called Top Cycles Rules (TCR). This algorithm yields the family of TCR mechanisms. Jaramillo and Manjunath prove that the TCR mechanisms are individually rational, Pareto-efficient, and strategyproof. Aziz and de Keijzer [3] prove that, like the TTAS family, each mechanism in the TCR family produces a core allocation, and produces a strict core allocation when the strict core is nonempty. Thus the TCR mechanisms match the properties mentioned above for the TTAS mechanisms. Furthermore, the TCR algorithm runs in polynomial time. (More specifically, Jaramillo and Manjunath establish an $O(n^6)$ bound on the time complexity of the TCR algorithm.) Like the TTAS algorithm, TCR proceeds iteratively. In each iteration, each remaining agent has a preferred set of houses amongst the remaining houses. Each agent selects a specific house from its preferred set. The cycle trading phase of an iteration then updates the tentative allocation in the usual TTC style, treating the selected houses as the unique top choices. Within this framework, the central question is how an agent should select

a specific house from its preferred set.

In the TTAS algorithm, each agent uses a local rule to select a specific house in each iteration. Due to the lack of global coordination, many bad cycles may occur. The TCR algorithm uses an entirely different method to perform the selection. The TCR method (see the description of the “pointing phase” in [6]) is somewhat involved; here we mention only two aspects of this method: (1) in order to achieve a certain “persistence” property, the TCR method takes into account the selections made in the previous iteration (i.e., it is not memoryless); (2) the TCR method involves multiple levels of tie-breakers. These aspects of the TCR method complicate the task of establishing strategyproofness.

We introduce and analyze a new, simpler method for an agent to select a specific house from its preferred set. In each iteration, we compute the shortest path distance of each house to an “unsatisfied” agent, i.e., an agent who is not tentatively assigned to a house in its preferred set (see the definition of $distance(G, v)$ in Section 2.1). Each agent selects the house in its preferred set with the smallest distance value, with ties broken according to a fixed total order over the set of houses (see the definition of $next(G, u)$ in Section 2.1). In contrast with the TCR method, our method is memoryless, and uses a single-level tie-breaking scheme. Our analogue of the persistence property of Jaramillo and Manjunath [6] is established as a consequence of the selection method (see Lemmas 3.4, 3.9, and 4.1), and not by complicating the selection method.

In Section 2 we define a nondeterministic algorithm, Algorithm 1. We prove that Algorithm 1 is confluent (see Lemma 4.5), implying that the output does not depend on the nondeterministic choices made during execution (see Theorem 1). In Section 3, we establish properties of the special class of bipartite digraphs that arise in modeling housing markets with indifferences; we use the term “configuration” (formally defined in Section 2.1) to refer to such a structure. In Section 4, we establish properties of the agent preferences, independent of the initial endowments; we use the term “wpp” (formally defined in Section 2.2) to refer to the preferences-related component of a problem instance. In Section 5 we exploit confluence to give a short proof that the family of mechanisms associated with Algorithm 1 is strategyproof (see Lemma 5.1 and Theorem 2). In Appendix F, we again exploit confluence, this time to obtain an $O(n^3)$ -time deterministic algorithm, Algorithm 2, with the same input-output behavior as Algorithm 1.

Aziz and de Keijzer [3] introduce a class of mechanisms called Generalized Absorbing Top Trading Cycle (GATTC), which is designed for housing markets with indifferences, and which includes the TTAS and TCR mechanisms. The GATTC family is quite broad, and includes mechanisms that are not strategyproof [3]. Aziz and de Keijzer [3] prove that every mechanism in this family produces a Pareto-efficient core allocation, and produces a strict core allocation when the strict core is nonempty. In Appendix E, we prove that the family of mechanisms associated with Algorithm 1 is contained in the GATTC family; this proof also exploits the confluence property of Algorithm 1. The primary focus of the remainder of the paper is to establish strategyproofness. (Remark: Our results do not strongly depend upon the work of Aziz and de Keijzer in the sense that it is straightforward to establish from first principles that our mechanisms produce a Pareto-efficient core allocation, and it is also straightforward to use the aforementioned characterization of Quint and Wako [8] to establish that our mechanisms produce a strict core allocation when the strict core is nonempty. The primary challenge is to show that the family of mechanisms defined by Algorithm 1 is strategyproof.)

Significance of our results. Housing markets are a fundamental model for the exchange of indivisible goods. The complications associated with accommodating weak preferences are already

recognized in Shapley and Scarf’s original paper [11, Section 5]. It is natural to ask whether there is a strategyproof mechanism for housing markets with weak preferences that produces a Pareto-efficient core allocation, and produces a strict core allocation when the strict core is nonempty. While the recent work of Alcalde-Unzu and Molis [2] answers this question in the affirmative, it is not entirely satisfactory since, as noted above, the associated TTAS algorithm has exponential time complexity. On the other hand, as indicated earlier, the TCR algorithm of Jaramillo and Manjunath [6] achieves polynomial time complexity. Our work improves on that of Jaramillo and Manjunath in several respects. First, as discussed above, the selection method — the essential ingredient in the definition of algorithms such as TTAS, TCR, and Algorithm 1 — of Algorithm 1 is substantially simpler than that of the TCR algorithm. Second, while our proof that Algorithm 1 defines a strategyproof family of mechanisms is not simple, it has a modular structure centered around a natural confluence property; the corresponding proof given in [6] is largely monolithic in nature, and is difficult to verify. Third, as a by-product of the simplicity of the selection rule, we obtain a substantially improved polynomial time bound ($O(n^3)$ versus $O(n^6)$). As discussed in Section 6, improving our $O(n^3)$ time bound by more than a \sqrt{n} factor would imply an improved time bound for the highly-studied bipartite maximum cardinality matching problem. Thus our $O(n^3)$ time bound appears to be close to optimal.

2 A Family of Mechanisms

In Section 1 we have provided an informal description of the family of mechanisms to be analyzed in the present paper. In this section, we provide a formal description. Recall that the main technical challenge of the paper is to establish the strategyproof property of our family of mechanisms. Accordingly, the style of our description is chosen to facilitate formal reasoning.

2.1 Configurations

A *configuration* is a bipartite digraph (U, V, E) where U is a set of agents, V is a totally ordered set of houses, and the following conditions hold: each agent u in U has indegree 1; each house v in V has outdegree 1. (Thus $|U| = |V|$.)

For any configuration $G = (U, V, E)$ and any house v in V , we define $agent(G, v)$ as the unique agent u such that edge (v, u) belongs to E . For any configuration $G = (U, V, E)$ and any agent u in U , we define $house(G, u)$ as the unique house v such that $agent(G, v) = u$.

For any configuration $G = (U, V, E)$, we define $allocation(G)$ as the allocation that assigns each house v in V to $agent(G, v)$.

For any configuration $G = (U, V, E)$, and any agent u in U , we define $\Gamma(G, u)$ as $\{v \mid (u, v) \in E\}$. A configuration $G = (U, V, E)$ is *initial* if $\Gamma(G, u)$ is empty for all agents u in U .

For any configuration $G = (U, V, E)$, we define $satisfied(G)$ as the set of all agents u in U such that $house(G, u)$ belongs to $\Gamma(G, u)$, and we define $unsatisfied(G)$ as $U \setminus satisfied(G)$. A configuration G is *final* if $unsatisfied(G)$ is empty.

For any configuration $G = (U, V, E)$ and any house v in V , we define $distance(G, v)$ as the length of a shortest path from v to an agent in $unsatisfied(G)$. If there is no such path, we define $distance(G, v)$ as ∞ .

For any configuration $G = (U, V, E)$ and any agent u in U , we define $next(G, u)$ as follows: if $distance(G, v) = \infty$ for all v in $\Gamma(G, u)$, then $next(G, u) = nil$; otherwise, letting V' denote the set of all v in $\Gamma(G, u)$ minimizing $distance(G, v)$, we define $next(G, u)$ as the minimum element of V' .

For any configuration $G = (U, V, E)$, we define $pruned(G)$ as the configuration $G' = (U, V, E \setminus E')$ where E' denotes

$$\{(u, v) \in E \mid u \in U \wedge v \neq next(G, u)\},$$

and we define $cycles(G)$ as the set of all directed cycles in $pruned(G)$.

For any configuration $G = (U, V, E)$ and any cycle C in $cycles(G)$, we define $trade(G, C)$ as the configuration $(U, V, (E \setminus E') \cup E'')$ where E' denotes

$$\{(v, u) \in V \times U \mid v \in C \wedge agent(G, v) = u\}$$

and E'' denotes

$$\{(v, u) \in V \times U \mid u \in C \wedge next(G, u) = v\}.$$

For any configuration G , we define $exhausted(G)$ as the set of all agents u in $unsatisfied(G)$ such that $next(G, u) = nil$.

2.2 Preferences

A *weak preference relation* is a total preorder (also known as a weak order). We define a *weak preference profile (wpp)* as a triple (U, V, \succsim) where U is a set of agents, V is a totally ordered set of houses such that $|U| = |V|$, and \succsim is a function from U to the set of weak preference relations over V . Notation: Given a wpp (U, V, \succsim) , and an agent u in U , we write \succsim_u to refer to the weak preference relation to which u is mapped by \succsim ; we use the symbol \sim to denote indifference and the symbol \succ to denote strict preference.

For any wpp $W = (U, V, \succsim)$, we define $configs(W)$ as the set of all configurations $G = (U, V, E)$ such that for any agent u in U , any house v in $\Gamma(G, u)$, and any house v' in $V \setminus \Gamma(G, u)$, we have $v \succ_u v'$.

Lemma 2.1. *For any wpp $W = (U, V, \succsim)$, any configuration $G = (U, V, E)$ in $configs(W)$, and any cycle C in $cycles(G)$, the configuration $trade(G, C)$ belongs to $configs(W)$.*

Proof. Straightforward. □

For any wpp $W = (U, V, \succsim)$, any agent u in U , and any subset V' of V , we define $top(W, u, V')$ as the set of houses v in V' such that $v \succ_u v'$ holds for all v' in V' .

