

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 HUN-REN 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 (1974). 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 in its core in which a receives either h or a house that a 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 are 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 [8, 12, 49] and on-campus housing [1].

In their seminal work Shapley and Scarf [52] 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 (with at least one agent strictly improving) by trading among themselves [48]. Although the strict core has some very appealing properties from a mathematical viewpoint (above all, that for strict preferences it contains a unique allocation that is easy to compute), there are arguments why the core is a more interesting solution concept. First, if we allow indifference between houses to appear in the preferences, then the strict core can be empty. Second, decision-makers in a real-world application have to deal with various constraints and optimization goals that may not be represented in the preferences. In kidney exchange programs such constraints can arise due to ethical issues [20], logistical considerations (e.g., limit on the length of the exchange cycles), and several optimization criteria that can improve the accessibility and

long term success of the programs (e.g., prioritization of highly sensitized patients, or keeping donors with blood type O for recipients with blood type O in order to avoid the accumulation of hard-to-match patients); see [12] for a survey on European practices. Thus, having a wider range of possible solutions to choose from is often preferable.

Biró et al. [9] conducted computer simulations on realistic kidney exchange instances to compare solutions of maximum size or maximum weight⁴, typically used in practice, with solutions contained in the core, in the set of competitive solutions, and in the strict core of the underlying housing market. Intuitively, as the latter three solution concepts become increasingly demanding to be satisfied, the “price of fairness” (roughly speaking, the reduction in the number of transplants due to taking preferences into account) also increases. Their results suggest that the core is a good compromise: they found that “core allocations for instances with 150 patient–donor pairs entail a less than 1% reduction in the number of transplants.” In contrast, regarding the strict core solutions, they observed that the likelihood of existence for instances with weak preferences is decreasing sharply as the kidney exchange pool grows, becoming almost zero for 150 patient–donor pairs.

Although the core of housing markets has been the subject of considerable research, there are still many challenges which have not been addressed. Our paper focuses on questions that may be raised by an agent who wants to decide whether to enter the market, and if so, on what conditions. To be able to judge their prospects correctly, agents are in need of information about the possible allocations the market may result in. Can they improve their situation by participating in the market? If so, how much? Which are the houses that they have a chance of receiving?

Assuming that our relevant solution concept is the core, an agent a may be interested in the following questions:

- Q1: Can agent a receive a house better than her own in some core allocation?
- Q2: Given some house h , can agent a obtain h in some core allocation?
- Q3: Given some house h , can agent a avoid obtaining h in some core allocation?

We believe that the above three questions are natural enough to warrant a quest for an efficient algorithm that can solve the underlying computational problems. However, their motivation is also clear from an economic point of view: a positive answer to question Q1 (or Q2) clearly provides a strong incentive for agent a to participate in the market. Hence, such an algorithm can be an important tool for an authority in charge of a centralized housing market where a larger market is more desirable (as is the case, e.g., in kidney exchange programs). Further motivation for studying questions Q2 and Q3 stems from the fact that, in some cases, realizing a given allocation requires certain investments that are a function of the allocation chosen: e.g., in the context of kidney exchange, additional compatibility tests are required before carrying out the planned transplantations. Narrowing down the set of possible donors whose kidney a given patient may obtain in a kidney exchange may allow for such tests to be performed in advance, sparing time for the patients, while keeping the costs incurred by such tests relatively low.⁵

⁴ In practice, solutions in a kidney exchange program are often sought as maximum-weight matchings between patient–donor pairs in a graph where weights reflect certain optimality criteria.

⁵ As an extreme example of additional expenditures necessary for realizing an allocation, consider Germany, where kidney exchange is only allowed when it involves people who have a “personal relationship.” However, the required personal relationship can be built for the purpose of the kidney exchange to take place—a process that takes precious time for the patients involved. See https://crossover-nierenspende.de/files/DIATRA_42021_26-29_Crossover_engl.pdf.

In the first part of our work, we focus on the computational complexity of the above questions. Similar questions have been extensively studied in the context of the STABLE MARRIAGE and the STABLE ROOMMATES problems [23, 24, 26, 28, 37], but have not yet been considered in relation to housing markets.

We also address questions that concern the possibility of an agent improving her situation by bringing a better endowment to the market. Assuming that agent a ensures that the value of her house increases, will this result in an improvement for a ? If the answer is positive, then this provides an incentive for the agent to invest in her house in order to obtain a preferable allocation. 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 is it at least true that by improving her house, a will not damage her own possibilities in the market? Can we determine whether when a strict improvement for a becomes possible?

