

The core of housing markets from an agent’s perspective: Is it worth sprucing up your home?*

Ildikó Schlotter^{1,2}, Péter Biró^{1,3}, and Tamás Fleiner^{1,2}

¹ Centre for Economic and Regional Studies, Budapest, Hungary,
{schlotter.ildiko, biro.peter, fleiner.tamas}@krtk.hu

² Budapest University of Technology and Economics, Budapest, Hungary

³ Corvinus University of Budapest, Budapest, Hungary

Abstract. We study housing markets as introduced by Shapley and Scarf [38]. We investigate the computational complexity of various questions regarding the situation of an agent a in a housing market H : we show that it is NP-hard to find an allocation in the core of H where (i) a receives a certain house, (ii) a does not receive a certain house, or (iii) a receives a house other than her own. We prove that the core of housing markets *respects improvement* in the following sense: given an allocation in the core of H where agent a receives a house h , if the value of the house owned by a increases, then the resulting housing market admits an allocation where a receives either h , or a house that she prefers to h ; moreover, such an allocation can be found efficiently. We further show an analogous result in the STABLE ROOMMATES setting by proving that stable matchings in a one-sided market also respect improvement.

1 Introduction

Housing markets is a classic model in economics where agents are initially endowed with one unit of an indivisible good, called a *house*, and agents may trade their houses according to their preferences without using monetary transfers. In such markets, trading results in a reallocation of houses in a way that each agent ends up with exactly one house. Motivation for studying housing markets comes from applications such as kidney exchange [35,8,12] and housing programs [1,43].

In their seminal work Shapley and Scarf [38] examined housing markets where agents’ preferences are weak orders. They proved that such markets always admit a *core* allocation, that is, an allocation where no coalition of agents can strictly improve their situation by trading only among themselves. They also described the Top Trading Cycles (TTC) algorithm, proposed by David Gale, and proved that the set of allocations that can be obtained through the TTC algorithm coincides with the set of competitive allocations; hence the TTC always produces an allocation in the core. When preferences are strict, the TTC produces the unique allocation in the *strict core*, that is, an allocation where no coalition of agents can weakly improve their situation by trading among themselves [34].

* Supported by the Hungarian Academy of Sciences (Momentum Programme LP2021-2) and the Hungarian Scientific Research Fund (NFKIH grants K128611, K124171).

Although the core of housing markets has been the subject of considerable research, there are still many challenges which have not been addressed. Consider the following question: given an agent a and a house h , does there exist an allocation in the core where a obtains h ? Or one where a does not obtain h ? Can we determine whether a may receive a house better than her own in some core allocation? Similar questions have been extensively studied in the context of the STABLE MARRIAGE and the STABLE ROOMMATES problems [31,23,21,22,20], but have not yet been considered in relation to housing markets.

Even less is known about the core of housing markets in cases where the market is not static. Although some researchers have addressed certain dynamic models, most of these either focus on the possibility of repeated allocation [34,28,29], or consider a situation where agents may enter and leave the market at different times [42,13,32]. Recently, Biró et al. [9] have investigated how a change in the preferences of agents affects the housing market. Namely, they considered how an improvement of the house belonging to agent a affects the situation of a . Following their lead, we aim to answer the following question: if the value of the house belonging to agent a increases, how does this affect the core of the market from the viewpoint of a ? Is such a change bound to be beneficial for a , as one would expect? This question is of crucial importance in the context of kidney exchange: if procuring a new donor with better properties (e.g., a younger or healthier donor) does not necessarily benefit the patient, then this could undermine the incentive for the patient to find a donor with good characteristics, damaging the overall welfare.

1.1 Our contribution

We consider the computational complexity of deciding whether the core of a housing market contains an allocation where a given agent a obtains a certain house. In Theorem 1 we prove that this problem is NP-complete, as is the problem of finding a core allocation where a does *not* receive a certain house. Even worse, it is already NP-complete to decide whether a core allocation can assign *any* house to a other than her own. Various generalizations of these questions can be answered efficiently in both the STABLE MATCHING and STABLE ROOMMATES settings [31,23,21,22,20], so we find these intractability results surprising.

Instead of asking for a core allocation where a given agent can trade her house, one can also look at the optimization problem which asks for an allocation in the core with the maximum number of agents involved in trading. This problem is known to be NP-complete [18]. We show in Theorem 2 that for any $\varepsilon > 0$, approximating this problem with ratio $|N|^{1-\varepsilon}$ for a set N of agents is NP-hard. We complement this strong inapproximability result in Proposition 3 by pointing out that a trivial approach yields an approximation algorithm with ratio $|N|$.

Turning our attention to the question of how an increase in the value of a house affects its owner, we show the following result in Theorem 4. If the core of a housing market contains an allocation where a receives a house h , and the market changes in a way such that some agents perceive an increased value for the house owned by a (and nothing else changes in the market), then the resulting

housing market admits an allocation in its core where a receives either h or a house that she prefers to h . We prove this in a constructive way, by presenting an algorithm that finds such an allocation. This settles an open question by Biró et al. [9] who ask whether the core *respects improvement* in the sense that the best allocation achievable for an agent a in a core allocation can only (weakly) improve for a as a result of an increase in the value of a 's house.

It is clear that an increase in the value of a 's house may not always yield a *strict* improvement for a (as a trivial example, some core allocation may assign a her top choice even before the change), but one may wonder if we can efficiently determine when a strict improvement for a becomes possible. This problem turns out to be closely related to the question whether a can obtain a given house in a core allocation; in fact, we were motivated to study the latter problem by our interest in determining the possibilities for a strict improvement. Although one can formulate several variants of the problem depending on what exactly one considers to be a strict improvement, by Theorem 11 each of them leads to computational intractability (NP-hardness or coNP-hardness).

Finally, we also answer a question raised by Biró et al. [9] regarding the property of respecting improvements in the context of the STABLE ROOMMATES problem. An instance of STABLE ROOMMATES contains a set of agents, each having preferences over the other agents; the usual task is to find a matching between the agents that is *stable*, i.e., no two agents prefer each other to their partners in the matching. It is known that a stable matching need not always exist, but if it does, then Irving's algorithm [26] finds one efficiently. In Theorem 14 we show that if some stable matching assigns agent a to agent b in a STABLE ROOMMATES instance, and the valuation of a increases (that is, if she moves upward in other agents' preferences, with anything else remaining constant), then the resulting instance admits a stable matching where a is matched either to b or to an agent she prefers to b . This result is a direct analog of the one stated in Theorem 4 for the core of housing markets; however, the algorithm we propose in order to prove it uses different techniques.

We remark that we use a model with partially ordered preferences (a generalization of weak orders), and describe a linear-time implementation of the TTC algorithm in such a model.

1.2 Related work

Most works relating to the core of housing markets aim for finding core allocations with some additional property that benefits global welfare, most prominently Pareto optimality [27,4,5,33,37]. Another line of research comes from kidney exchange where the length of trading cycles is of great importance and often plays a role in agents' preferences [19,16,15,7,17] or is bounded by some constant [2,10,11,25,18]. None of these papers deal with problems where a core allocation is required to fulfill some constraint regarding a given agent or set of agents—that they be trading, or that they obtain (or not obtain) a certain house. Nevertheless, some of them focus on finding a core allocation where the number of agents involved in trading is as large as possible. Cechlárová and Repiský [18]

proved that this problem is NP-hard in the classical housing market model, while Biró and Cechlárová [7] considered a special model where agents care first about the house they receive and after that about the length of their trading cycle (shorter being better); they prove that for any $\varepsilon > 0$, it is NP-hard to approximate the number of agents trading in a core allocation with a ratio $|N|^{1-\varepsilon}$ (where N is the set of agents).

The property of respecting improvement has first been studied in a paper by Balinski and Sönmez [6] on college admission, who proved that the student-optimal stable matching algorithm respects the improvement of students, so a better test score for a student always results in an outcome weakly preferred by the student (assuming other students' scores remain the same). Hatfield et al. [24] contrasted their findings by showing that no stable mechanism respects the improvement of school quality. Sönmez and Switzer [39] applied the model of matching with contracts to the problem of cadet assignment in the United States Military Academy, and have proved that the cadet-optimal stable mechanism respects improvement of cadets. Recently, Klaus and Klijn [30] have obtained results of a similar flavor in a school-choice model with minimal-access rights.

Roth et al. [36] deal with the property of respecting improvement in connection to kidney exchange: they show that in a setting with dichotomous preferences and pairwise exchanges priority mechanisms are donor monotone, meaning that a patient can only benefit from bringing an additional donor on board. Biró et al. [9] focus on the classical Shapley-Scarf model and investigate how different solution concepts behave when the value of an agent's increases. They prove that both the strict core and the set of competitive allocations satisfy the property of respecting improvements, however, this is no longer true when the lengths of trading cycles are bounded by some constant.

2 Preliminaries

Preferences as partial orders. In the majority of the existing literature, preferences of agents are usually considered to be either strict or, if the model allows for indifference, weak linear orders. Weak orders can be described as lists containing *ties*, a set of alternatives considered equally good for the agent. Partial orders are a generalization of weak orders that allow for two alternatives to be *incomparable* for an agent. Incomparability may not be transitive, as opposed to indifference in weak orders. Formally, an (irreflexive)⁴ *partial ordering* \prec on a set of alternatives is an irreflexive, antisymmetric and transitive relation.

Partially ordered preferences arise by many natural reasons; we give two examples motivated by kidney exchanges. For example, agents may be indifferent between goods that differ only slightly in quality. Indeed, recipients might be indifferent between two organs if their expected graft survival times differ by less than one year. However, small differences may add up to a significant contrast:

⁴ Throughout the paper we will use the term *partial ordering* in the sense of an irreflexive (or strict) partial ordering.

an agent may be indifferent between a and b , and also between b and c , but strictly prefer a to c . Partial preferences also emerge in multiple-criteria decision making. The two most important factors for estimating the quality of a kidney transplant are the HLA-matching between donor and recipient, and the age of the donor.⁵ An organ is considered better than another if it is better with respect to both of these factors, leading to partial orders.