For any wpp $W = (U, V, \succsim)$, any configuration $G = (U, V, E)$ in $configs(W)$, and any agent u in $exhausted(G)$, we define $reveal(W, G, u)$ as the configuration $(U, V, E \cup E')$ where

$$E' = \{(u, v) \mid v \in top(W, u, V \setminus \Gamma(G, u))\}.$$

Lemma 2.2. *For any wpp $W = (U, V, \succsim)$, any configuration $G = (U, V, E)$ in $configs(W)$, and any agent u in $exhausted(G)$, the configuration $reveal(W, G, u)$ belongs to $configs(W)$.*

Proof. Straightforward. □

```

while  $\Gamma(W, G) \neq \emptyset$ 
     $G :=$  a nondeterministically chosen element of  $\Gamma(W, G)$ 
return  $allocation(G)$ 

```

Figure 1: We refer to the above nondeterministic algorithm as Algorithm 1. Initially, (W, G) is a housing market instance.

For any wpp $W = (U, V, \succ)$, let $moves(W)$ denote the edge-labeled digraph with vertex set $configs(W)$ and edge set determined as follows. First, for any G in $configs(W)$ and any C in $cycles(G)$, there is an edge (G, G') with label C , where $G' = trade(G, C)$ belongs to $configs(W)$ by Lemma 2.1. Second, for any G in $configs(W)$ and any agent u in $exhausted(G)$, there is an edge (G, G') with label u , where $G' = reveal(W, G, u)$ belongs to $configs(W)$ by Lemma 2.2.

For any wpp W and any G in $configs(W)$, we define $\Gamma(W, G)$ as the set of all configurations G' such that edge (G, G') belongs to $moves(W)$.

2.3 A Nondeterministic Algorithm

A *housing market instance* is a pair (W, G) where W is a wpp and G is an initial configuration in $configs(W)$.

We will first study the simple nondeterministic algorithm of Figure 1, which we refer to as Algorithm 1. We show that Algorithm 1 terminates with a configuration G that is final (Lemma 4.3). We also show that the output allocation is uniquely determined (Theorem 1). In Appendix F, we present a simple deterministic algorithm for computing the output allocation in $O(n^3)$ time.

Since Algorithm 1 assumes a total ordering over the set of houses (see the tie-breaker in the definition of $next(G, u)$ in Section 2.1), it defines a family of house allocation mechanisms, as opposed to a single mechanism. The related works [2, 3, 6] discussed in Section 1 share the same characteristic.

Remark: In this paper we assume that we are given a total ordering over the set of houses. If instead we are given a total ordering over the set of agents, then we can use this ordering, together with the initial allocation, to induce a total ordering over the set of houses.

3 Properties of Configurations

Lemma 3.1. *Let $G = (U, V, E)$ be a configuration and let $G' = pruned(G)$. Then $distance(G', v) = distance(G, v)$ for all houses v in V , and $next(G', u) = next(G, u)$ for all agents u in U .*

Proof. Let $P(i)$ denote the predicate “for any house v in V such that $distance(G, v) = 2i + 1$, we have $distance(G', v) \leq 2i + 1$ ”. We use induction to prove that $P(i)$ holds for all nonnegative integers i .

Base case: $i = 0$. Let v be a house such that $distance(G, v) = 1$. Let u denote $agent(G, v)$. Thus u belongs to $unsatisfied(G)$, and hence also belongs to $unsatisfied(G')$. Since G' includes the edge (v, u) , we conclude that $distance(G', v) = 1$. Hence $P(0)$ holds.

Induction step. Let i be a nonnegative integer, and assume that $P(i)$ holds. Let v be a house such that $distance(G, v) = 2i + 3$. Let u denote $agent(G, v)$. Let v' denote $next(G, u)$. Thus

$distance(G, v') = 2i + 1$, and the induction hypothesis implies that $distance(G', v') \leq 2i + 1$. Since the edges (v, u) and (u, v') belong to $pruned(G')$, we conclude that $distance(G', v) \leq 2i + 3$, as required. Hence $P(i + 1)$ holds, completing our proof by induction.

Since $P(i)$ holds for all nonnegative integers i , and since $distance(G, v) = \infty$ implies $distance(G', v) \leq distance(G, v)$, we conclude that $distance(G', v) \leq distance(G, v)$ for all houses v in V . On the other hand, since G' is a subgraph of G , $distance(G', v) \geq distance(G, v)$ for all houses v in V . Thus $distance(G', v) = distance(G, v)$ for all houses v in V .

Let u be an agent in U . We prove that $next(G', u) = next(G, u)$ by considering two cases.

Case 1: $next(G, u) = nil$. Thus $\Gamma(G', u)$ is empty, and hence $next(G', u) = nil$.

Case 2: $next(G, u) \neq nil$. Let v denote $next(G, u)$. Thus $distance(G, v)$ is finite. Since $distance(G', v) = distance(G, v)$, we conclude that $distance(G', v)$ is finite. Since $\Gamma(G', u) = \{v\}$ and $distance(G', v)$ is finite, we have $next(G', u) = v$, as required. \square

Lemma 3.2. *Let G be a configuration, and let C belong to $cycles(G)$. Then at least one agent in $unsatisfied(G)$ is on C .*

Proof. Immediate from Lemma 3.1. \square

Lemma 3.3. *Let G be a configuration, let C belong to $cycles(G)$, and let u be an agent on C . Then u belongs to $satisfied(trade(G, C))$.*

Proof. Immediate from the definition of $trade(G, C)$. \square

Lemma 3.4. *Let $G = (U, V, E)$ be a configuration, let C belong to $cycles(G)$, let G' denote $trade(G, C)$, and let v belong to V . Then $distance(G', v) \geq distance(G, v)$.*

Proof. Immediate from Lemma A.1, which is proven in Appendix A. \square

For any configuration G , we define $frozen(G)$ as the set of all agents u in $satisfied(G)$ such that $next(G, u) = nil$.

Lemma 3.5. *Let $G = (U, V, E)$ be a configuration, let u belong to U , and let v denote $house(G, u)$. Then $distance(G, v) = \infty$ if and only if u belongs to $frozen(G)$.*

Proof. If u belongs to $frozen(G)$, then $distance(G, v) = \infty$. Assume that $distance(G, v) = \infty$. We consider two cases.

Case 1: u belongs to $satisfied(G)$. Since $distance(G, v) = \infty$, we deduce that $next(G, u) = nil$. Hence u belongs to $frozen(G)$, as required.

Case 2: u belongs to $unsatisfied(G)$. Then $distance(G, v) = 1$, a contradiction. \square

Lemma 3.6. *Let G be a configuration. Then G is final if and only if $cycles(G)$ and $exhausted(G)$ are empty.*

Proof. For the “only if” direction, assume that G is final. Hence $unsatisfied(G)$ is empty. Thus Lemma 3.2 implies that $cycles(G)$ is empty, and the definition of $exhausted(G)$ implies that $exhausted(G)$ is empty.

For the “if” direction, assume that $cycles(G)$ and $exhausted(G)$ are empty. Let G be of the form (U, V, E) and let U' denote $U \setminus frozen(G)$. Since U' is disjoint from $exhausted(G) \cup frozen(G)$, we deduce that $next(G, u) \neq nil$ for all agents u in U' . Thus there is a cycle in $pruned(G)$ unless U' is empty. Since $cycles(G)$ is empty, we conclude that U' is empty, and hence $frozen(G) = U$. Thus $unsatisfied(G)$ is empty, and G is final. \square

For any configurations $G = (U, V, E)$ and $G' = (U', V', E')$, we write $G \lesssim G'$ to mean that the following conditions are satisfied: $U = U'$; $V = V'$; for all agents u in U , $\Gamma(G', u)$ contains $\Gamma(G, u)$; for all agents u in $\text{satisfied}(G)$, $\Gamma(G', u) = \Gamma(G, u)$; $\text{satisfied}(G')$ contains $\text{satisfied}(G)$; for all houses v in V , $\text{distance}(G', v) \geq \text{distance}(G, v)$; for all agents u in $\text{frozen}(G) \cup \text{unsatisfied}(G')$, $\text{house}(G', u) = \text{house}(G, u)$. Thus for any configuration G , we have $G \lesssim G$. Lemma 3.8 below establishes that the relation \lesssim defines a preorder over the set of all configurations.