We investigate the following question: is an increase in the value of some agent a 's house beneficial for a in terms of the possible core allocations? More precisely, we consider two slightly different versions of this question:

- Q4: Given a core allocation for the original market where a obtains some house h , can agent a obtain a house *at least as good* as h in some core allocation after an increase in the value of a 's house?
- Q5: Given a core allocation for the original market where a obtains some house h , can an agent a obtain a house *strictly better* than h in some core allocation after an increase in the value of a 's house?

Questions Q4 and Q5 are of crucial importance when we consider agents' incentives to choose the endowment with which they enter the market. In the context of kidney exchange: if procuring a new donor with better properties (e.g., a younger or healthier donor), or registering an additional willing donor (which is possible in most European programs) 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 in a system where inefficiency directly leads to loss of lives. Being able to answer the above questions is therefore paramount, and has direct consequences on agents' incentives.

1.1 Our contribution.

Regarding questions Q1, Q2, and Q3 raised above, we show in Theorem 1 that each of them is computationally intractable. Remarkably, it is already NP-complete to decide whether a core allocation can assign *any* house to a other than her own. Similarly, deciding whether the core of a housing market contains an allocation where a given agent a obtains a certain house (or where a does *not* receive a certain house) is also NP-complete. Various generalizations of these questions can be answered efficiently in both the STABLE MARRIAGE and STABLE ROOMMATES settings [23, 24, 26, 28, 37], so we find these intractability results surprising.

We shall note that our complexity results do not mean that finding core allocations that fulfill some additional requirement would be impossible in practice, it only means that a polynomial-time algorithm is highly unlikely to exist for such problems. However, this does not preclude the use of heuristics or robust optimization techniques such as integer programming (IP) methods for computing core allocations with additional requirements. In fact, Biró et al. [9] developed and tested such methods by conducting simulations, and demonstrated the possibility of using the solution concept of core in practice.

Turning our attention to the question of how an increase in the value of a house affects its owner, we present Theorem 2, our main technical result, which answers question Q4 affirmatively as follows:

If the core of a housing market contains an allocation where a receives some house h , and the market changes in a way 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 a prefers to h .

Using the terminology of Biró et al. [9], the above result shows that 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. We prove Theorem 2 by presenting a polynomial-time algorithm that finds an allocation as promised by the sentence highlighted above; the ideas and techniques on which this algorithm relies is our main technical contribution. This settles an open question asked explicitly by Biró et al. [9].

This result has important implications for the practice of kidney exchanges. Biró et al. [9] conducted simulations for measuring how often the property of respecting improvement is violated when using solutions of maximum size/weight (as done in practice) compared to using core, competitive, or strict core solutions. They observed a significant number of violations for solutions of maximum size/weight, but none for core solutions, conjecturing the theorem that we have proved theoretically. Therefore, Theorem 2 gives a theoretical foundation for their observation, and it implies that the usage of core solutions provides good individual incentives for recipients to bring better or more donors. In the meantime, these simulations show that the current practice of focusing only on the number of transplants and other weighted optimality criteria may not provide a compelling incentive for participants to bring valuable donors to the pool.

The significance of our positive result in Theorem 2 is especially pronounced in view of the intractability results of Theorem 1: even though we cannot efficiently compute information about the possible houses an agent may obtain in some core allocation, we *do know* that entering the market with a more valuable house, an agent can only improve (and never damage) their situation. Hence, improving the value of the house with which the agent plans to enter the market is always a safe choice. We believe that this aspect of the core is a very strong argument for considering it as a good solution concept to be used in centralized housing markets, since it provides a motivation for agents to increase the value of their initial endowment.

Contrasting our positive result for question Q4, the slightly different question Q5 turns out to be significantly harder: although one can formulate several variants of this 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. An instance of STABLE ROOMMATES can therefore be considered as a housing market with the additional requirement that (i) trading can only happen along cycles of length 2, and (ii) only blocking cycles of length 2 can cause instability; then stable matchings correspond exactly to core allocations. We examine the following question, which is the direct analog of question Q4 for the STABLE ROOMMATES model:

- Q6: In an instance of STABLE ROOMMATES, does increasing the value of an agent a (as manifested in the preferences of others) lead to a (weak) improvement in the situation of a ?

Again, we are able to assert a positive answer, although only in a conditional form: in Theorem 16 we show that if some stable matching assigns agent a to agent b in a STABLE ROOMMATES instance, and the value of a increases (that is, if a moves upward in other agents' preferences, with everything else remaining constant), and the resulting instance admits a stable matching, then it necessarily admits a stable matching where a is matched either to b or to an agent preferred by a to b . This result is an analog of the one stated in Theorem 2 for the core of housing markets; however, the algorithm we propose in order to prove it uses different techniques.