Housing markets. Let $H = (N, \{\prec_a\}_{a \in N})$ be a *housing market* with agent set N and with the preferences of each agent $a \in N$ represented by a partial ordering \prec_a of the agents. For agents a , b , and c , we interpret $a \prec_c b$ as agent c *preferring* the house owned by agent b to the house of agent a . We will write $a \preceq_c b$ as equivalent to $b \not\prec_c a$, and we write $a \sim_c b$ if $a \not\prec_c b$ and $b \not\prec_c a$. We say that agent a finds the house of b *acceptable*, if $a \preceq_a b$, and we denote by $A(a) = \{b \in N : a \preceq_a b\}$ the set of agents whose house is acceptable for a . We define the *acceptability graph* of the housing market H as the directed graph $G^H = (N, E)$ with $E = \{(a, b) \mid b \in A(a)\}$; we let $|G^H| = |N| + |E|$. Note that $(a, a) \in E$ for each $a \in N$. The *submarket* of H on a set $W \subseteq N$ of agents is the housing market $H_W = (W, \{\prec_a^{|W}\}_{a \in W})$ where $\prec_a^{|W}$ is the partial order \prec_a restricted to W ; the acceptability graph of H_W is the subgraph of G^H induced by W , denoted by $G^H[W]$. For a set W of agents, let $H - W$ be the submarket $H_{N \setminus W}$ obtained by *deleting* W from H ; for $W = \{a\}$ we may write simply $H - a$.

For a set $X \subseteq E$ of arcs in G^H and an agent $a \in N$ we let $X(a)$ denote the set of agents b such that $(a, b) \in X$; whenever $X(a)$ is a singleton $\{b\}$ we will abuse notation by writing $X(a) = b$. We also define $\delta_X^-(a)$ and $\delta_X^+(a)$ as the number of in-going and out-going arcs of a in X , respectively. For a set $W \subseteq N$ of agents, we let $X[W]$ denote the set of arcs in X that run between agents of W .

We define an *allocation* X in H as a subset $X \subseteq E$ of arcs in G^H such that $\delta_X^-(a) = \delta_X^+(a) = 1$ for each $a \in N$, that is, X forms a collection of cycles in G^H containing each agent exactly once. Then $X(a)$ denotes the agent whose house a obtains according to allocation X . If $X(a) \neq a$, then a is *trading* in X . For allocations X and X' , we say that a *prefers* X to X' if $X'(a) \prec_a X(a)$.

For an allocation X in H , an arc $(a, b) \in E$ is *X -augmenting*, if $X(a) \prec_a b$. We define the *envy graph* $G_{X \prec}^H$ of X as the subgraph of G^H containing all X -augmenting arcs. A *blocking cycle* for X in H is a cycle in $G_{X \prec}^H$, that is, a cycle C where each agent a on C prefers $C(a)$ to $X(a)$. An allocation X is contained in the *core* of H , if there does not exist a blocking cycle for it, i.e., if $G_{X \prec}^H$ is acyclic. A *weakly blocking cycle* for X is a cycle C in G^H where $X(a) \preceq_a C(a)$ for each agent a on C and $X(a) \prec_a C(a)$ for at least one agent a on C . The *strict core* of H contains allocations that do not admit weakly blocking cycles.

Organization. Section 3 contains an adaptation of the TTC algorithm for partially ordered preferences, followed by our results on finding core allocations with various arc restrictions and on maximizing the number of agents involved in trading. In Section 4 we present our results on the property of respecting

⁵ In fact, these are the two factors for which acceptability thresholds can be set by the patients in the UK program [8].

improvements in relation to the core of housing markets, including our main technical result, Theorem 4. In Section 5 we study the respecting improvement property in the context of the STABLE ROOMMATES problem. We conclude with some questions for future research in Section 6.

3 The core of housing markets: some computational problems

We investigate a few computational problems related to the core of housing markets. In Section 3.1 we describe our adaptation of TTC to partially ordered preferences. In Section 3.2 we turn our attention to the problem of finding an allocation in the core of a housing market that satisfies certain arc restrictions, requiring that a given arc be contained or, just the opposite, not be contained in the desired allocation. In Section 3.3 we look at the most prominent optimization problem in connection to the core: given a housing market, find an allocation in its core where the number of agents that are trading is as large as possible.

3.1 Top Trading Cycles for preferences with incomparability

Strict preferences. If agents’ preferences are represented by strict orders, then the TTC algorithm [38] produces the unique allocation in the strict core. TTC creates a directed graph D where each agent a points to her top choice, that is, to the agent owning the house most preferred by a . In the graph D each agent has out-degree exactly 1, since preferences are assumed to be strict. Hence, D contains at least one cycle, and moreover, the cycles in D do not intersect. TTC selects all cycles in the graph D as part of the desired allocation, deletes from the market all agents trading along these cycles, and repeats the whole process until there are no agents left.

Preferences as partial orders. When preferences are represented by partial orders, one can modify the TTC algorithm by letting each agent a in D point to her *undominated* choices: b is undominated for a , if there is no agent c such that $b \prec_a c$. Notice that an agent’s out-degree is then *at least* 1. Thus, D contains at least one cycle, but in case it contains more than one cycle, these may overlap.

A simple approach is to select a set of mutually vertex-disjoint cycles in each round, removing the agents trading along them from the market and proceeding with the remainder in the same manner. It is not hard to see that this approach yields an algorithm that produces an allocation in the core: by the definition of undominated choices, any arc of a blocking cycle leaving an agent a necessarily points to an agent that was already removed from the market at the time when a cycle containing a got selected. Clearly, no cycle may consist of such “backward” arcs only, proving that the computed allocation is indeed in the core.

Implementation in linear time. Abraham et al. [3] describe an implementation of the TTC algorithm for strict preferences that runs in $O(|G^H|)$ time. We extend their ideas to the case when preferences are partial orders as follows.

For each agent $a \in N$ we assume that a 's preferences are given using a *Hasse diagram* which is a directed acyclic graph H_a that can be thought of as a compact representation of \prec_a . The vertex set of H_a is $A(a)$, and it contains an arc (b, c) if and only if $b \prec_a c$ and there is no agent c' with $b \prec_a c' \prec_a c$. Then the description of our housing market H has length $\sum_{a \in A} |H_a|$ which we denote by $|H|$. If preferences are weak or strict orders, then $|H| = O(|G^H|)$.

Throughout our variant of TTC, we will maintain a list $U(a)$ containing the undominated choices of a among those that still remain in the market, as well as a subgraph D of G^H spanned by all arcs (a, b) with $b \in U(a)$. Furthermore, for each agent a in the market, we will keep a list of all occurrences of a as someone's undominated choice. Using H_a we can find the undominated choices of a in $O(|H_a|)$ time, so initialization takes $O(|H|)$ time in total.

Whenever an agent a is deleted from the market, we find all agents b such that $a \in U(b)$, and we update $U(b)$ by replacing a with its in-neighbors in H_b . Notice that the total time required for such deletions (and the necessary replacements) to maintain $U(b)$ is $O(|H_b|)$. Hence, we can efficiently find the undominated choices of each agent at any point during the algorithm, and thus traverse the graph D consisting of arcs (a, b) with $b \in U(a)$.

To find a cycle in D , we simply keep building a path using arcs of D , until we find a cycle (perhaps a loop). After recording this cycle and deleting its agents from the market (updating the lists $U(a)$ as described above), we simply proceed with the last agent on our path. Using the data structures described above the total running time of our variant of TTC is $O(|N| + \sum_{a \in N} |H_a|) = O(|H|)$.

3.2 Allocations in the core with arc restrictions

We now focus on the problem of finding an allocation in the core that fulfills certain arc constraints. The simplest such constraints arise when we require a given arc to be included in, or conversely, be avoided by the desired allocation.

We define the ARC IN CORE problem as follows: given a housing market $H = (N, \{\prec_a\}_{a \in N})$ and an arc (a, b) in G^H , decide whether there exists an allocation in the core of H that contains (a, b) , or in other words, where agent a obtains the house of agent b . Analogously, the FORBIDDEN ARC IN CORE problem asks to decide if there exists an allocation in the core of H *not* containing (a, b) .

By giving a reduction from ACYCLIC PARTITION [14], we show in Theorem 1 that both of these problems are computationally intractable, even if agents have a strict ordering over the houses. In fact, we cannot even hope to decide for a given agent a in a housing market H whether there exists an allocation in the core of H where a is trading; we call this problem AGENT TRADING IN CORE.

Theorem 1 (\star^6). *Each of the following problems is NP-complete, even if agents' preferences are strict orders:*

- ARC IN CORE,
- FORBIDDEN ARC IN CORE, and
- AGENT TRADING IN CORE.

⁶ Proofs marked by an asterisk can be found in Appendix A.

3.3 Maximizing the number of agents trading in a core allocation

Perhaps the most natural optimization problem related to the core of housing markets is the following: given a housing market H , find an allocation in the core of H whose *size*, defined as the number of trading agents, is maximal among all allocations in the core of H ; we call this the MAX CORE problem. MAX CORE is NP-hard by a result of Ceclárová and Repiský [18]. In Theorem 2 below we show that even approximating MAX CORE is NP-hard. Our result is tight in the following sense: we prove that for any $\varepsilon > 0$, approximating MAX CORE with a ratio of $|N|^{1-\varepsilon}$ is NP-hard, where $|N|$ is the number of agents in the market. By contrast, a very simple approach yields an approximation with ratio $|N|$.

We remark that Biró and Ceclárová [7] proved a similar inapproximability result, but since they considered a special model where agents not only care about the house they receive but also about the length of their exchange cycle, their result cannot be translated to our model, and so does not imply Theorem 2. Instead, our reduction relies on the ideas we use to prove Theorem 1.

Theorem 2 (\star). *For any constant $\varepsilon > 0$, the MAX CORE problem is NP-hard to approximate within a ratio of $\alpha_\varepsilon(N) = |N|^{1-\varepsilon}$ where N is the set of agents, even if agents' preferences are strict orders.*

We contrast Theorem 2 with the observation that an algorithm that outputs *any* allocation in the core yields an approximation for MAX CORE with ratio $|N|$.