Lemma 3.7. *Let G and G' be configurations such that $G \lesssim G'$. Then $\text{frozen}(G')$ contains $\text{frozen}(G)$.*

Proof. Let u be an agent in $\text{frozen}(G)$ and let v denote $\text{house}(G, u)$. Since $G \lesssim G'$, we have $\text{distance}(G', v) \geq \text{distance}(G, v)$. Since Lemma 3.5 implies that $\text{distance}(G, v) = \infty$, we deduce that $\text{distance}(G', v) = \infty$. Hence Lemma 3.5 implies that u belongs to $\text{frozen}(G')$. \square

Lemma 3.8. *Let G , G' , and G'' be configurations such that $G \lesssim G'$ and $G' \lesssim G''$. Then $G \lesssim G''$.*

Proof. Immediate from the definition of \lesssim and Lemma 3.7. \square

Lemma 3.9. *Let $G = (U, V, E)$ be a configuration, let u be an agent such that $\text{next}(G, u) \neq \text{nil}$, let v denote $\text{next}(G, u)$, and let G' be a configuration such that $G \lesssim G'$, $\text{distance}(G', v) = \text{distance}(G, v)$, and $\Gamma(G', u) = \Gamma(G, u)$. Then $\text{next}(G', u) = \text{next}(G, u)$.*

Proof. Since $G \lesssim G'$, we have $\text{distance}(G', v') \geq \text{distance}(G, v')$ for all v' in V . Since $\Gamma(G', u) = \Gamma(G, u)$ and $\text{distance}(G', v) = \text{distance}(G, v)$, we conclude that $\text{next}(G', u) = \text{next}(G, u)$, as required. \square

Lemma 3.10. *Let $G = (U, V, E)$ be a configuration, let C belong to $\text{cycles}(G)$, and let G' denote $\text{trade}(G, C)$. Then $G \lesssim G'$ holds.*

Proof. The configuration G' is of the form (U, V, E') , and for all agents u in U , we have $\Gamma(G', u) = \Gamma(G, u)$. Lemmas 3.2 and 3.3 imply that $\text{satisfied}(G')$ properly contains $\text{satisfied}(G)$. Lemma 3.4 implies that $\text{distance}(G', v) \geq \text{distance}(G, v)$ for all v in V . Since $\text{next}(G, u) = \text{nil}$ for any agent u in $\text{frozen}(G)$, no agent on C belongs to $\text{frozen}(G)$. Hence for all agents u in $\text{frozen}(G')$, $\text{house}(G', u) = \text{house}(G, u)$. Lemma 3.3 implies that all agents u on C belong to $\text{satisfied}(G')$. Thus for all agents u in $\text{unsatisfied}(G')$, we have $\text{house}(G', u) = \text{house}(G, u)$. The claim of the lemma follows. \square

4 Properties of Wpps

Lemma 4.1. *Let $W = (U, V, \succ)$ be a wpp, let G belong to $\text{configs}(W)$, let u belong to $\text{exhausted}(G)$, let G' denote $\text{reveal}(W, G, u)$, and let u' belong to $U - u$. Then $G \lesssim G'$. Furthermore, if u belongs to $\text{unsatisfied}(G')$ or u' belongs to a cycle in $\text{cycles}(G)$, then $\text{next}(G', u') = \text{next}(G, u')$.*

Proof. The claim that $G \lesssim G'$ holds is immediate from Lemma B.2, which is proven in Appendix B.

Assume that u belongs to $\text{unsatisfied}(G')$. Hence Lemma B.2 implies that $\text{distance}(G', v) = \text{distance}(G, v)$ for all houses v in V . Since $\Gamma(G', u')$ is equal to $\Gamma(G, u')$, we deduce that $\text{next}(G', u')$ is equal to $\text{next}(G, u')$, as required.

Assume that u' belongs to a cycle C in $\text{cycles}(G)$. Thus $\text{next}(G, u')$ is a house on C ; let v denote $\text{next}(G, u')$. Lemma 3.2 implies that $\text{attractor}(G, v)$ (see Appendix B for the definition of $\text{attractor}(G, v)$) is an agent on C , and hence is not equal to u . Thus Lemma B.2 implies $\text{distance}(G', v) = \text{distance}(G, v)$. Since $G \lesssim G'$, $\text{distance}(G', v) = \text{distance}(G, v)$, and $\Gamma(G', u') = \Gamma(G, u')$, Lemma 3.9 implies that $\text{next}(G', u') = \text{next}(G, u')$, as required. \square

For any wpp $W = (U, V, \succsim)$, any subset U' of U , and any G in $\text{configs}(W)$, we define $\Gamma(W, G, U')$ as the set of all configurations G' such that (G, G') is an edge in $\text{moves}(W)$ and the label of edge (G, G') is neither an agent in U' nor a cycle that includes an agent in U' , and we define $\Gamma^*(W, G, U')$ as the set of all configurations G' in $\text{configs}(W)$ for which there exists a nonnegative integer k and a sequence of configurations G_i , $0 \leq i \leq k$, such that the following conditions hold: $G_0 = G$; $G_k = G'$; G_{i+1} belongs to $\Gamma(W, G_i, U')$ for $0 \leq i < k$.

For any wpp W , and any configuration G in $\text{configs}(W)$, we define $\Gamma^*(W, G)$ as $\Gamma^*(W, G, \emptyset)$.

Lemma 4.2. *Let W be a wpp, let G belong to $\text{configs}(W)$, and let G' belong to $\Gamma^*(W, G)$. Then $G \lesssim G'$.*

Proof. Immediate from Lemmas 3.8, 3.10, and 4.1. \square

The next lemma shows that the nondeterministic algorithm of Figure 1 terminates within a polynomial number of iterations. We present a faster implementation in Appendix F.

Lemma 4.3. *Consider an execution of the while loop of Algorithm 1 on a wpp $W = (U, V, \succsim)$ and an initial configuration G in $\text{configs}(W)$. Then the while loop terminates with a final configuration within at most $|V|^2 + |V|$ iterations.*

Proof. For any configuration $G' = (U, V, E)$, let $f(G')$ denote $|\text{satisfied}(G')|$, and let $g(G')$ denote $|E| - |V|$. Thus $0 \leq f(G') \leq |V|$ and $0 \leq g(G') \leq |V|^2$. Let G_i denote the configuration associated with the variable G after i iterations of the while loop have been completed. Thus $f(G_0) = g(G_0) = 0$. Lemma 4.2 implies that $f(G_{i+1}) \geq f(G_i)$ and $g(G_{i+1}) \geq g(G_i)$ for all $i \geq 0$. Lemmas 3.2, 3.3, and 4.2 together imply that if $G_{i+1} = \text{trade}(G_i, C)$ for some C in $\text{cycles}(G_i)$, then $f(G_{i+1}) \geq f(G_i) + 1$; thus we can have at most $|V|$ iterations in this category. If $G_{i+1} = \text{reveal}(W, G_i, u)$ for some u in $\text{exhausted}(G_i)$, then $g(G_{i+1}) \geq g(G_i) + 1$; thus we can have at most $|V|^2$ iterations in this category.

Lemma 3.6 implies that the configuration corresponding to program variable G is final when the while loop of Algorithm 1 terminates. \square

Lemma 4.4. *Let $W = (U, V, \succsim)$ be a wpp, let G be a configuration in $\text{configs}(W)$, let G' belong to $\Gamma(W, G)$, let ℓ be the label of edge (G, G') in $\text{moves}(W)$, and let G'' be a configuration in $\Gamma^*(W, G)$ such that there is no label- ℓ edge outgoing from G'' . Then G'' belongs to $\Gamma^*(W, G')$.*

Proof. Immediate from Lemmas C.1, C.2, C.3, C.4, C.5, and C.6, which are proven in Appendix C. \square

For any wpp $W = (U, V, \succsim)$, any subset U' of U , and any configuration G in $\text{configs}(W)$, we define $\text{sinks}(W, G, U')$ as the set of all configurations G' in $\Gamma^*(W, G, U')$ such that $\Gamma(W, G', U')$ is empty.

Lemma 4.5. *Let $W = (U, V, \succsim)$ be a wpp, let U' be a subset of U , and let G belong to $\text{configs}(W)$. Then $|\text{sinks}(W, G, U')| = 1$.*

Proof. Lemma 4.3 implies that $\text{sinks}(W, G, U')$ is nonempty. Let G^* belong to $\text{sinks}(W, G, U')$. Thus G^* is contained in $\Gamma^*(W, G, U')$ and $\text{sinks}(W, G^*, U') = \{G^*\}$.

Lemma 4.4 implies that $\text{sinks}(W, G, U') = \text{sinks}(W, G', U')$ for all G' in $\Gamma(W, G, U')$. Thus, by repeated application of Lemma 4.4, we deduce that $\text{sinks}(W, G, U') = \text{sinks}(W, G', U')$ for all G' in $\Gamma^*(W, G, U')$. Since G^* is contained in $\Gamma^*(W, G, U')$, we find that $\text{sinks}(W, G, U') = \text{sinks}(W, G^*, U') = \{G^*\}$. The claim of the lemma follows. \square

For any wpp $W = (U, V, \succsim)$, any configuration G in $\text{configs}(W)$, any agent u in U , and any house v in V , we define the predicate $\text{bottom}(W, G, u, v)$ to mean that v belongs to $\Gamma(G, u)$ and $v' \succsim_u v$ holds for all houses v' in $\Gamma(G, u)$.

For any wpp $W = (U, V, \succsim)$, we define $\text{admissible}(W)$ as the set of all configurations G in $\text{configs}(W)$ such that the following conditions hold: for any agent u in $\text{frozen}(G)$, we have $\text{bottom}(W, G, u, \text{house}(G, u))$; for any agent u in $U \setminus \text{frozen}(G)$ and any house v in $\Gamma(G, u)$ such that $\text{agent}(G, v)$ does not belong to $\text{frozen}(G)$, we have $\text{bottom}(W, G, u, v)$.

Lemma 4.6. *Let $W = (U, V, \succsim)$ be a wpp, let G belong to $\text{admissible}(W)$, let G' belong to $\Gamma^*(W, G)$, let u belong to U , let v denote $\text{house}(G, u)$, and let v' denote $\text{house}(G', u)$. Then G' belongs to $\text{admissible}(W)$ and $v' \succsim_u v$. Furthermore, if u belongs to $\text{satisfied}(G)$ then $v' \sim_u v$.*

Proof. By induction using Lemmas D.1 and D.2, which are proven in Appendix D. \square

Lemma 4.7. *Let W be a wpp and let G be an initial configuration in $\text{configs}(W)$. Then G belongs to $\text{admissible}(W)$.*

Proof. Straightforward. \square

For any wpp $W = (U, V, \succsim)$, any G in $\text{configs}(W)$, and any subset U' of U , we define $\text{sink}(W, G, U')$ as the unique (by Lemma 4.5) element of $\text{sinks}(W, G, U')$.

For any wpp W and any G in $\text{configs}(W)$, we define $\text{sink}(W, G)$ as $\text{sink}(W, G, \emptyset)$.