Proposition 3 (\star). *MAX CORE can be approximated with a ratio of $|N|$ in polynomial time, where $|N|$ is the number of agents in the input.*

4 The effect of improvements in housing markets

Let $H = (N, \{\prec_a\}_{a \in N})$ be a housing market containing agents p and q . We consider a situation where the preferences of q are modified by “increasing the value” of p for q without altering the preferences of q over the remaining agents. If the preferences of q are given by a strict or weak order, then this translates to *shifting* the position of p in the preference list of q towards the top. Formally, a housing market $H' = (N, \{\prec'_a\}_{a \in N})$ is called a (p, q) -*improvement* of H , if $\prec_a = \prec'_a$ for any $a \in N \setminus \{q\}$, and \prec'_q is such that (i) $a \prec'_q b$ iff $a \prec_q b$ for any $a, b \in N \setminus \{p\}$, and (ii) if $a \prec_q p$, then $a \prec'_q p$ for any $a \in N$. We will also say that a housing market is a p -*improvement* of H , if it can be obtained by a sequence of (p, a_i) -improvements for a series a_1, \dots, a_k of agents for some $k \in \mathbb{N}$.

To examine how p -improvements affect the situation of p in the market, one may consider several solution concepts such as the core, the strict core, and so on. We regard a solution concept as a function Φ that assigns a set of allocations to each housing market. Based on the preferences of p , we can compare allocations in Φ . Let $\Phi_p^+(H)$ denote the set containing the best houses p can obtain in $\Phi(H)$:

$$\Phi_p^+(H) = \{X(p) \mid X \in \Phi(H), \forall X' \in \Phi(H) : X'(p) \preceq_p X(p)\}.$$




Fig. 1. The housing markets H and H' in the proof of Proposition 6. Here and everywhere else we depict markets through their acceptability graphs with all loops omitted; preferences are indicated by numbers along the arcs. For both H and H' , the allocation represented by red (and bold) arcs yields the worst possible outcome for p in any core allocation of the given market.

Similarly, let $\Phi_p^-(H)$ be the set containing the worst houses p can obtain in $\Phi(H)$.

Following the notation used by Biró et al. [9], we say that Φ *respects improvement for the best available house* or simply *satisfies the RI-best property*, if for any housing markets H and H' such that H' is a p -improvement of H for some agent p , $a \preceq_p a'$ for every $a \in \Phi_p^+(H)$ and $a' \in \Phi_p^+(H')$. Similarly, Φ *respects improvement for the worst available house* or simply *satisfies the RI-worst property*, if for any housing markets H and H' such that H' is a p -improvement of H for some agent p , $a \preceq_p a'$ for every $a \in \Phi_p^-(H)$ and $a' \in \Phi_p^-(H')$.

Notice that the above definition does not take into account the possibility that a solution concept Φ may become empty as a result of a p -improvement. To exclude such a possibility, we may require the condition that an improvement does not destroy all solutions. We say that Φ *strongly satisfies the RI-best (or RI-worst) property*, if besides satisfying the RI-best (or, respectively, RI-worst) property, it also guarantees that whenever $\Phi(H) \neq \emptyset$, then $\Phi(H') \neq \emptyset$ also holds where H' is a p -improvement of H for some agent p .

We prove that the core of housing markets satisfies the RI-best property. In fact, Theorem 4 (proved in Section 4.2) states a slightly stronger statement. By contrast, Proposition 6 shows that the core of housing markets violates the RI-worst property.

Theorem 4. *Given an allocation X in the core of the housing market H and a p -improvement H' of H , there exists an allocation X' in the core of H' such that either $X(p) = X'(p)$ or p prefers X' to X . Moreover, given H , H' and X , it is possible to find such an allocation X' in polynomial time.*

Corollary 5. *The core of housing markets strongly satisfies the RI-best property.*

Proposition 6. *The core of housing markets violates the RI-worst property.*

Proof. Let $N = \{a, b, c, p, q\}$ be the set of agents. The preferences indicated in Figure 1 define a housing market H and a (p, q) -improvement H' of H .

We claim that in every allocation in the core of H , agent p obtains the house of a . To see this, let X be an allocation where $(p, a) \notin X$. If agent a is not trading in X , then a and p form a blocking cycle; therefore, $(b, a) \in X$. Now, if $(c, b) \notin X$, then c and b form a blocking cycle for X ; otherwise, q and b form a blocking cycle for X . Hence, p obtains her top choice in all core allocations of H .

However, it is easy to verify that the core of H' contains an allocation where p obtains only her second choice (q 's house), as shown in Figure 1. \square

We describe our algorithm for Theorem 4 in Section 4.1, and prove its correctness in Section 4.2. In Section 4.3 we look at the problem of deciding whether a p -improvement leads to a situation strictly better for p .

4.1 Description of algorithm HM-Improve

Before describing our algorithm for Theorem 4, we need some notation.

Pre-allocations and their envy graphs. Given a housing market $H = (N, \{\prec_a\}_{a \in N})$ and two distinct agents u and v in N , we say that a set Y of arcs in $G^H = (N, E)$ is a *pre-allocation* from u to v in H , if

- $\delta_Y^-(u) = \delta_Y^+(v) = 0$,
- $\delta_Y^+(a) = 1$ for each $a \in N \setminus \{v\}$, and
- $\delta_Y^-(a) = 1$ for each $a \in N \setminus \{u\}$.

Note that Y is a collection of vertex-disjoint cycles and a unique path P in G^H , with P leading from u to v . We call u the *source* of Y and v its *sink*.

Given a pre-allocation Y from u to v in H , an arc $(a, b) \in E$ is *Y -augmenting*, if $a \neq v$ and $Y(a) \prec_a b$. We define the *envy graph* of Y as $G_{Y, \prec}^H = (N, E_Y)$ where E_Y is the set of Y -augmenting arcs in E . A blocking cycle for Y is a cycle in $G_{Y, \prec}^H$; notice that such a cycle cannot contain the sink v , since no Y -augmenting arc leaves v . We say that the pre-allocation Y is *stable*, if no blocking cycle exists for Y , that is, if its envy graph is acyclic.

We are now ready to propose an algorithm called HM-Improve that given an allocation X in the core of H outputs an allocation X' as required by Theorem 4. Observe that we can assume w.l.o.g. that H' is a (p, q) -improvement of H for some agent q , as we can apply such a single-agent version of Theorem 4 repeatedly to obtain the theorem for p -improvements involving multiple agents.

Algorithm HM-Improve. First, HM-Improve checks whether X belongs to the core of H' , and if so, outputs $X' = X$. Hence, we may assume that X admits a blocking cycle in H' . Observe such a cycle must contain the arc (q, p) , as otherwise it would block X in H as well. This implies that $X(q) \prec'_q p$.

HM-Improve proceeds by modifying the housing market: it adds a new agent \tilde{q} to H' , with \tilde{q} taking the place of p in the preferences of q ; the only house that agent \tilde{q} prefers to her own will be the house of p . Let \tilde{H} be the housing market obtained. Then the acceptability graph \tilde{G} of \tilde{H} can be obtained from the acceptability graph of H' by subdividing the arc (q, p) with a new vertex corresponding to agent \tilde{q} . Let $\tilde{N} = N \cup \{\tilde{q}\}$, and let \tilde{E} be the set of arcs in \tilde{G} .

Initialization. Let $Y = X \setminus \{(q, X(q))\} \cup \{(q, \tilde{q})\}$ in \tilde{G} . Observe that Y is a pre-allocation from the source $X(q)$ to the sink \tilde{q} in \tilde{H} . Additionally, we define a set R of *irrelevant* agents, initially empty. We may think of irrelevant agents as temporarily deleted from the market.

Iteration. Next, algorithm HM-Improve iteratively modifies the pre-allocation Y and the set R of irrelevant agents. It will maintain the property that Y is a pre-allocation in $\tilde{H} - R$; we denote its envy graph by $\tilde{G}_{Y \prec}$, having vertex set $\tilde{N} \setminus R$. While the source of Y changes during the iteration, the sink \tilde{q} remains fixed.

At each iteration, HM-Improve performs the following steps:

1. Let u be the source of Y . If $u \in \{p, \tilde{q}\}$, then the iteration stops.
2. Otherwise, if there exists a Y -augmenting arc (s, u) in $\tilde{G}_{Y \prec}$ entering u (note that $s \in \tilde{N} \setminus R$), then let $u' = Y(s)$. The algorithm modifies Y by deleting the arc (s, u') and adding the arc (s, u) to Y . Note that Y thus becomes a pre-allocation from u' to \tilde{q} in $\tilde{H} - R$.
3. Otherwise, no arc in $\tilde{G}_{Y \prec}$ enters u ; let $u' = Y(u)$. The algorithm adds u to the set R of irrelevant agents, and modifies Y by deleting the arc (u, u') . Again, Y becomes a pre-allocation from u' to \tilde{q} in $\tilde{H} - R$.

Output. Let Y be the pre-allocation at the end of the above iteration, u its source, and R the set of irrelevant agents. HM-Improve applies the variant of the TTC algorithm described in Section 3.1 to the submarket H'_R of H' when restricted to the set of irrelevant agents. Let X_R denote the obtained allocation in the core of H'_R . Then HM-Improve outputs an allocation X' defined as

$$X' = \begin{cases} X_R \cup Y \setminus \{(q, \tilde{q})\} \cup \{(q, p)\} & \text{if } u = p, \\ X_R \cup Y & \text{if } u = \tilde{q}. \end{cases}$$

4.2 Correctness of algorithm HM-Improve

We begin proving the correctness of algorithm HM-Improve with the following.