For any wpp $W = (U, V, \succsim)$ and any agent u in U , we define $\text{sinks}(W, u)$ as the set of all configurations G in $\text{admissible}(W)$ such that $\Gamma(W, G, \{u\})$ is empty and u belongs to $\text{unsatisfied}(G)$.

For any configuration G and any agent u in $\text{unsatisfied}(G)$, we define $\text{reach}(G, u)$ as the set of all houses v such that there is a path from v to u in $\text{pruned}(G)$.

Lemma 4.8. *Let $W = (U, V, \succsim)$ be a wpp, let u belong to U , let G be a configuration in $\text{sinks}(W, u)$, and let v belong to V . Then exactly one of the following two conditions is satisfied: $\text{agent}(G, v)$ belongs to $\text{frozen}(G)$; v belongs to $\text{reach}(G, u)$.*

Proof. Let u^* denote $\text{agent}(G, v)$. We consider two cases.

Case 1: u^* belongs to $\text{frozen}(G)$. Thus u^* belongs to $\text{satisfied}(G)$ and $\text{next}(G, u^*) = \text{nil}$. Since u^* belongs to $\text{satisfied}(G)$ and G belongs to $\text{sinks}(W, u)$, we deduce that $u^* \neq u$. Since $\text{next}(G, u^*) = \text{nil}$, agent u^* is the only agent reachable via a path from v in $\text{pruned}(G)$. Since $u \neq u^*$, we conclude that v does not belong to $\text{reach}(G, u)$.

Case 2: u^* does not belong to $\text{frozen}(G)$. Consider the path P in $\text{pruned}(G)$ obtained by starting at v and repeatedly following the outgoing edge of the current vertex until (1) a cycle C

is formed or (2) an agent u' is reached such that $next(G, u') = nil$. If (1) occurs, then since G belongs to $sinks(W, u)$, we deduce that u belongs to C and hence that v belongs to $reach(G, u)$.

Now assume that (2) occurs, and let v' denote $house(G, u')$. We claim that $distance(G, v')$ is finite. If $v = v'$, the claim follows from Lemma 3.5 and the Case 2 condition. Assume $v \neq v'$ and let u'' denote the agent preceding v' on P . Then $v' = next(G, u'')$ and hence $distance(G, v')$ is finite, completing the proof of the claim. The claim, together with Lemma 3.5, implies that u' does not belong to $frozen(G)$. Since $next(G, u') = nil$ and u' does not belong to $frozen(G)$, we deduce that u' belongs to $unsatisfied(G)$ and hence that u' belongs to $exhausted(G)$. Since G belongs to $sinks(W, u)$ and u' belongs to $exhausted(G)$, we deduce that $u' = u$ and hence that v belongs to $reach(G, u)$. \square

Lemma 4.9. *Let $W = (U, V, \succ)$ be a wpp, let u belong to U , and let G be a configuration in $sinks(W, u)$. Then $|\Gamma(W, G)| = 1$ and the following claims hold, where G' denotes the unique configuration in $\Gamma(W, G)$.*

1. *If u belongs to $unsatisfied(G')$, then $reach(G', u) = reach(G, u)$ and G' belongs to $sinks(W, u)$.*
2. *If u belongs to $satisfied(G')$, then $house(G', u)$ belongs to $top(W, u, reach(G, u))$.*

Proof. We first argue that $|\Gamma(W, G)| = 1$. Since G belongs to $sinks(W, u)$, we find that $\Gamma(W, G, \{u\})$ is empty. Hence u appears on any cycle C in $cycles(G)$. Since the cycles in $cycles(G)$ are disjoint, we have $|cycles(G)| \leq 1$. We consider two cases.

Case 1: $next(G, u) = nil$. Since u belongs to $unsatisfied(G)$, we deduce that u belongs to $exhausted(G)$ and that $cycles(G)$ is empty. It follows that $\Gamma(W, G) = \{reveal(W, G, u)\}$. Hence $|\Gamma(W, G)| = 1$.

Case 2: $next(G, u) \neq nil$. Thus u does not belong to $exhausted(G)$, and hence $exhausted(G)$ is empty. Since u belongs to $unsatisfied(G)$, Lemma 3.6 implies that $cycles(G)$ is nonempty. Since $|cycles(G)| \leq 1$, we deduce that $|cycles(G)| = 1$ and hence $|\Gamma(W, G)| = 1$.

We now address Claim 1. Assume that u belongs to $unsatisfied(G')$. Thus $G' = reveal(W, G, u)$ and u belongs to $exhausted(G)$. Hence Lemma 4.1 implies that $G \lesssim G'$ and $next(G', u')$ is equal to $next(G, u')$ for all agents u' in $U - u$. It follows that $reach(G', u)$ is equal to $reach(G, u)$ and $exhausted(G') \cap (U - u)$ is equal to $exhausted(G) \cap (U - u)$. Since G belongs to $sinks(W, u)$, we find that $exhausted(G) \cap (U - u)$ is empty. Thus $exhausted(G') \cap (U - u)$ is empty. Furthermore, if C belong to $cycles(G')$ and u is not on C , then C belongs to $cycles(G)$; since G belongs to $sinks(W, u)$, we deduce that u belongs to every cycle in $cycles(G')$. Since $exhausted(G') \cap (U - u)$ is empty, we conclude that $\Gamma(W, G', \{u\})$ is empty.

Since G belongs to $sinks(W, u)$, we know that G belongs to $admissible(W)$. Hence Lemma 4.6 (or simply Lemma D.1) implies that G' belongs to $admissible(W)$.

Since G' is a configuration in $admissible(W)$ such that $\Gamma(W, G', \{u\})$ is empty and u belongs to $unsatisfied(G')$, we conclude that G' belongs to $sinks(W, u)$.

It remains to address Claim 2. Assume that u belongs to $satisfied(G')$. Let v denote $house(G, u)$, let v' denote $house(G', u)$, and let u' denote $agent(G, v')$. Let V_0 denote the set of all houses v_0 such that $agent(G, v_0)$ belongs to $frozen(G)$. Lemma 4.2 implies $G \lesssim G'$, and hence that v' does not belong to V_0 . By Lemma 4.8, $V_0 = V \setminus reach(G, u)$. Thus v' belongs to $reach(G, u)$. We consider two cases.

Case 1: u does not belong to $exhausted(G)$. Thus $G' = trade(G, C)$ for some C in $cycles(G)$ such that u is on C , $\Gamma(G, u) = \Gamma(G', u)$, and v' belongs to $\Gamma(G, u)$. Let V_1 denote $reach(G, u) \cap$

$\Gamma(G, u)$, and let V_2 denote $reach(G, u) \setminus V_1$. Thus v' belongs to V_1 . Since G belongs to $admissible(W)$, we have $v' \sim_u v''$ for all houses v'' in V_1 . Since G belongs to $configs(W)$, we have $v' \succ_u v''$ for all houses v'' in V_2 . We conclude that v' belongs to $top(W, u, V_1 \cup V_2) = top(W, u, reach(G, u))$, as required.

Case 2: u belongs to $exhausted(G)$. Thus $G' = reveal(W, G, u)$, $v' = v$, and it follows that $v' \succ_u v''$ holds for all houses v'' in $V \setminus V_0 = reach(G, u)$. Thus v' belongs to $top(W, u, reach(G, u))$, as required. \square

Lemma 4.10. *Let $W = (U, V, \succ)$ be a wpp, let u belong to U , and let G belong to $sinks(W, u)$. Then $house(sink(W, G), u)$ belongs to $top(W, u, reach(G, u))$.*

Proof. By repeated application of Lemma 4.9. \square

5 Main Results

For any housing market instance $I = (W, G)$, we define $sink(I)$ as $sink(W, G)$. Theorem 1 below establishes that Algorithm 1 defines a deterministic mechanism.

Theorem 1. *Let I be a housing market instance. Then any execution of Algorithm 1 on instance I produces the same allocation.*

Proof. Any execution of Algorithm 1 on instance I returns $allocation(sink(I))$. \square

For any housing market instance $I = (W, G)$ where $W = (U, V, \succ)$ and any agent u in U , we define $sink(I, u)$ as $sink(W, G, \{u\})$, we define $houses(I, u)$ as $reach(sink(I, u), u)$, and we define $house(I, u)$ as $house(sink(I, u), u)$.

Lemma 5.1. *Let $I = (W, G)$ be a housing market instance with $W = (U, V, \succ)$, and let u belong to U . Then $house(I, u)$ belongs to $top(W, u, houses(I, u))$.*

Proof. Let G' denote $sink(I, u)$. Thus G' belongs to $\Gamma^*(W, G, \{u\})$ and $\Gamma(W, G', \{u\})$ is empty. Since Lemma 4.7 implies that G belongs to $admissible(W)$, Lemma 4.6 implies that G' belongs to $admissible(W)$. Since I is a housing market instance, the configuration G is initial, and hence $\Gamma(G, u)$ is empty. Since $\Gamma(G, u)$ is empty and G' belongs to $\Gamma^*(W, G, \{u\})$, we find that $\Gamma(G', u)$ is empty and hence u belongs to $unsatisfied(G')$. Since G' belongs to $admissible(W)$, $\Gamma(W, G', \{u\})$ is empty, and u belongs to $unsatisfied(G')$, we conclude that G' belongs to $sinks(W, u)$. Thus Lemma 4.10 implies that $house(sink(W, G'), u)$ belongs to $top(W, u, reach(G', u))$. Theorem 1 implies that $sink(I) = sink(W, G')$. Hence $house(I, u)$ belongs to $top(W, u, reach(G', u))$. By definition, $houses(I, u)$ is equal to $reach(G', u)$, and the claim of the lemma follows. \square

For any wpps $W = (U, V, \succ)$ and $W' = (U, V, \succ')$, and any u in U , we define the predicate $equiv(W, W', u)$ to hold if $v \succ_{u'} v'$ is logically equivalent to $v \succ'_u v'$ for all agents u' in $U - u$ and all houses v and v' in V .

Lemma 5.2. *Let $G = (U, V, E)$ be an initial configuration, let u belong to U , and let $I = (W, G)$ and $I' = (W', G)$ be housing market instances such that $equiv(W, W', u)$ holds. Then $sink(I, u) = sink(I', u)$ and $houses(I, u) = houses(I', u)$.*

Proof. The first equation holds because $\text{sink}(I, u)$ and $\text{sink}(I', u)$ are each independent of the preferences of u . The second equation follows since $\text{houses}(I, u) = \text{reach}(\text{sink}(I, u), u) = \text{reach}(\text{sink}(I', u), u) = \text{houses}(I', u)$. \square

Theorem 2. *Each mechanism in the family associated with Algorithm 1 is strategyproof, produces a Pareto-efficient core allocation, and produces a strict core allocation when the strict core is nonempty.*

Proof. Lemma E.1 implies that the family of mechanisms associated with Algorithm 1 lies within the class of GATTC mechanisms introduced by Aziz and de Keijzer [3]. Every mechanism in the GATTC family produces a Pareto-efficient core allocation, and produces a strict core allocation when the strict core is nonempty [3]. It remains only to establish strategyproofness.

Let $I = (W, G)$ be a housing market instance with $W = (U, V, \succsim)$, and let u be an agent in U such that \succsim reflects u 's true preferences over the houses in V . Let $I' = (W', G)$ be a housing market instance such that $\text{equiv}(W, W', u)$ holds. Let V' denote $\text{houses}(I, u)$, which is equal to $\text{houses}(I', u)$ by Lemma 5.2. Lemma 5.1 implies that $\text{house}(I, u)$ belongs to $\text{top}(W, u, V')$ and that $\text{house}(I', u)$ belongs to V' . Hence $\text{house}(I, u) \succsim_u \text{house}(I', u)$. \square

6 Concluding Remarks

Abraham et al. [1] show how to use the $O(n^{2.5})$ -time Hopcroft-Karp algorithm [5] for computing a maximum cardinality matching in a bipartite graph to obtain the same asymptotic time bound for computing maximum cardinality Pareto-efficient matchings for house allocation problems. For the setting of housing markets with indifferences considered in the present paper, it is straightforward to argue (by considering instances in which each agent has at most two tiers of preference, where the first tier represents edges, and the second tier represents non-edges) that if a Pareto-efficient mechanism admits an $O(f(n))$ -time implementation, then the complexity of computing a maximum cardinality matching in a bipartite graph is $O(f(n))$. Given the foregoing remarks, it is natural to ask whether the Hopcroft-Karp algorithm can be adapted to our setting to improve the $O(n^3)$ time bound established in Appendix F to $O(n^{2.5})$. Based on a preliminary investigation, it remains unclear whether such an improvement is achievable, so we leave this question open.

Acknowledgments

The author is grateful to Himanshu Chauhan, Onur Domaniç, Pravesh Kothari, and Chris Tosh for their many helpful comments on earlier versions of this work.

References

- [1] D. J. Abraham, K. Cechlárová, D. F. Manlove, and K. Mehlhorn. Pareto-optimality in housing allocation problems. In *Proceedings of the 15th International Symposium on Algorithms and Computation, LNCS 3341*, pages 3–15, 2004. Also appears in LNCS 3827, with publisher typesetting errors corrected.

- [2] J. Alcalde-Unzu and E. Molis. Exchange of indivisible goods and indifferences: The Top Trading Absorbing Sets mechanisms. *Games and Economic Behavior*, 73:1–16, 2011.
- [3] A. Aziz and B. de Keijzer. Housing markets with indifferences: A tale of two mechanisms. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1249–1255, July 2012.
- [4] C. G. Bird. Group incentive compatibility in a market with indivisible goods. *Economic Letters*, 14:309–313, 1984.
- [5] J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matching in bipartite graphs. *SIAM Journal on Computing*, 2:225–231, 1973.
- [6] P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory*, 147:1913–1946, 2012.
- [7] J. Ma. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23:75–83, 1994.
- [8] T. Quint and J. Wako. On houseswapping, the strict core, segmentation, and linear programming. *Mathematics of Operations Research*, 29:861–877, 2004.
- [9] A. E. Roth. Incentive compatibility in a market with indivisible goods. *Economic Letters*, 9:127–132, 1982.
- [10] A. E. Roth and A. Postlewaite. Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, 4:131–137, 1977.
- [11] L. S. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1:23–37, 1974.
- [12] J. Wako. A note on the strong core of a market with indivisible goods. *Journal of Mathematical Economics*, 13:189–194, 1984.

A Trading and Distance

The purpose of this section is to establish Lemma A.1, which immediately implies Lemma 3.4.

Lemma A.1. *Let $G = (U, V, E)$ be a configuration, let C belong to $\text{cycles}(G)$, and let G' denote $\text{trade}(G, C)$. For any nonnegative integer k , let $P(k)$ denote the predicate “for any house v in V such that $\text{distance}(G', v) = 2k + 1$, the following two conditions hold: (1) if v does not belong to C then $\text{distance}(G, v) \leq 2k + 1$; (2) if v belongs to C then $\text{distance}(G, v) \leq 2k - 1$.”*

Proof. We use induction to prove that $P(k)$ holds for all nonnegative integers k .

Base case ($k = 0$). Let v be a house in V such that $\text{distance}(G', v) = 1$. Lemma 3.3 implies that v does not belong to C . Let u denote $\text{agent}(G', v)$. Since $\text{distance}(G', v) = 1$, we know that u belongs to $\text{unsatisfied}(G')$. Since v does not belong to C , we have $\text{agent}(G, v) = u$. Since u

belongs to $unsatisfied(G')$, $\Gamma(G', u) = \Gamma(G, u)$, and $house(G', u) = house(G, u)$, we deduce that u belongs to $unsatisfied(G)$. Hence $distance(G, v) = 1$, as required.

Induction step. Let k be a nonnegative integer, and assume that $P(k)$ holds. We need to prove that $P(k + 1)$ holds. Let v be a house in V such that $distance(G', v) = 2(k + 1) + 1 = 2k + 3$. Let u denote $agent(G', v)$. Since $distance(G', v) = 2k + 3$ and $agent(G', v) = u$, there is a house v' in $\Gamma(G', u)$ such that $distance(G', v') = 2k + 1$. The induction hypothesis implies that $distance(G, v') \leq 2k + 1$. We now consider two cases.

Case 1: v does not belong to C . Thus $agent(G, v) = agent(G', v) = u$. Since $\Gamma(G', u) = \Gamma(G, u)$, there is a path of length two from v to v' in G , and hence $distance(G, v) \leq 2k + 3$, as required.

Case 2: v belongs to C . Thus $next(G, u) = v$. Since $\Gamma(G', u) = \Gamma(G, u)$, we know that v' belongs to $\Gamma(G, u)$. Since $next(G, u) = v$ and v' belongs to $\Gamma(G, u)$, we deduce that $distance(G, v) \leq distance(G, v')$. Since $distance(G, v') \leq 2k + 1$, we have $distance(G, v) \leq 2k + 1$, as required. \square

B Edge Revelation

The purpose of this section is to establish Lemma B.2, which we use to prove Lemma 4.1.

Based on Lemma 3.1, we know that if $distance(G, v)$ is finite, then the unique outgoing path of length $distance(G, v)$ starting at v in $pruned(G)$ is a shortest path in G from v to an agent in $unsatisfied(G)$. We define this agent as $attractor(G, v)$. If $distance(G, v) = \infty$ then we define $attractor(G, v)$ as nil .

Lemma B.1. *Let $G = (U, V, E)$ be a configuration, let u belong to $unsatisfied(G)$, let v belong to V , and let G' denote $(U, V, E + (u, v))$. Then $G \lesssim G'$ and the following claims hold.*

1. *If u belongs to $unsatisfied(G')$, then for any house v' in V , we have $attractor(G', v') = attractor(G, v')$ and $distance(G', v') = distance(G, v')$.*
2. *If u belongs to $satisfied(G')$, then for any house v' in V such that $attractor(G, v') \neq u$, we have $distance(G', v') = distance(G, v')$.*

Proof. In order to establish that $G \lesssim G'$ holds, the only nontrivial conjunct to be shown is that $distance(G', v') \geq distance(G, v')$ for all houses v' in V .

Case 1: u belongs to $unsatisfied(G')$. Hence $u \neq agent(G, v)$ and the edge (u, v) does not belong to a shortest path in G' from any house v' to an agent in $unsatisfied(G')$. It follows that for any house v' in V , we have $distance(G', v') = distance(G, v')$. Hence $G \lesssim G'$ holds. For any agent u' in $U - u$, we have $next(G', u') = next(G, u')$ since $\Gamma(G', u') = \Gamma(G, u')$ and $distance(G', v') = distance(G, v')$ for all v' in V . Hence $pruned(G')$ is the same as $pruned(G)$ except that $next(G', u)$ may differ from $next(G, u)$. Since u belongs to $unsatisfied(G')$, we deduce that $attractor(G', v') = attractor(G, v')$ for all houses v' in V .

Case 2: u belongs to $satisfied(G')$. Hence $u = agent(G, v)$. Let v' be an arbitrary house in V . We need to prove that $distance(G', v') \geq distance(G, v')$ and that this inequality is tight if $attractor(G, v') \neq u$.

Case 2.1: $distance(G', v') = \infty$. Thus $distance(G', v') \geq distance(G, v')$. It remains to prove that if $attractor(G, v') \neq u$, then $distance(G, v') = \infty$.

Case 2.1.1: $\text{attractor}(G, v') = \text{nil}$. Thus $\text{distance}(G, v') = \infty$, as required.