Lemma 7. *At each iteration, pre-allocation Y is stable in $\tilde{H} - R$.*

Proof. The proof is by induction on the number n of iterations performed. For $n = 0$, observe that initially $Y(a) = X(a)$ for each agent $a \in N \setminus \{q\}$, and by $X(q) \prec'_q p$ we know that q prefers $Y(q) = \tilde{q}$ to $X(q)$. Note also that neither (q, p) nor the arcs (q, \tilde{q}) and (\tilde{q}, p) are contained in the envy graph $\tilde{G}_{Y \prec}$. Thus, a cycle in $\tilde{G}_{Y \prec}$ would be present in the envy graph of X in H as well. Since X is in the core of H , it follows that Y is stable in \tilde{H} . Note that initially $R = \emptyset$.

For $n \geq 1$, assume that the algorithm has performed $n - 1$ iterations so far. Let Y and R be as defined at the beginning of the n -th iteration, with u being the source of Y , and let Y' and R' be the pre-allocation and the set of irrelevant agents obtained after the modifications in this iteration. Assume that Y is stable in $\tilde{H} - R$, so $\tilde{G}_{Y \prec}$ is acyclic. In case HM-Improve does not stop in Step 1 but modifies Y and possibly R , we distinguish between two cases:

- (a) the algorithm modifies Y in Step 2, by using a Y -augmenting arc (s, u) ; then $R' = R$. Note that s prefers Y' to Y , and for any other agent $a \in N \setminus R'$ we know $Y(a) = Y'(a)$. Hence, this modification amounts to deleting all arcs (s, a) from the envy graph $\tilde{G}_{Y \prec}$ where $Y(s) \prec_s a \preceq_s Y'(s)$.
- (b) the algorithm modifies Y in Step 3, by adding the source u to the set of irrelevant agents, i.e., $R' = R \cup \{u\}$. Then $Y'(a) = Y(a)$ for each agent $a \in N \setminus R'$, so the envy graph $\tilde{G}_{Y' \prec}$ is obtained from $\tilde{G}_{Y \prec}$ by deleting u .

Since deleting some arcs or a vertex from an acyclic graph results in an acyclic graph, the stability of Y' is clear. \square

We proceed with the observation that an agent's situation in Y may only improve, unless it becomes irrelevant: this is a consequence of the fact that the algorithm only deletes arcs and agents from the envy graph $\tilde{G}_{Y \prec}$.

Proposition 8. *Let Y_1 and Y_2 be two pre-allocations computed by algorithm HM-Improve, with Y_1 computed at an earlier step than Y_2 , and let a be an agent that is not irrelevant at the end of the iteration when Y_2 is computed. Then either $Y_1(a) = Y_2(a)$ or a prefers Y_2 to Y_1 .*

We need an additional lemma that will be useful for arguing why irrelevant agents may not become the cause of instability in the housing market.

Lemma 9. *At the end of algorithm HM-Improve, there does not exist an arc $(a, b) \in \tilde{E}$ such that $a \in N \setminus R$, $b \in R$ and $Y(a) \prec'_a b$.*

Proof. Suppose for contradiction that (a, b) is such an arc, and let Y and R be as defined at the end of the last iteration. Let us suppose that HM-Improve adds b to R during the n -th iteration, and let Y_n be the pre-allocation at the beginning of the n -th iteration. By Proposition 8, either $Y_n(a) = Y(a)$ or $Y_n(a) \prec'_a Y(a)$. The assumption $Y(a) \prec'_a b$ yields $Y_n(a) \prec'_a b$ by the transitivity of \prec'_a . Thus, (a, b) is a Y_n -augmenting arc entering b , contradicting our assumption that the algorithm put b into R in Step 3 of the n -th iteration. \square

The following lemma, the last one necessary to prove Theorem 4, shows that HM-Improve runs in linear time; the proof relies on the fact that in each iteration but the last either an agent or an arc is deleted from the envy graph, thus limiting the number of iterations by $|E| + |N|$.

Lemma 10 (\star). *Algorithm HM-Improve runs in $O(|H|)$ time.*

Proof (of Theorem 4). By Lemma 10 it suffices to show that algorithm HM-Improve is correct. Let Y and R be the pre-allocation and the set of irrelevant agents, respectively, at the end of algorithm HM-Improve, and let u be the source of Y . To begin, we prove it formally that X' is an allocation for H' .

First assume $u = \tilde{q}$. This means that Y is the union of disjoint cycles covering each agent in $N \setminus R$ exactly once; note that no arc of Y enters or leaves \tilde{q} . Hence, Y is an allocation not only in $\tilde{H} - R$, but also in the submarket of H' on agent

set $N \setminus R$, i.e., $H'_{N \setminus R}$. Second, assume that $u = p$; in this case $(q, \tilde{q}) \in Y$, because \tilde{q} can be entered only through (q, \tilde{q}) . So the arc set $Y \setminus \{(q, \tilde{q})\} \cup \{(q, p)\}$ is an allocation in $H'_{N \setminus R}$. Consequently, X' is indeed an allocation in H' in both cases.

Now, let us prove that the allocation X' is in the core of H' by showing that the envy graph $G_{X', \prec}^{H'}$ of X' is acyclic. First, the subgraph $G_{X', \prec}^{H'}[R]$ is exactly the envy graph of X_R in H'_R and hence is acyclic.

Claim. *Let $a \in N \setminus R$ and let (a, b) be an X' -augmenting arc in H' . Then (a, b) is Y -augmenting as well, i.e., $Y(a) \prec'_a b$.*

Proof (of Claim). If $(a, b) \neq (q, p)$, then (a, b) is an arc in $G^{\tilde{H}}$, and thus the claim follows immediately from $X'(a) = Y(a)$ except for the case $a = q$ and $Y(q) = \tilde{q}$; in this latter case $X'(q) = p \prec'_q b$ implies that q prefers b to $Y(q) = \tilde{q}$ in \tilde{H} as well, that is, (q, b) is Y -augmenting.

We finish the proof of the claim by showing that (q, p) is not X' -augmenting if $q \notin R$. Let u be the source of Y . If $u = p$, then this is clear by $(q, p) \in X'$. If $u = \tilde{q}$, then let us now consider the penultimate iteration in which the source of Y is moved to \tilde{q} either in Step 2 or in Step 3. Recall that the only arc entering \tilde{q} is (q, \tilde{q}) . If \tilde{q} became the source of Y in Step 2, then we know $\tilde{q} \prec_q Y(q)$. By the construction of \tilde{H} , this means that q prefers $Y(q) = X'(q)$ to p in H' , so (q, p) is not X' -augmenting, a contradiction. Finally, if \tilde{q} became the source of Y in Step 3, then we get $q \in R$, which contradicts our assumption $q \notin R$. ■

As a consequence of our claim, we obtain that $G_{X', \prec}^{H'}[N \setminus R]$ is a subgraph of $\tilde{G}_{Y, \prec}$ and therefore it is acyclic by Lemma 7. Hence, any cycle in $G_{X', \prec}^{H'}$ must contain agents both in R and in $N \setminus R$ (recall that $G_{X', \prec}^{H'}[R]$ is acyclic as well). However, $G_{X', \prec}^{H'}$ contains no arcs from $N \setminus R$ to R , since such arcs cannot be Y -augmenting by Lemma 9. Thus $G_{X', \prec}^{H'}$ is acyclic and X' is in the core of H' . □

4.3 Strict improvement

Looking at Theorem 4 and Corollary 5, one may wonder whether it is possible to detect efficiently when a p -improvement leads to a situation that is strictly better for p . For a solution concept Φ and housing markets H and H' such that H' is a p -improvement of H for some agent p , one may ask the following questions:

1. POSSIBLE STRICT IMPROVEMENT FOR BEST HOUSE or PSIB:
is it true that $a \prec_p a'$ for some $a \in \Phi(H)_p^+$ and $a' \in \Phi(H')_p^+$?
2. NECESSARY STRICT IMPROVEMENT FOR BEST HOUSE or NSIB:
is it true that $a \prec_p a'$ for every $a \in \Phi(H)_p^+$ and $a' \in \Phi(H')_p^+$?
3. POSSIBLE STRICT IMPROVEMENT FOR WORST HOUSE or PSIW:
is it true that $a \prec_p a'$ for some $a \in \Phi(H)_p^-$ and $a' \in \Phi(H')_p^-$?
4. NECESSARY STRICT IMPROVEMENT FOR WORST HOUSE or NSIW:
is it true that $a \prec_p a'$ for every $a \in \Phi(H)_p^-$ and $a' \in \Phi(H')_p^-$?

Focusing on the core of housing markets, it turns out that all of the above four problems are computationally intractable, even in the case of strict preferences.

Theorem 11 (\star). *With respect to the core of housing markets, PSIB and NSIB are NP-hard, while PSIW and NSIW are coNP-hard, even if agents' preferences are strict orders.*

5 The effect of improvements in Stable Roommates

In the STABLE ROOMMATES problem we are given a set N of agents, and a preference relation \prec_a over N for each agent $a \in N$; the task is to find a stable matching M between the agents. A matching is *stable* if it admits no *blocking pair*, that is, a pair of agents such that each prefers the other over her partner in the matching. Notice that an input instance for STABLE ROOMMATES is in fact a housing market. Viewed from this perspective, a stable matching in a housing market can be thought of as an allocation that (i) contains only cycles of length at most 2, and (ii) does not admit a blocking cycle of length at most 2.

For an instance of STABLE ROOMMATES, we assume mutual acceptability, that is, for any two agents a and b , we assume that $a \prec_a b$ holds if and only if $b \prec_b a$ holds. Consequently, it will be more convenient to define the acceptability graph G^H of an instance H of STABLE ROOMMATES as an undirected simple graph where agents a and b are connected by an edge $\{a, b\}$ if and only if they are acceptable to each other and $a \neq b$. A *matching* in H is then a set of edges in G^H such that no two of them share an endpoint.

Biró et al. [9] have shown the following statements.

Proposition 12 ([9]). *Stable matchings in the STABLE ROOMMATES model*

- *violate the RI-worst property (even if agents' preferences are strict), and*
- *violate the RI-best property, if agents' preferences may include ties.*

Complementing Proposition 12, we show that a (p, q) -improvement can lead to an instance where no stable matching exists at all. This may happen even in the case when preferences are strict orders; hence, stable matchings do not strongly satisfy the RI-best property. For an illustration of Propositions 12 and 13 by simple examples see Appendices A.3 and A.4, respectively.