Case 2.1.2: $\text{attractor}(G, v') = u'$ for some agent u' in $U - u$. Then there is a path P in G from v' to u' such that u does not appear on P . It follows that path P also exists in G' . Since u' belongs to the set $\text{unsatisfied}(G')$, which is contained in the set $\text{unsatisfied}(G)$, we conclude that $\text{distance}(G', v')$ is finite, contradicting the Case 2.1 assumption.

Case 2.2: $\text{distance}(G', v')$ is finite. Let P be a shortest path in G' from v' to an agent in $\text{unsatisfied}(G')$; thus P is of length $\text{distance}(G', v')$. We need to argue that $\text{distance}(G, v')$ is at most the length of P . If u does not belong to path P , then P is a path in G from v' to an agent in $\text{unsatisfied}(G)$, and so $\text{distance}(G, v')$ is at most the length of P . If u belongs to path P , then let P' denote the prefix of P terminating at u , and observe that P' is a path in G from v' to an agent in $\text{unsatisfied}(G)$; hence $\text{distance}(G, v')$ is at most the length of P' , which is at most the length of P . Hence $G \lesssim G'$ holds.

It remains to argue that if $\text{attractor}(G, v') \neq u$, then $\text{distance}(G', v') = \text{distance}(G, v')$. We have already established that $\text{distance}(G, v')$ is finite; hence $\text{attractor}(G, v') \neq \text{nil}$. Let u^* denote $\text{attractor}(G, v')$, and assume that $u^* \neq u$. Let P denote the unique path of length $\text{distance}(G, v')$ in $\text{pruned}(G)$ from v' to u^* . Since u^* belongs to $\text{unsatisfied}(G)$ and $u^* \neq u$, we conclude that u^* belongs to $\text{unsatisfied}(G')$. Since P is a path in $\text{pruned}(G)$, which is a subgraph of G' , we conclude that P exists in G' . Since P exists in G' , we deduce that $\text{distance}(G', v') \leq \text{distance}(G, v')$. Since we have already established that $\text{distance}(G', v') \geq \text{distance}(G, v')$, we conclude that $\text{distance}(G', v') = \text{distance}(G, v')$, as required. \square

Lemma B.2. *Let $G = (U, V, E)$ be configuration, let u be an agent in $\text{unsatisfied}(G)$, let E' be a set of edges in $\{u\} \times V$, and let G' denote the configuration $(U, V, E \cup E')$. Then $G \lesssim G'$ and the following claims hold.*

1. *If u belongs to $\text{unsatisfied}(G')$, then for any house v in V , we have $\text{distance}(G', v) = \text{distance}(G, v)$.*
2. *If u belongs to $\text{satisfied}(G')$, then for any house v in V such that $\text{attractor}(G, v) \neq u$, we have $\text{distance}(G', v) = \text{distance}(G, v)$.*

Proof. Let E'' denote $E' - (u, \text{house}(G, u))$, and let G'' denote the configuration $(U, V, E \cup E'')$. Then $G \lesssim G''$ holds by repeated application of Lemmas B.1 and 3.8. Furthermore, u belongs to $\text{unsatisfied}(G'')$, and by repeated application of Lemma B.1, we have $\text{attractor}(G'', v) = \text{attractor}(G, v)$ and $\text{distance}(G'', v) = \text{distance}(G, v)$ for all houses v in V . If $E'' = E'$, this completes the proof.

It remains to consider the case where $E'' \neq E'$. In this case, u belongs to $\text{satisfied}(G')$, and by an additional application of Lemma B.1, we find that $G'' \lesssim G'$ holds and $\text{distance}(G', v) = \text{distance}(G'', v)$ for all houses v in V such that $\text{attractor}(G'', v) \neq u$. Since we have established above that $\text{attractor}(G'', v) = \text{attractor}(G, v)$ for all houses v in V , we conclude that $\text{distance}(G', v) = \text{distance}(G'', v)$ for all houses v in V such that $\text{attractor}(G, v) \neq u$.

Since $G \lesssim G''$ and $G'' \lesssim G'$, Lemma 3.8 implies $G \lesssim G'$, as required. Let v be a house in V such that $\text{attractor}(G, v) \neq u$. In the foregoing we have established that $\text{distance}(G'', v) = \text{distance}(G, v)$ and $\text{distance}(G', v) = \text{distance}(G'', v)$. Thus $\text{distance}(G', v) = \text{distance}(G, v)$, as required. \square

C Confluence

The six lemmas below are used to prove Lemma 4.4.

Lemma C.1. *Let $G = (U, V, E)$ be a configuration, let C and C' be distinct cycles in $\text{cycles}(G)$, and let G' denote $\text{trade}(G, C)$. Then C' belongs to $\text{cycles}(G')$.*

Proof. Since each house v in V has outdegree one in $\text{pruned}(G)$, the cycles C and C' are disjoint. Let u be an arbitrary agent on C' . Thus $\text{next}(G, u) \neq \text{nil}$. Let v denote the house $\text{next}(G, u)$, which is on C' . Lemma 3.2 implies that there is at least one agent in $\text{unsatisfied}(G)$ on C' , and hence that $\text{distance}(G, v)$ is finite. Grow a path P in C' by starting at v and following edges of C' until an agent in $\text{unsatisfied}(G)$ is reached. Lemma 3.1 implies that P is of length $\text{distance}(G, v)$. Since P is a portion of the cycle C' , we deduce that P is disjoint from C , and hence that P exists in G' . It follows that $\text{distance}(G', v) \leq \text{distance}(G, v)$. On the other hand, Lemma 3.4 implies that $\text{distance}(G', v) \geq \text{distance}(G, v)$. We conclude that $\text{distance}(G', v) = \text{distance}(G, v)$. By Lemma 3.10, we have $G \lesssim G'$. Since $G \lesssim G'$, $\text{distance}(G', v) = \text{distance}(G, v)$, and $\Gamma(G', u) = \Gamma(G, u)$, Lemma 3.9 implies that $\text{next}(G', u) = \text{next}(G, u)$.

Since $\text{next}(G', u) = \text{next}(G, u)$ for all agents u on C' , and $\text{agent}(G', v) = \text{agent}(G, v)$ for all houses v on C' , the claim of the lemma follows. \square

Lemma C.2. *Let W be a wpp, let G belong to $\text{configs}(W)$, let C belong to $\text{cycles}(G)$, and let G' belong to $\Gamma(W, G) - \text{trade}(G, C)$. Then C belongs to $\text{cycles}(G')$.*

Proof. There are two cases to consider.

Case 1: The configuration G' is of the form $\text{trade}(G, C')$ for some C' in $\text{cycles}(G) - C$. In this case, the claim of the lemma follows from Lemma C.1.

Case 2: The configuration G' is of the form $\text{reveal}(W, G, u_0)$ for some agent u_0 in U . Thus u_0 belongs to $\text{exhausted}(G)$, and hence u_0 is not on C . Thus Lemma 4.1 implies that $\text{next}(G', u) = \text{next}(G, u)$ holds for all agents u on C . Furthermore, $\text{agent}(G', v) = \text{agent}(G, v)$ for all houses v in V . The claim of the lemma follows. \square

Lemma C.3. *Let W be a wpp, let G belong to $\text{configs}(W)$, let u belong to $\text{exhausted}(G)$, and let G' belong to $\Gamma(W, G) - \text{reveal}(W, G, u)$. Then u belongs to $\text{exhausted}(G')$.*

Proof. We claim that $\Gamma(G', u) = \Gamma(G, u)$. If the configuration G' is of the form $\text{trade}(G, C)$ for some C in $\text{cycles}(G)$, then $\Gamma(G', u') = \Gamma(G, u')$ for all agents u' in U , and the claim holds. Otherwise, the configuration G' is of the form $\text{reveal}(W, G, u_0)$ for some agent u_0 in $U - u$, and since $u \neq u_0$, the claim holds.

Since u belongs to $\text{exhausted}(G)$, we have $\text{distance}(G, v) = \infty$ for all houses v in $\Gamma(G, u)$. Lemma 4.2 implies $G \lesssim G'$, and hence that $\text{distance}(G', v) \geq \text{distance}(G, v)$ for all houses v in V . Since $\Gamma(G', u) = \Gamma(G, u)$, we conclude that $\text{distance}(G', v) = \infty$ for all houses v in $\Gamma(G', u)$, and hence that $\text{next}(G', u) = \text{nil}$.

We claim that $\text{house}(G', u) = \text{house}(G, u)$. If the configuration G' is of the form $\text{trade}(G, C)$ for some C in $\text{cycles}(G)$, then u is not on C since $\text{next}(G, u) = \text{nil}$; hence the claim holds. Otherwise, the configuration G' is of the form $\text{reveal}(W, G, u_0)$ for some agent u_0 in $U - u$; hence $\text{allocation}(G') = \text{allocation}(G)$ and the claim holds.

Since $\text{house}(G', u) = \text{house}(G, u)$, $\Gamma(G', u) = \Gamma(G, u)$, and u belongs to $\text{unsatisfied}(G)$, we conclude that u belongs to $\text{unsatisfied}(G')$. Since u belongs to $\text{unsatisfied}(G')$ and $\text{next}(G', u) = \text{nil}$, we conclude that u belongs to $\text{exhausted}(G')$, as required. \square

Lemmas C.1, C.2 and C.3 ensure that all of the expressions appearing in the next three lemma statements are well-defined.

Lemma C.4. *Let $G = (U, V, E)$ be a configuration, and let C and C' be distinct cycles in $\text{cycles}(G)$. Then $\text{trade}(\text{trade}(G, C), C') = \text{trade}(\text{trade}(G, C'), C)$.*

Proof. Since each agent or house in $\text{pruned}(G)$ has outdegree 1, the cycles C and C' are disjoint. The claim of the lemma follows easily. \square

Lemma C.5. *Let W be a wpp, let G belong to $\text{configs}(W)$, let u belong to $\text{exhausted}(G)$, and let C belong to $\text{cycles}(G)$. Then $\text{trade}(\text{reveal}(W, G, u), C) = \text{reveal}(W, \text{trade}(G, C), u)$.*

Proof. Straightforward. \square

Lemma C.6. *Let W be a wpp, let G belong to $\text{configs}(W)$, and let u and u' be distinct agents in $\text{exhausted}(G)$. Then $\text{reveal}(W, \text{reveal}(W, G, u), u') = \text{reveal}(W, \text{reveal}(W, G, u'), u)$.*

Proof. Straightforward. \square

D Admissibility

The two lemmas below are used to prove Lemma 4.6.

Lemma D.1. *Let $W = (U, V, \succsim)$ be a wpp, let G belong to $\text{admissible}(W)$, and let G' belong to $\Gamma(W, G)$. Then G' belongs to $\text{admissible}(W)$.*

Proof. Let u belong to U , let v denote $\text{house}(G, u)$, and let v' denote $\text{house}(G', u)$. Lemma 4.2 implies $G \lesssim G'$.

Case 1: u belongs to $\text{frozen}(G)$. Since G belongs to $\text{admissible}(W)$, we find that $\text{bottom}(W, G', u, v)$ holds. Since $G \lesssim G'$, we find that u belongs to $\text{frozen}(G')$, $\Gamma(G', u) = \Gamma(G, u)$, and $v' = v$. Hence $\text{bottom}(W, G', u, v')$ holds.

Case 2: u does not belong to $\text{frozen}(G)$. Let V_0 denote the set of all houses v_0 in $\Gamma(G, u)$ such that $\text{agent}(G, v_0)$ does not belong to $\text{frozen}(G)$. Since G belongs to $\text{admissible}(W)$, we have $\text{bottom}(W, G, u, v_0)$ for all houses v_0 in V_0 . We consider two subcases.

Case 2.1: u does not belong to $\text{frozen}(G')$. We need to prove that for any house v in $\Gamma(G', u)$ such that $\text{agent}(G', v)$ does not belong to $\text{frozen}(G')$, we have $\text{bottom}(W, G', u, v)$. Since $G \lesssim G'$, we find that $\Gamma(G', u)$ contains $\Gamma(G, u)$ and $\text{frozen}(G')$ contains $\text{frozen}(G)$. If $\Gamma(G', u) = \Gamma(G, u)$, the desired claim follows immediately. Otherwise, $G' = \text{reveal}(W, G, u)$ and hence u belongs to $\text{exhausted}(G)$. Thus $\text{agent}(G, v_1)$ belongs to $\text{frozen}(G)$ (and hence also $\text{frozen}(G')$) for all houses v_1 in $\Gamma(G, u)$, and the desired claim follows easily from the definition of $\text{reveal}(W, G, u)$.

Case 2.2: u belongs to $\text{frozen}(G')$. Thus u belongs to $\text{satisfied}(G')$. We need to establish $\text{bottom}(W, G', u, v')$. We consider two subcases.

Case 2.2.1: V_0 is nonempty. Thus $\Gamma(G', u) = \Gamma(G, u)$. Since u belongs to $\text{satisfied}(G')$ and $G \lesssim G'$, we deduce that v' belongs to V_0 . Since v' belongs to V_0 , we have $\text{bottom}(W, G, u, v')$. We conclude that $\text{bottom}(W, G', u, v')$ holds, as required.

Case 2.2.2: V_0 is empty. Thus u belongs to $\text{exhausted}(G)$. Since u belongs to $\text{satisfied}(G')$, we conclude that $G' = \text{reveal}(W, G, u)$ and $v' = v$. The desired claim follows easily from the definition of $\text{reveal}(W, G, u)$. \square

Lemma D.2. *Let $W = (U, V, \succsim)$ be a wpp, let G belong to $\text{admissible}(W)$, let G' belong to $\Gamma(W, G)$, let u belong to U , let v denote $\text{house}(G, u)$, and let v' denote $\text{house}(G', u)$. Then $v' \succsim_u v$. Furthermore, if u belongs to $\text{satisfied}(G)$ then $v' \sim_u v$.*

Proof. If $v = v'$ then the claim of the lemma is trivial. For the remainder of the proof, assume that $v \neq v'$. Thus $G' = \text{trade}(G, C)$ for some C in $\text{cycles}(G)$ such that agent u is on C . Let u' denote $\text{agent}(G, v')$. Since no agents in $\text{frozen}(G)$ appear on C , neither u nor u' belongs to $\text{frozen}(G)$. Since v' is equal to $\text{next}(G, u)$, we conclude that v' belongs to $\Gamma(G, u)$. Since v' belongs to $\Gamma(G, u)$, u' does not belong to $\text{frozen}(G)$, and G belongs to $\text{admissible}(W)$, we conclude that $\text{bottom}(W, G, u, v')$ holds. If u belongs to $\text{unsatisfied}(G)$, then since G belongs to $\text{configs}(W)$ and $\text{bottom}(W, G, u, v')$ holds, we have $v' \succ_u v$. If u belongs to $\text{satisfied}(G)$, then v belongs to $\Gamma(G, u)$. Since v belongs to $\Gamma(G, u)$, u does not belong to $\text{frozen}(G)$, and G belongs to $\text{admissible}(W)$, we conclude that $\text{bottom}(W, G, u, v)$ holds, and hence that $v \sim_u v'$. \square

E Containment in the GATTC Family

The purpose of this section is to establish that the family of mechanisms defined by Algorithm 1 is contained in the GATTC family of Aziz and de Keijzer [3]. The GATTC family was designed to include the TTAS family of Alcalde-Unzu and Molis [2] and the TCR family of Jaramillo and Manjunath [6] as special cases. As discussed in Section 1, the family of mechanisms defined by Algorithm 1 is closely related to the TTAS and TCR families. As such it is not surprising that this family is also contained in the GATTC family. Indeed, no major ideas are needed to carry out the proof of this claim. The main issue is merely to reconcile the notations of the present paper with those of [3].

We begin by presenting Algorithm GATTC [3], which defines the GATTC family. The notations used in the description, which are formally defined in [3], are clarified below. The algorithm maintains a digraph; we use program variable H to refer to this digraph. The initial digraph is defined in terms of the given housing market instance in the same manner as in the original TTC mechanism. The following steps are repeated until digraph H is empty.

1. Repeat the following a finite number of times on H :
 - (a) Either implement a non-good cycle (if H is not empty), or do nothing.
 - (b) Either remove a paired-symmetric absorbing set and adjust H (if a paired-symmetric absorbing set exists), or do nothing.
2. Repeatedly remove paired-symmetric absorbing sets and adjust H , until there are no paired-symmetric absorbing sets in H .
3. If H is not empty, implement a good cycle.

Our goal is to prove that for any housing markets instance I , any execution of Algorithm 1 on I produces the same allocation as some execution of Algorithm GATTC on I . We will actually prove a stronger version of this statement with Algorithm GATTC replaced by Algorithm GATTC', where Algorithm GATTC' corresponds to the special case of Algorithm GATTC in which step 1 above is always executed zero times. Thus Algorithm GATTC' repeats the following two steps until digraph H is empty.

1. Repeatedly remove paired-symmetric absorbing sets and adjust H , until there are no paired-symmetric absorbing sets in H .
2. If H is not empty, implement a good cycle.

The digraph H of Algorithm GATTC' plays essentially the same role as the configuration G in Algorithm 1. The digraph H has a vertex for each remaining (i.e., non-removed) house and agent. Each house in H has a directed edge to its tentative owner. Each agent in H has directed edges to its top preferences amongst the remaining houses. A paired-symmetric absorbing set A is a strongly connected component of H with no outgoing edges, and such that each agent in A is tentatively assigned to one of its top choices. In a corresponding configuration G , all of the agents in a paired-symmetric absorbing set belong to $frozen(G)$. (The notion of a “corresponding” configuration is formalized in Section E.1; see the definition of the term “compatible”.)

The notion of “adjusting” the digraph H is defined in the same way as in the original TTC mechanism, that is, for each agent u with no outgoing edges to any of the remaining houses, we adjust u by adding outgoing edges from u to its top preferences amongst the remaining houses.

The notion of “implementing a cycle C ” is the same as in the original TTC algorithm, which in turn coincides with our $trade(G, C)$ formalism. A cycle is considered to be “good” if at least one of the agents on the cycle is not tentatively assigned to one of its top preferences. In a corresponding configuration G , this means that some agent on C belongs to $unsatisfied(G)$.

We now state the main result of this section.

Lemma E.1. *Let I be an arbitrary housing market instance. Then any allocation computed by some execution of Algorithm 1 on instance I is also computed by some execution of Algorithm GATTC' on instance I .*

Proof. By Lemma E.5, which is proven in Section E.1 below, there exist executions of Algorithms 1 and GATTC' on instance I that produce the same allocation. The claim of the lemma then follows from Theorem 1. □

E.1 Relating Algorithms 1 and GATTC'

The purpose of this section is to establish Lemma E.5, which is used in the proof of Lemma E.1 above. To simplify our notations, throughout this section we fix an arbitrary housing market instance I with associated wpp W .

For any nonnegative integer k , we define S_k as the set of all configurations that can be reached by performing exactly k iterations of the main loop of Algorithm 1 on instance I . (If an execution terminates before performing k iterations, it does not contribute a configuration to S_k .)

We view the *state* of an execution of Algorithm GATTC' on instance I as being comprised of the digraph associated with program variable H and the set of removed agent-house pairs. We define the *rank* of such a state as the number of agents in H . Algorithm GATTC' terminates when a state of rank zero is reached, at which point the output allocation is given by the set of removed agent-house pairs.