Proposition 13 (\star). *Stable matchings in the STABLE ROOMMATES model do not strongly satisfy the RI-best property, even if agents' preferences are strict.*

Contrasting Propositions 12 and 13, it is somewhat surprising that if agents' preferences are strict, then the RI-best property holds for the STABLE ROOMMATES setting. Thus, the situation of p cannot deteriorate as a consequence of a p -improvement unless instability arises. The proof of Theorem 14 is provided at the end of this section.

Theorem 14. *Let $H = (N, \{\prec_a\}_{a \in N})$ be a housing market where agents' preferences are strict orders. Given a stable matching M in H and a (p, q) -improvement H' of H for two agents $p, q \in N$, either H' admits no stable matchings at all, or there exists a stable matching M' in H' such that $M(p) \preceq_i M'(p)$. Moreover, given H , H' and M it is possible to find such a matching M' in polynomial time.*

Corollary 15. *Stable matchings in the STABLE ROOMMATES model satisfy the RI-best property.*

Structural ingredients. To prove Theorem 14 we are going to rely on the concept of proposal-rejection alternating sequences introduced by Tan and Hsueh [41], originally used as a tool for finding a stable partition in an incremental fashion by adding agents one-by-one to a STABLE ROOMMATES instance. We somewhat tailor their definition to fit our current purposes.

Let $\alpha_0 \in N$ be an agent in a housing market H , and let M_0 be a stable matching in $H - \alpha_0$. A sequence S of agents $\alpha_0, \beta_1, \alpha_1, \dots, \beta_k, \alpha_k$ is a *proposal-rejection alternating sequence* starting from M_0 , if there exists a sequence of matchings M_1, \dots, M_k such that for each $i \in \{1, \dots, k\}$

- (i) β_i is the agent most preferred by α_{i-1} among those who prefer α_{i-1} to their partner in M_{i-1} or are unmatched in M_{i-1} ,
- (ii) $\alpha_i = M_{i-1}(\beta_i)$, and
- (iii) $M_i = M_{i-1} \setminus \{\{\alpha_i, \beta_i\}\} \cup \{\{\alpha_{i-1}, \beta_i\}\}$ is a matching in $H - \alpha_i$.

We say that the sequence S starts from M_0 , and that the matchings M_1, \dots, M_k are induced by S . We say that S stops at α_k , if there does not exist an agent fulfilling condition (i) in the above definition for $i = k + 1$, that is, if no agent prefers α_k to her current partner in M_k and no unmatched agent in M_k finds α_k acceptable. We will also allow a proposal-rejection alternating sequence to take the form $\alpha_0, \beta_1, \alpha_1, \dots, \beta_k$, in case conditions (i), (ii), and (iii) hold for each $i \in \{1, \dots, k-1\}$, and β_k is an unmatched agent in M_{k-1} satisfying condition (i) for $i = k$. In this case we define the last matching induced by the sequence as $M_k = M_{k-1} \cup \{\{\alpha_{k-1}, \beta_k\}\}$, and we say that the sequence stops at agent β_k .

We summarize the most important properties of proposal-rejection alternating sequences in Lemma 16 as observed and used by Tan and Hsueh.⁷

Lemma 16 ([41] \star). *Let $\alpha_0, \beta_1, \alpha_1, \dots, \beta_k, (\alpha_k)$ be a proposal-rejection alternating sequence starting from a stable matching M_0 and inducing the matchings M_1, \dots, M_k in a housing market H . Then the following hold.*

1. M_i is a stable matching in $H - \alpha_i$ for each $i \in \{1, \dots, k-1, (k)\}$.
2. If $\beta_j = \alpha_i$ for some i and j , then H does not admit a stable matching; in such a case we say that sequence S has a return.
3. If the sequence stops at α_k or β_k , then M_k is a stable matching in H .
4. For any $i \in \{1, \dots, k-1\}$ agent α_i prefers $M_{i-1}(\alpha_i)$ to $M_{i+1}(\alpha_i)$.
5. For any $i \in \{1, \dots, k-1\}$ agent β_i prefers $M_i(\beta_i)$ to $M_{i-1}(\beta_i)$.

Description of algorithm SR-Improve. Let $M = (N, \{\prec_a\}_{a \in N})$ be the stable matching given for the housing market H , and let $H' = (N, \{\prec'_a\}_{a \in N})$ be a (p, q) -improvement of H for two agents p and q in N (recall that $\prec'_a = \prec_a$ unless $a = q$). We now propose algorithm SR-Improve that computes a stable matching M' in H' with $M(p) \preceq_p M'(p)$, whenever H' admits some stable matching.

⁷ The first claim of the lemma is only implicit in the paper by Tan and Hsueh [41], we prove it for the sake of completeness in Appendix A.4.

First, SR-Improve checks whether M is stable in H' , and if so, returns the matching $M' = M$. Otherwise, $\{p, q\}$ must be a blocking pair for M in H' .

Second, the algorithm checks whether H' admits a stable matching and if so, computes *any* stable matching M^* in H' using Irving's algorithm [26]; if no stable matching exists for H' , algorithm SR-Improve stops. Now, if $M(p) \preceq'_p M^*(p)$, then SR-Improve returns $M' = M^*$, otherwise proceeds as follows.

Let \tilde{H} be the housing market obtained from H' by deleting all agents $\{a \in N : a \preceq'_q p\}$ from the preference list of q (and vice versa, deleting q from the preference list of these agents). Notice that in particular this includes the deletion of p as well as of $M(q)$ from the preference list of q (recall that $M(q) \prec'_q p$).

Let us define $\alpha_0 = M(q)$ and $M_0 = M \setminus \{q, \alpha_0\}$. Notice that M_0 is a stable matching in $\tilde{H} - \alpha_0$: clearly, any possible blocking pair must contain q , but any blocking pair $\{q, a\}$ that is blocking in \tilde{H} would also block H by $M(q) \prec_q a$. Observe also that q is unmatched in M_0 .

Finally, SR-Improve builds a proposal-rejection alternating sequence S of agents $\alpha_0, \beta_1, \alpha_1, \dots, \beta_k, \alpha_k$ in \tilde{H} starting from M_0 , and inducing matchings M_1, \dots, M_k until one of the following cases occurs:

- (a) $\alpha_k = p$: in this case SR-Improve outputs $M' = M_k \cup \{\{p, q\}\}$;
- (b) S stops: in this case SR-Improve outputs $M' = M_k$.

Correctness of algorithm SR-Improve. The proof that algorithm SR-Improve is correct relies on the following two facts.

Lemma 17 (\star). *The sequence S cannot have a return. Furthermore, if S stops, then it stops at β_k with $\beta_k = q$.*

Lemma 18 (\star). *If SR-Improve outputs a matching M' , then M' is stable in H' and $M(p) \preceq'_p M'(p)$.*

Proof (of Theorem 14). From the description of SR-Improve and Lemma 18 it is immediate that any output the algorithm produces is correct. It remains to show that it does not fail to produce an output. By Lemma 17 we know that the sequence S built by the algorithm cannot have a return and can only stop at q , implying that SR-Improve will eventually produce an output. Considering the fifth statement of Lemma 16, we also know that the length of S is at most $2|E|$. Thus, the algorithm finishes in $O(|E|)$ time. \square

6 Further research

Even though the property of respecting improvement is important in exchange markets, many solution concepts have not been studied from this aspect. For instance, in the STABLE ROOMMATES setting with weakly or partially ordered preferences, do strongly stable matchings satisfy the RI-best property? What about stable half-matchings (or equivalently, stable partitions) in instances of STABLE ROOMMATES without a stable matching? Although Appendix A.5 contains an example about stable half-matchings where improvement of an agents' house damages her situation, perhaps a more careful investigation may shed light on some interesting monotonicity properties.

References

1. A. Abdulkadiroğlu and T. Sönmez. House allocation with existing tenants. *J. Econ. Theory*, 88(2):233–260, 1999.
2. D. J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *EC '07: Proc. of the 8th ACM Conference on Electronic Commerce*, pages 295–304, 2007.
3. D. J. Abraham, K. Cechlárová, D. F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *ISAAC 2004*, pages 3–15, Berlin, Heidelberg, 2004.
4. J. Alcalde-Unzu and E. Molis. Exchange of indivisible goods and indifferences: The Top Trading Absorbing Sets mechanisms. *Game. Econ. Behav.*, 73(1):1–16, 2011.
5. H. Aziz and B. de Keijzer. Housing markets with indifferences: A tale of two mechanisms. In *AAAI '12*, pages 1249–1255, 2012.
6. M. Balinski and T. Sönmez. A tale of two mechanisms: Student placement. *J. Econ. Theory*, 84(1):73–94, 1999.
7. P. Biró and K. Cechlárová. Inapproximability of the kidney exchange problem. *Inform. Process. Lett.*, 101(5):199–202, 2007.
8. P. Biró, B. Haase-Kromwijk, T. Andersson, E. I. Ásgeirsson, T. Baltsová, I. Boletis, et al. Building kidney exchange programmes in Europe: an overview of exchange practice and activities. *Transplantation*, 103(7):1514–1522, 2019.
9. P. Biró, F. Klijn, X. Klimentova, and A. Viana. Shapley-Scarf housing markets: Respecting improvement, integer programming, and kidney exchange. *CoRR*, abs/2102.00167, 2021. arXiv:2102.00167 [econ.TH].
10. P. Biró, D. Manlove, and R. Rizzi. Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Math. Algorithms Appl.*, 1(4):499–517, 2009.
11. P. Biró and E. McDermid. Three-sided stable matchings with cyclic preferences. *Algorithmica*, 58(1):5–18, 2010.
12. P. Biró, J. van de Klundert, D. Manlove, and et al. Modelling and optimisation in European Kidney Exchange Programmes. *Eur. J. Oper. Res.*, 291(2):447–456, 2021.
13. F. Bloch and D. Cantala. Markovian assignment rules. *Soc. Choice Welfare*, 40:1–25, 2003.
14. D. Bokal, G. Fijavž, M. Juvan, P. M. Kayll, and B. Mohar. The circular chromatic number of a digraph. *J. Graph Theor.*, 46(3):227–240, 2004.
15. K. Cechlárová, T. Fleiner, and D. F. Manlove. The kidney exchange game. In *SOR '05*, pages 77–83, 2005.
16. K. Cechlárová and J. Hajduková. Computational complexity of stable partitions with b-preferences. *Int. J. Game Theory*, 31(3):353–364, 2003.
17. K. Cechlárová and V. Lacko. The kidney exchange problem: How hard is it to find a donor? *Ann. Oper. Res.*, 193:255–271, 2012.
18. K. Cechlárová and M. Repiský. On the structure of the core of housing markets. Technical report, P. J. Šafárik University, 2011. IM Preprint, series A, No. 1/2011.
19. K. Cechlárová and A. Romero-Medina. Stability in coalition formation games. *Int. J. Game Theory*, 29(4):487–494, 2001.
20. Á. Cseh and D. F. Manlove. Stable marriage and roommates problems with restricted edges: Complexity and approximability. *Discrete Optim.*, 20:62–89, 2016.
21. V. Dias, G. da Fonseca, C. Figueiredo, and J. Szwarcfiter. The stable marriage problem with restricted pairs. *Theor. Comput. Sci.*, 306:391–405, 01 2003.

22. T. Fleiner, R. W. Irving, and D. F. Manlove. Efficient algorithms for generalized stable marriage and roommates problems. *Theor. Comput. Sci.*, 381(1):162–176, 2007.
23. D. Gusfield and R. W. Irving. *The Stable Marriage problem: Structure and Algorithms*. MIT press, 1989.
24. J. W. Hatfield, F. Kojima, and Y. Narita. Improving schools through school choice: A market design approach. *J. Econ. Theory*, 166(C):186–211, 2016.
25. C.-C. Huang. Circular stable matching and 3-way kidney transplant. *Algorithmica*, 58(1):137–150, 2010.
26. R. W. Irving. An efficient algorithm for the stable roommates problem. *J. Algorithms*, 6(4):577–595, 1985.
27. P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *J. Econ. Theory*, 147(5):1913–1946, 2012.
28. Y. Kamijo and R. Kawasaki. Dynamics, stability, and foresight in the Shapley-Scarf housing market. *J. Math. Econ.*, 46(2):214–222, 2010.
29. R. Kawasaki. Roth–Postlewaite stability and von Neumann–Morgenstern stability. *J. Math. Econ.*, 58:1–6, 2015.
30. B. Klaus and F. Klijn. Minimal-access rights in school choice and the deferred acceptance mechanism. Cahiers de Recherches Economiques du Département d’économie 21.11, Université de Lausanne, 2021.
31. D. E. Knuth. *Mariages stables et leurs relations avec d’autres problèmes combinatoires*. Les Presses de l’Université de Montréal, Montreal, Que., 1976.
32. M. Kurino. House allocation with overlapping generations. *Am. Econ. J.-Microecon.*, 6(1):258–289, 2014.
33. C. G. Plaxton. A simple family of Top Trading Cycles mechanisms for housing markets with indifferences. In *ICGT 2013*, 2013.
34. A. E. Roth and A. Postlewaite. Weak versus strong domination in a market with indivisible goods. *J. Math. Econ.*, 4:131–137, 1977.
35. A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *Quarterly J. of Econ.*, 119:457–488, 2004.
36. A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *J. Econ. Theory*, 125(2):151–188, 2005.
37. D. Saban and J. Sethuraman. House allocation with indifferences: A generalization and a unified view. In *EC ’13: Proc. of the 14th ACM Conference on Electronic Commerce*, pages 803–820, 2013.
38. L. Shapley and H. Scarf. On cores and indivisibility. *J. Math. Econ.*, 1:23–37, 1974.
39. T. Sönmez and T. Switzer. Matching with (Branch-of-Choice) contracts at the United States Military Academy. *Econometrica*, 81:451–488, 2013.
40. J. J. M. Tan. Stable matchings and stable partitions. *Int. J. Comput. Math.*, 39(1-2):11–20, 1991.
41. J. J. M. Tan and Y.-C. Hsueh. A generalization of the stable matching problem. *Discrete Appl. Math.*, 59(1):87–102, 1995.
42. M. U. Ünver. Dynamic kidney exchange. *Rev. Econ. Stud.*, 77(1):372–414, 2010.
43. Y. Yuan. Residence exchange wanted: A stable residence exchange problem. *Eur. J. Oper. Res.*, 90(3):536–546, 1996.

A Appendix

We present all proofs missing from Sections 3 and 4 in Sections A.1 and A.2, respectively. We provide examples for Proposition 12 in Section A.3. Section A.4 contains all proofs missing from Section 5. We close the appendix by some notes on the respecting improvement property in relation to stable half-matchings in a STABLE ROOMMATES instance in Section A.5.

A.1 Missing proofs from Section 3

Theorem 1. *Each of the following problems is NP-complete, even if agents' preferences are strict orders:*

- ARC IN CORE,
- FORBIDDEN ARC IN CORE, and
- AGENT TRADING IN CORE.

Proof. It is easy to see that all of these problems are in NP, since given an allocation X for H , we can check in linear time whether it admits a blocking cycle: taking the envy graph $G_{X \leftarrow}^H$ of X , we only have to check that it is *acyclic*, i.e., contains no directed cycles (this can be decided using, e.g., some variant of the depth-first search algorithm).

To prove the NP-hardness of ARC IN CORE, we present a polynomial-time reduction from the ACYCLIC PARTITION problem: given a directed graph D , decide whether it is possible to partition the vertices of D into two acyclic sets V_1 and V_2 . Here, a set W of vertices is *acyclic*, if $D[W]$ is acyclic. This problem was proved to be NP-complete by Bokal et al. [14].

Given our input $D = (V, A)$, we construct a housing market H as follows (see Fig. 2 for an illustration). We denote the vertices of D by v_1, \dots, v_n , and we define the set of agents in H as

$$N = \{a_i, b_i, c_i, d_i \mid i \in \{1, \dots, n\}\} \cup \{a^*, b^*, a_0, b_0\}.$$

The preferences of the agents' are as shown below; for each agent $a \in N$ we only list those agents whose house a finds acceptable. Here, for any set W of agents we let $[W]$ denote an arbitrary fixed ordering of W .

$$\begin{array}{ll} a^* : b^*; & \\ b^* : a_0, a_1, \dots, a_n, a^*; & \\ a_i : b_i, b^* & \text{where } i \in \{0, 1, \dots, n\}; \\ b_i : c_{i+1}, d_{i+1} & \text{where } i \in \{0, 1, \dots, n-1\}; \\ b_n : a_0; & \\ c_i : d_i, [\{c_j \mid (v_i, v_j) \in A\}], a_i & \text{where } i \in \{1, \dots, n\}; \\ d_i : c_i, [\{d_j \mid (v_i, v_j) \in A\}], a_i & \text{where } i \in \{1, \dots, n\}. \end{array}$$

We finish the construction by defining our instance of ARC IN CORE as the pair $(H, (a^*, b^*))$. We claim that there exists an allocation in the core of H




Fig. 2. Illustration of the housing market H constructed in the NP-hardness proof for ARC in CORE. The symbol ∞ indicates the least-preferred choice of an agent. The example assumes that (v_1, v_2) and (v_n, v_2) are arcs of the directed input graph D , as indicated by the dashed arcs.

containing (a^*, b^*) if and only if the vertices of D can be partitioned into two acyclic sets.

“ \Rightarrow ”: Let us suppose that there exists an allocation X that does not admit any blocking cycles and contains (a^*, b^*) .

We first show that X contains every arc (a_i, b_i) for $i \in \{0, 1, \dots, n\}$. To see this, observe that the only possible cycle in X that contains (a^*, b^*) is the cycle (a^*, b^*) of length 2, because the arc (b^*, a^*) is the only arc going into a^* . Hence, if for some $i \in \{0, 1, \dots, n\}$ the arc (a_i, b_i) is not in X , then the cycle (a_i, b^*) is a blocking cycle. As a consequence, exactly one of the arcs (b_i, c_{i+1}) and (b_i, d_{i+1}) must be contained in X for any $i \in \{0, 1, \dots, n-1\}$, and similarly, exactly one of the arcs (c_i, a_i) and (d_i, a_i) is contained in X for any $i \in \{1, \dots, n\}$.

Next consider the agents c_i and d_i for some $i \in \{1, \dots, n\}$. As they are each other’s top choice, it must be the case that either (c_i, d_i) or (d_i, c_i) is contained in X , as otherwise they both prefer to trade with each other as opposed to their allocation according to X , and the cycle (c_i, d_i) would block X . Using the facts of the previous paragraph, we obtain that for each $v_i \in V$ exactly one of the following conditions holds:

- X contains the arcs (b_{i-1}, c_i) , (c_i, d_i) , and (d_i, a_i) , in which case we put v_i into V_1 ;
- X contains the arcs (b_{i-1}, d_i) , (d_i, c_i) , and (c_i, a_i) , in which case we put v_i into V_2 .

We claim that both V_1 and V_2 are acyclic in D . For a contradiction, let C_1 be a cycle within vertices of V_1 in D . Note that any arc (v_i, v_j) of C_1 corresponds to an arc (d_i, d_j) in the acceptability graph $G = G^H$ for H . Moreover, since $v_i \in V_1$, by definition we know that d_i prefers d_j to $X(d_i) = a_i$. This yields that the agents $\{d_i \mid v_i \text{ appears on } C_1\}$ form a blocking cycle for H . The same

argument works to show that any cycle C_2 within V_2 corresponds to a blocking cycle formed by the agents $\{c_i \mid v_i \text{ appears on } C_2\}$, proving the acyclicity of V_2 .