We say that state σ of an execution of Algorithm GATTC' on instance I is *compatible* with configuration G if the set of removed agents in σ is contained in $frozen(G)$, each removed agent in σ is assigned to the same house as in G , and the digraph H associated with σ is equal to the

subgraph of configuration G induced by the agents that do not belong to $frozen(G)$ and their tentatively allocated houses (i.e., the non-removed agents and houses in σ).

For any nonnegative integer k , we define T_k as the set of all states σ reached during some execution of Algorithm GATTC' on instance I such that σ is compatible with a configuration in S_k .

Lemma E.2. *The initial state of an execution of Algorithm GATTC' on instance I belongs to T_k for some nonnegative integer k .*

Proof. In the initial state of an execution of Algorithm GATTC', call it σ , each agent in H has edges to its top preferences, and each house in H has an edge to its initial owner. In the initial configuration G of an execution of Algorithm 1, each house has an edge to its initial owner, and each agent has no outgoing edges. Thus each agent belongs to $exhausted(G)$, and a revelation action can be applied to each agent, i.e., for each agent u , we can assign G to $reveal(W, G, u)$. The configuration obtained by applying these revelation actions does not depend on the order in which they are performed. The resulting configuration G belongs to S_k where k is equal to the number of agents, and is compatible with σ . Thus σ belongs to T_k . \square

Lemma E.3. *Let k be a nonnegative integer and let σ be a state in T_k with positive rank. Then there exists a state σ' such that at least one of the following two conditions holds: (1) σ' belongs to T_k and the rank of σ' is less than that of σ , or (2) σ' belongs to $T_{k+\ell}$ for some $\ell > 0$ and the rank of σ' is equal to that of σ .*

Proof. Let H be the digraph associated with state σ . Let G be a configuration in S_k that is compatible with σ .

Case 1: Either there is a paired-symmetric absorbing set in H , or some agent in H has no outgoing edges. Fix an arbitrary execution E of Algorithm GATTC' on instance I that passes through state σ .

Case 1.1: Immediately after reaching state σ , execution E removes a paired-symmetric absorbing set, thereby reaching state σ' . Thus the rank of σ' is less than that of σ . It is straightforward to verify that σ' is compatible with G . Hence σ' belongs to T_k , and we conclude that Condition (1) in the statement of the lemma is satisfied.

Case 1.2: Immediately after reaching state σ , execution E adjusts some agent u . Since G is compatible with σ and u has no outgoing edges in H , all of u 's outgoing edges in G go to houses assigned to agents in $frozen(G)$. Since Lemma 3.5 implies that $distance(G, v) = \infty$ for every such house v , we conclude that $next(G, u) = nil$. Since G is compatible with σ and u is an agent in H , we conclude that u does not belong to $frozen(G)$, and hence that u belongs to $unsatisfied(G)$. Since $next(G, u) = nil$ and u belongs to $unsatisfied(G)$, we find that u belongs to $exhausted(G)$. Now define a positive integer ℓ and a sequence of configurations G_1, \dots, G_ℓ as follows. Define G_1 as $reveal(W, G, u)$. If u does not belong to $exhausted(G_1)$, then define ℓ to be 1. Otherwise, define G_2 as $reveal(W, G_1, u)$ and if u does not belong to $exhausted(G_2)$, define ℓ to be 2. Otherwise, define G_3 as $reveal(W, G_2, u)$, and so on, until we reach G_i such that u does not belong to $exhausted(G_i)$ (it is easy to see that such an i exists), at which point we define ℓ to be i . It is straightforward to verify that G_ℓ is compatible with σ' . Since G_ℓ belongs to $S_{k+\ell}$, we conclude that σ' belongs to $T_{k+\ell}$. Since σ and σ' are of equal rank, Condition (2) in the statement of the lemma is satisfied.

Case 2: There is no paired-symmetric absorbing set in H , and every agent in H has at least one outgoing edge. Let U_0 denote the set of removed agents in state σ . Since G is compatible with σ , $frozen(G)$ contains U_0 . Furthermore, if $frozen(G)$ properly contains U_0 , then it is straightforward to argue that there are one or more paired-symmetric absorbing sets in H , contradicting the case condition. We conclude that $frozen(G)$ is equal to U_0 . Since every agent in H has at least one outgoing edge, G is compatible with σ , and $frozen(G)$ is equal to U_0 , we deduce from Lemma 3.5 that for each agent u in $unsatisfied(G)$ there is at least one edge from u to a house v such that $distance(G, v)$ is finite, and hence that $next(G, u) \neq nil$. It follows that $exhausted(G)$ is empty. Since σ has positive rank, G is compatible with σ , and $exhausted(G)$ is empty, we deduce from Lemma 3.6 that $cycles(G)$ is nonempty. Let C be a cycle in $cycles(G)$. Thus the configuration $G' = trade(G, C)$ belongs to S_{k+1} . Lemma 3.2 implies that at least one agent on C belongs to $unsatisfied(G)$. Since G is compatible with σ , $frozen(G)$ is equal to U_0 , C belongs to $cycles(G)$, and C contains at least one agent in $unsatisfied(G)$, we deduce that C is a good cycle in H . The case condition implies that from state σ , an execution of Algorithm GATTC' can proceed by trading along cycle C ; let σ' denote the resulting state. It is straightforward to prove that state σ' is compatible with configuration G' . Hence σ' belongs to T_{k+1} . Furthermore, the rank of σ' is equal to that of σ . Thus Condition (2) in the statement of the lemma holds. \square

Lemma E.4. *There is a nonnegative integer k such that T_k contains a state with rank zero.*

Proof. Let σ_0 denote the initial state of Algorithm GATTC'. By Lemma E.2, state σ_0 belongs to T_{k_0} for some nonnegative integer k_0 . If σ_0 has rank zero, the lemma holds. Otherwise, by applying Lemma E.3, we know that there exists a state a nonnegative integer k_1 and a state σ_1 in T_{k_1} such that at least one of the following two conditions holds: (1) $k_1 = k_0$ and the rank of σ_1 is less than that of σ_0 , or (2) $k_1 > k_0$ and the rank of σ_1 is equal to that of σ_0 . If σ_1 has rank zero, the lemma holds. Otherwise we continue applying Lemma E.3 in a similar manner to obtain a sequence of states σ_i in T_{k_i} until we reach a state σ_i with rank zero. Termination is guaranteed since Lemma 4.3 implies that S_k is empty for all sufficiently large k , and hence there is a finite upper bound on the k_i 's, thereby bounding the number of times that Condition (2) can occur in our sequence of applications of Lemma E.3. \square

Lemma E.5. *There exist executions of Algorithms 1 and GATTC' on instance I that produce the same allocation.*

Proof. Lemma E.4 implies that there is a nonnegative integer k such that T_k contains a state σ with rank zero. Let G be a configuration in S_k that is compatible with σ . Thus every agent in G belongs to $frozen(G)$, and hence G is final. Furthermore, since G is compatible with σ , the allocation associated with σ is equal to $allocation(G)$. \square

F A Fast Implementation

In this section we describe an $O(n^3)$ -time deterministic algorithm, Algorithm 2. It will be evident that any execution of Algorithm 2 corresponds to a possible execution of Algorithm 1. Thus Theorem 1 implies that Algorithm 2 has the same input-output behavior as Algorithm 1.

Let $W = (U, V, \succ)$ be a wpp and let G be a non-final configuration in $admissible(W)$. Below we describe an $O(|V|^2)$ time subroutine to compute a configuration G' in $\Gamma^*(W, G)$ such that

$|unsatisfied(G')| < |unsatisfied(G)|$. By Lemma 4.6, we find that G' belongs to $admissible(W)$. Thus we can iteratively apply this subroutine at most $|V|$ times, yielding an overall time bound of $O(|V|^3)$.

The $O(|V|^2)$ -time subroutine works in three phases, as follows. In the first phase, we use a breadth-first traversal in the reversal of G (i.e., the graph obtained by reversing the direction of the edges in G), starting from the set of agents in $unsatisfied(G)$ (which are easy to identify), to compute $distance(G, v)$ for all houses v in V , and $next(G, u)$ for all agents u in U . Then we identify the set of agents in $exhausted(G)$, call it U' , and the set of all houses v such that $distance(G, v)$ is finite, call it V' . The time complexity of the first phase is easily seen to be $O(|V|^2)$.

In the second phase, we process the agents in U' in arbitrary order. To process such an agent u , we repeatedly update G to $reveal(W, G, u)$ until $\Gamma(G, u) \cap V'$ is nonempty. Each such update merely involves adding some new outgoing edges to agent u . The total number of new edges added to u is at most $|V|$, and the time complexity of processing u is $O(|V|)$. Once u has been processed, we check whether u now belongs to $satisfied(G)$. If so, we can terminate the subroutine and take the desired configuration G' to be the current configuration G . (As an optimization, we can go ahead and process any remaining agents in U' before terminating the subroutine.) If we process all of the agents in U' without arriving at a configuration G' such that $|unsatisfied(G')| < |unsatisfied(G)|$, then we proceed to the third phase. Since at most $|V|$ agents are processed in the second phase, the time complexity of the second phase is $O(|V|^2)$.

In the third phase, since $exhausted(G)$ is empty and G is not final (since all of the agents in U' continue to belong to $unsatisfied(G)$), Lemma 3.6 implies that $cycles(G)$ is nonempty. As in a standard TTC algorithm for the simple case of strict preferences, we can easily identify all of the cycles in $cycles(G)$ in $O(|V|)$ time. We can then update G to $trade(G, C)$ for some C in $cycles(G)$. As argued in the proof of Lemma 4.3, such an update yields a configuration G' such that $|unsatisfied(G')| < |unsatisfied(G)|$. The time complexity of the third phase is $O(|V|)$. (As an optimization, we can process all of the cycles in $cycles(G)$; the time complexity of the third phase remains $O(|V|)$.)