“ \Leftarrow ”: Assume now that V_1 and V_2 are two acyclic subsets of V forming a partition. We define an allocation X to contain the cycle (a^*, b^*) , and a cycle consisting of the arcs in

$$\begin{aligned} X_o = & \{(b_n, a_0)\} \cup \{(a_i, b_i) \mid v \in \{0, 1, \dots, n\}\} \\ & \cup \{(b_{i-1}, c_i), (c_i, d_i), (d_i, a_i) \mid v_i \in V_1\} \\ & \cup \{(b_{i-1}, d_i), (d_i, c_i), (c_i, a_i) \mid v_i \in V_2\}. \end{aligned}$$

Observe that X_o is indeed a cycle, and that X is an allocation containing the arc (a^*, b^*) . We claim that the core of H contains X . Assume for the sake of contradiction that X admits a blocking cycle C . Now, since a^* as well as each agent a_i , $i \in \{0, 1, \dots, n\}$, is allocated its first choice by X , none of these agents appears on C . This implies that neither b^* , nor any of the agents b_i , $i \in \{0, 1, \dots, n\}$, appears on C , since these agents have no in-neighbors that could possibly appear on C . Furthermore, every agent in $\{c_i \mid v_i \in V_1\} \cup \{d_i \mid v_i \in V_2\}$ is allocated its first choice by X . It follows that C may contain only agents from $D_1 = \{d_i \mid v_i \in V_1\}$ and $C_2 = \{c_i \mid v_i \in V_2\}$. Observe that there is no arc in G from D_1 to C_2 or vice versa, hence C is either contained in $G[D_1]$ or $G[C_2]$. Now, since any cycle within $G[D_1]$ or $G[C_2]$ would correspond to a cycle in D , the acyclicity of V_1 and V_2 ensures that X admits no blocking cycle, proving the correctness of our reduction for the ARC IN CORE problem.

Observe that the same reduction proves the NP-hardness of AGENT TRADING IN CORE, since agent a^* is trading in an allocation X for H if and only if the arc (a^*, b^*) is used in X .

Finally, we modify the above construction to give a reduction from ACYCLIC PARTITION to FORBIDDEN ARC IN CORE. We simply add a new agent s^* to H , with the house of s^* being acceptable only for a^* as its second choice (after b^*), and with s^* preferring only a^* to its own house. We claim that the resulting market H' together with the arc (a^*, s^*) is a yes-instance of FORBIDDEN ARC IN CORE if and only if H with (a^*, b^*) constitutes a yes-instance of ARC IN CORE. To see this, it suffices to observe that any allocation for H' not containing (a^*, s^*) is either blocked by the cycle (a^*, s^*) of length 2, or contains the arc (a^*, b^*) . Hence, any allocation in the core of H' contains (a^*, b^*) if and only if it does not contain (a^*, s^*) , proving the theorem. \square

Theorem 2. *For any constant $\varepsilon > 0$, the MAX CORE problem is NP-hard to approximate within a ratio of $\alpha_\varepsilon(N) = |N|^{1-\varepsilon}$ where N is the set of agents, even if agents' preferences are strict orders.*

Proof. Let $\varepsilon > 0$ be a constant. Assume for the sake of contradiction that there exists an approximation algorithm \mathcal{A}_ε that given an instance H of MAX CORE with agent set N computes in time polynomial in $|N|$ an allocation in the core of H having size at least $\text{OPT}(H)/\alpha_\varepsilon(N)$, where $\text{OPT}(H)$ is the maximum size of (i.e., number of agents trading in) any allocation in the core of H . We can

prove our statement by presenting a polynomial-time algorithm for the NP-hard ACYCLIC PARTITION problem using \mathcal{A}_ε .

We are going to re-use the reduction presented in the proof of Theorem 1 from ACYCLIC PARTITION to ARC IN CORE. Recall that the input of this reduction is a directed graph D on n vertices, and it constructs a housing market H containing a set N of $4n + 4$ agents and a pair (a^*, b^*) of agents such that the vertices of D can be partitioned into two acyclic sets if and only if some allocation in the core of H contains the arc (a^*, b^*) . Moreover, such an allocation (if existent) must have size $4n + 4$, by our arguments in the proof of Theorem 1.

Let us now define a housing market $H' = (N', \{\prec_a\}_{a \in N'})$ that can be obtained by subdividing the arc (a^*, b^*) with K newly introduced agents p_1, \dots, p_K where

$$K = \lceil (4n + 4)^{1/\varepsilon} \rceil.$$

Let $N' = N \cup \{p_1, \dots, p_K\}$. Formally, we define preferences \prec'_a for each agent $a \in N'$ as follows: \prec'_a is identical to \prec_a if $a \in N \setminus \{a^*\}$, \prec'_{a^*} is obtained from \prec_{a^*} by p_1 taking the role of b^* , and each agent $p_i \in N' \setminus N$ prefers only the house of agent p_{i+1} to her own house (where we set $p_{K+1} = b^*$). Clearly, the allocations in the core of H correspond to the allocations in the core of H' in a bijective manner. Hence, it is easy to see that if there is an allocation in the core of H that contains (a^*, b^*) and where every agent of N is trading, then there is an allocation in the core of H' where each agent of N' is trading. Conversely, if there is no allocation in the core of H that contains (a^*, b^*) , then the agents p_1, \dots, p_K cannot be trading in any allocation in the core of H' . Thus, we have that if D is a yes-instance of ACYCLIC PARTITION, then $\text{OPT}(H') = |N'| = 4n + 4 + K$; otherwise $\text{OPT}(H') \leq 4n + 4$.

Now, after constructing H' we apply algorithm \mathcal{A}_ε with H' as its input; let X' be its output. If the size of X' is greater than $4n + 4$, then we conclude that D must be a yes-instance of ACYCLIC PARTITION, as implied by the previous paragraph. Otherwise, we conclude that D is a no-instance of ACYCLIC PARTITION. To show that this is correct, it suffices to see that if D is a yes-instance, that is, if $\text{OPT}(H') = |N'|$, then the size of X' is greater than $4n + 4$. And indeed, the definition of K implies

$$(4n + 4)^{1/\varepsilon} < 4n + 4 + K = |N'|$$

which raised to the power of ε yields

$$4n + 4 < |N'|^\varepsilon = \frac{|N'|}{|N'|^{1-\varepsilon}} = \frac{\text{OPT}(H')}{\alpha_\varepsilon(N')}$$

as required.

It remains to observe that the above reduction can be computed in polynomial time, because ε is a constant and so K is a polynomial of n of fixed degree. \square

Proposition 3. *MAX CORE can be approximated with a ratio of $|N|$ in polynomial time, where $|N|$ is the number of agents in the input.*

Proof. An approximation algorithm for MAX CORE has ratio $|N|$, if for any housing market H with agent set N it outputs an allocation with at least $\text{OPT}(H)/|N|$ agents trading, where $\text{OPT}(H)$ is the maximum number of trading agents in a core allocation of H . Thus, it suffices to decide whether $\text{OPT}(H) \geq 1$, and if so produce an allocation in which at least one agent is trading. Observe that $\text{OPT}(H) = 0$ is only possible if G^H is acyclic, as any cycle in G^H blocks the allocation where each agent gets her own house. Hence, computing *any* allocation in the core of H is an $|N|$ -approximation for MAX CORE; this can be done in linear time using the variant of the TTC algorithm described in Section 3.1. \square

A.2 Missing proofs from Section 4

Lemma 10. *Algorithm HM-Improve runs in $O(|H|)$ time.*

Proof. Observe that the initialization takes $O(|E| + |N|) = O(|E|)$ time; note that E contains every loop (a, a) where $a \in N$, so we have $|E| \geq |N|$. We can maintain the envy graph $\tilde{G}_{Y \prec}$ in a way that deleting an arc from it when it ceases to be Y -augmenting can be done in $O(1)$ time, and detecting whether a given agent is entered by a Y -augmenting arc also takes $O(1)$ time. Observe that there can be at most $|E| + |N|$ iterations, since at each step but the last, either an agent or an arc is deleted from the envy graph. Thus, the whole iteration takes $O(|E|)$ time. Finally, the allocation X_R for irrelevant agents by the variant of TTC described in Section 3.1 can be computed in $O(|H|)$ time. Hence, the overall running time of our algorithm is $O(|H|) + O(|E|) = O(|H|)$. \square

Theorem 11. *With respect to the core of housing markets, PSIB and NSIB are NP-hard, while PSIW and NSIW are coNP-hard, even if agents' preferences are strict orders.*

Proof. Since agents' preferences are strict orders, we get that PSIB and NSIB are equivalent, and similarly, PSIW and NSIW are equivalent as well, since there is a unique best and a unique worst house that an agent may obtain in a core allocation. Therefore, we are going to present two reductions, one for PSIB and NSIB, and one for PSIW and NSIW. Since both reductions will be based on those presented in the proof of Theorem 1, we are going to re-use the notation defined there.

The reduction for PSIB (and NSIB) is obtained by slightly modifying the reduction from ACYCLIC PARTITION to ARC IN CORE which, given a directed graph D constructs the housing market H . We define a housing market \hat{H} by simply deleting the arc (b^*, a^*) from the acceptability graph of H . Then H is an a^* -improvement of \hat{H} . Clearly, as the house of a^* is not acceptable to any other agent in \hat{H} , the best house that a^* can obtain in any allocation in the core of \hat{H} is her own. Moreover, the best house that a^* can obtain in any allocation in the core of H is either the house of b^* or her own. This immediately implies that (\hat{H}, H) is a yes-instance of PSIB (and of NSIB) with respect to the core if and only if there exists an allocation in the core of H that contains the arc




Fig. 3. The housing markets H and H' as instances of STABLE ROOMMATES with ties, in the Example 19. For both H and H' , the matching represented by bold arcs yields the best possible partner for p in any stable matching of the given market.

(a^*, b^*) . Therefore, (\widehat{H}, H) is a yes-instance of PSIB and of NSIB with respect to the core if and only if D is a yes-instance of ACYCLIC PARTITION, finishing our proof for PSIB (and NSIB).

The reduction for PSIW (and NSIW) is obtained analogously, by slightly modifying the reduction from ACYCLIC PARTITION to FORBIDDEN ARC IN CORE which, given a directed graph D constructs the housing market H' . We define a housing market \widehat{H}' by deleting the arc (a^*, s^*) from the acceptability graph of H' . Then H' is an s^* -improvement of \widehat{H}' . Clearly, as the house of s^* is not acceptable to any other agent in \widehat{H}' , the worst house that s^* can obtain in any allocation in the core of \widehat{H}' is her own. Moreover, the worst house that s^* can obtain in any allocation in the core of H' is either the house of a^* or her own. Therefore, (\widehat{H}', H') is a no-instance of PSIW (and of NSIW) with respect to the core if and only if there exists an allocation in the core of H' where s^* is not trading, i.e., that does not contain the arc (a^*, s^*) . So (\widehat{H}', H') is a no-instance of PSIW and of NSIW with respect to the core if and only if D is a yes-instance of ACYCLIC PARTITION, finishing our proof for PSIW (and NSIW). \square

A.3 Examples for Proposition 12

Example 19. Let $N = \{a, b, c, d, e, p, q\}$ be the set of agents. The preferences indicated in Figure 3 define two housing markets H and H' such that H' is a (p, q) -improvement of H . Note that d is indifferent between her two possible partners. Looking at H and H' in the context of STABLE ROOMMATES, it is easy to see that the best partner that p might obtain in a stable matching for H is her second choice b , while in H' the only stable matching assigns a to p , which is her third choice.

Example 20. Let $N = \{a, b, p, q\}$ be the set of agents. The preferences indicated in Figure 4 define two housing markets H and H' such that H' is a (p, q) -improvement of H . The worst partner that p might obtain in a stable matching for H is her top choice a , while in H' there exists a stable matching that assigns b to p , which is her second choice.




Fig. 4. The housing markets H and H' in Example 20. For both H and H' , the matching represented by bold arcs yields the worst possible partner for p in any stable matching of the given market.




Fig. 5. Housing markets H and H' illustrating the proof of Proposition 13. For H , the bold arcs represent a stable matching, while the instance H' , which is a (p, q) -improvement of H , does not admit any stable matchings.

A.4 Missing proofs from Section 5

Proposition 13. *Stable matchings in the STABLE ROOMMATES model do not strongly satisfy the RI-best property, even if agents' preferences are strict.*

Proof. Let $N = \{a, b, p, q\}$ be the set of agents. The preferences indicated in Figure 5 define housing markets H and H' where H' is an (p, q) -improvement of H . The best partner that p might obtain in a stable matching for H is her second choice a , while H' does not admit any stable matchings at all. \square

Lemma 16 (statement 1). *Let $\alpha_0, \beta_1, \alpha_1, \dots, \beta_k, \alpha_k$ be a proposal-rejection alternating sequence starting from a stable matching M_0 and inducing matchings M_1, \dots, M_k in a housing market H . Then M_i is a stable matching in $H - \alpha_i$ for each index $i \in \{1, \dots, k - 1, k\}$.*

Proof. We prove the lemma by induction on i ; the case $i = 0$ is clear. Assume that $i \geq 1$ and M_{i-1} is stable in $H - \alpha_{i-1}$. Since $M_i \Delta M_{i-1} = \{\{\alpha_i, \beta_i\}, \{\alpha_{i-1}, \beta_i\}\}$ we know that any blocking pair for M_i in $H - \alpha_i$ must contain either β_i or α_{i-1} . By our choice of β_i , it is clear that α_{i-1} cannot be contained in a blocking pair. Moreover, since β_i prefers α_{i-1} to $M_{i-1}(\beta_i) = \alpha_i$, any blocking pair for M_i would also be blocking in M_{i-1} , a contradiction. \square

Lemma 17. *The sequence S cannot have a return. Furthermore, if S stops, then it stops at β_k with $\beta_k = q$.*

Proof. Recall that M^* is a stable matching in H' with $M^*(p) \prec_p M(p)$. Since $\{p, q\}$ is a blocking pair for M in H' , we know $M(p) \prec_p q$, yielding $M^*(p) \prec_p q$. By the stability of M^* , this implies that q is matched in M^* and $p \prec'_q M^*(q)$. As a consequence, M^* is a stable matching not only in H' but also in \tilde{H} , since deleting agents less preferred by q than $M^*(q)$ from q 's preference list cannot compromise the stability of M^* .

By the second claim of Lemma 16, we know that if S has a return, then \tilde{H} admits no stable matching, contradicting the existence of M^* . Furthermore, since q is matched in M^* , it must be matched in every stable matching of \tilde{H} , by the well-known fact that in an instance of STABLE ROOMMATES where agents' preferences are strict all stable matchings contain exactly the same set of agents [23, Theorem 4.5.2]. Now, if S stops with the last induced matching M_k , then by the third statement of Lemma 16 we get that M_k is a stable matching in \tilde{H} , and thus q must be matched in M_k . Clearly, as q is unmatched in M_0 , this can only occur if $\beta_k = q$ and S stops at q . \square

Lemma 18. *If SR-Improve outputs a matching M' , then M' is stable in H' and $M(p) \preceq'_p M'(p)$.*

Proof. First, assume that the algorithm stops when $\alpha_k = p$. Then by the first statement of Lemma 16, M_k is stable in $\tilde{H} - p$. Note also that q must be unmatched in M_k , as q can only obtain a partner in the sequence of matchings induced by S if $q = \beta_k$, which cannot happen when $\alpha_k = p$. So $M' = M_k \cup \{\{p, q\}\}$ is indeed a matching in H' .

Let us prove that M' is stable in H' . Since q is unmatched in M_k , and M_k is stable in $\tilde{H} - p$, no agent acceptable for q prefers q to her partner in M_k or is left unmatched in M_q . Hence, q cannot be contained in a blocking pair for M' . Thus, any blocking pair for M' must contain p . Suppose that $\{p, a\}$ blocks M' in H' ; then $q \prec'_p a$. Since S cannot have a return by Lemma 17, we know that p is not among the agents $\alpha_0, \beta_1, \alpha_1, \dots, \beta_{k-1}$. Therefore, $M_{k-1}(p) = M_0(p) = M(p)$. Recall that $M(p) \prec'_p q$, which implies $M_{k-1}(p) \prec'_p q$. Since $M_{k-1}(a) = M_k(a)$ (because $a \notin \{\alpha_{k-1}, \beta_k, p\}$), we get that $\{p, a\}$ must also block M_{k-1} in $\tilde{H} - \alpha_{k-1}$, a contradiction. This shows that M' is stable in H' . By $M(p) \prec'_p q = M'(p)$, the lemma follows in this case.

Second, assume that SR-Improve outputs $M' = M_k$ after finding that the sequence S stops with q being matched in M_k . By the first statement of Lemma 16, we know that M' is stable in \tilde{H} , and by the definition of \tilde{H} , we know that $p \prec_q M'(q)$. Therefore, M' is also stable in H' (as adding agents less preferred by q than $M'(q)$ to q 's preference list cannot compromise the stability of M'). To show that $M(p) \preceq'_p M'(p)$, it suffices to observe that $p = \alpha_i$ is not possible for any $i \in \{1, \dots, k\}$ (as in this case q would be unmatched, as argued in the first paragraph of this proof), and hence by the fifth claim of Lemma 16 the partner that p receives in the matchings M_0, M_1, \dots, M_k can only get better for p , and thus $M(p) = M_0(p) \preceq'_p M_k(p) = M'(p)$. \square

A.5 Stable half-matchings in Stable Roommates

Though an instance of STABLE ROOMMATES may not contain a stable matching, Tan [40] has proved that a so-called stable half-matching can always be found. Given an instance $H = (N, \{\prec_a\}_{a \in N})$ of STABLE ROOMMATES with an underlying graph $G = (N, E)$, a *fractional matching* is a function $f : E \rightarrow [0, 1]$ that assigns to each edge of G a number between 0 and 1 such that for each agent $a \in N$, the total value of f summed over all edges incident to a is at most 1. In other words, the constraint

$$\sum_{e \in \mathcal{N}(a)} f(e) \leq 1 \quad (1)$$

is satisfied for each agent a , where $\mathcal{N}(a)$ denotes the set of edges in E incident to a . A *half-matching* for H is a fractional matching that can only take half-integral values, that is, a function $f : E \rightarrow \{0, \frac{1}{2}, 1\}$ that satisfies Inequality (1). Then f can be viewed as pair (M_f, H_f) where M_f is the set of *matched edges* to which f assigns value 1, and H_f is the set of *half-matched edges* to which f assigns value $\frac{1}{2}$. We say that an edge $\{a, b\} \in \mathcal{N}(a)$ is *beneficial* for agent a , if either $\sum_{e \in \mathcal{N}(a)} f(e) < 1$ or there exists some agent $b' \in N$ such $b' \prec_a b$ and $f(\{a, b'\}) > 0$. A *blocking edge* for f is an edge in E that is beneficial for both of its endpoints; we consider a half-matching f to be *stable*, if there does not exist a blocking edge for f .




Fig. 6. The housing markets H and H' in Example 21. We depict half-matchings in blue, with double lines for matched edges and single bold lines for half-matched edges. For H , the half-matching f depicted leaves p more satisfied than the half-matching f' depicted for H' .

Example 21. Let $N = \{a, b, c, p, q\}$ be the set of agents. The preferences indicated in Figure 6 define two housing markets H and H' such that H' is a (p, q) -improvement of H . Note that the half-matching f that assigns $\frac{1}{2}$ to each edge except for $\{p, q\}$ is stable in H . However, the only stable half-matching that H' admits is f' , assigning 1 to $\{a, b\}$ and assigning $\frac{1}{2}$ to each of the edges $\{p, q\}$, $\{p, c\}$ and $\{c, q\}$. Thus, p is covered by two half-matched edges, $\{p, a\}$ and $\{p, c\}$, in f , while it is covered by the two half-matched edges $\{p, q\}$ and $\{p, c\}$ in f' . Although we have not defined a formal way to compare half-matchings from the viewpoint of an agent, it is clear that the situation of p is strictly better in f than in f' , since p prefers a to q .