We remark that throughout the paper we use a model with partially ordered preferences (a generalization of weak orders). Although partially ordered preferences have been studied in the context of the STABLE MARRIAGE and STABLE ROOMMATES problems [22, 25, 26, 27, 45], we are not aware of any paper on housing markets featuring preferences that are expressed as partial orders.

1.2 Related work.

Most works relating to the core of housing markets aim to find core allocations with some additional property that benefits global welfare, most prominently Pareto optimality [4, 5, 33, 46, 51]. 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 [7, 15, 16, 17, 19] or is bounded by some constant [2, 10, 11, 18, 30]. 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. Although, to the best of our knowledge none of the questions Q1–Q3 have been studied so far, some papers have focused on finding a core allocation where the number of agents involved in trading is as large as possible, obtaining mostly intractability results [7, 18].

In the context of the STABLE MARRIAGE and the STABLE ROOMMATES problems, it is known that the problem of finding a stable matching with edge restrictions, i.e., a stable matching that contains a given set of *forced* edges but is disjoint from a given set of *forbidden* edges, can be found in polynomial time [24, 26]. These results strongly contrast Theorem 1, which shows that the analogous problems in the context of house allocation are NP-hard, even if there is only a single arc that we require to be included in (or excluded from) the desired allocation.

Questions Q4 and Q5 can be considered as inquiries about housing markets where preferences are subject to change. Although some researchers have addressed certain dynamic models, most of these either focus on the possibility of repeated allocation [34, 35, 48], or consider a situation where agents may enter and leave the market at different times [13, 41, 56].

The line of research that concerns questions akin to Q4 and Q5 was initiated by Balinski and Sönmez in their paper on the property of respecting improvement in the context of college admission [6]. They 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)—this means that the analog of question Q4 for the college admission problem (when viewed from the students' side) can always be answered affirmatively. Hatfield et al. [29] contrasted the findings of Balinski and Sönmez by showing that no stable mechanism respects the improvement of school quality. Sönmez and Switzer [53] 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 [36] have obtained results of a similar flavor in a school-choice model with minimal-access rights.

Roth et al. [50] deal with the property of respecting improvement in connection with 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.

Closest to our work is the paper by Biró et al. [9] who focused on the classical Shapley–Scarf model and investigated how different solution concepts behave when the value of an agent’s house increases. They proved that both the strict core and the set of competitive allocations satisfy the property of respecting improvements, although this is no longer true when the lengths of trading cycles are bounded by some constant. We remark that Biró et al. [9] were not able to show that the property of respecting improvement holds for the core of housing markets. In fact, they posed questions Q4 and Q6 as open problems. We answer both of these questions affirmatively.

1.3 Organization.

Section 2 contains all definitions necessary for our model. Section 3 deals with the decision problems associated with questions Q1, Q2, and Q3, and their computational complexity. In Section 4 we present our results on the property of respecting improvements in relation to the core of housing markets, that is, questions Q4 and Q5: Sections 4.1 and 4.2 contain our main technical result, Theorem 2, while Section 4.3 deals with the computational complexity of the decision problem associated with question Q5. In Section 5 we study the respecting improvement property in the context of STABLE ROOMMATES, i.e., question Q6. We close in Section 6 with some questions for future research.

In an appendix we further provide some loosely related results: Appendix A contains an adaptation of the TTC algorithm for partially ordered preferences. Appendix B deals with the variants of questions Q1–Q3 for the strict core in a setting where agents’ preferences are weak orders. Finally, Appendix C contains an inapproximability result on the problem of maximizing the number of agents involved in trading in some core allocation.

2 Preliminaries.

Here we describe our model, and provide all the necessary notation. Information about the organization of this paper can be found at the end of this section.

2.1 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)⁶ *partial ordering* \prec on a set of alternatives is an irreflexive, antisymmetric and transitive relation.

⁶ Throughout the paper we will use the term *partial ordering* in the sense of an irreflexive (or strict) partial ordering.

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: an agent may be indifferent between a and b , and also between b and c , but strictly prefer a to c . Such preferences result in so-called *semiorders*, a special case of partial orders.

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.⁷ An organ is considered better than another if it is better with respect to both of these factors, leading to partial orders.

2.2 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 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 interpret $a \prec_c b$ (or $a \preceq_c b$) as agent c *preferring* (or *weakly preferring*, respectively) the house owned by agent b to the house of agent 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.

⁷ In fact, these are the two factors for which patients in the UK program can set acceptability thresholds [8].

3 Allocations in the core with arc restrictions.

We 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.

The input of the ARC IN CORE problem is a housing market $H = (N, \{\prec_a\}_{a \in N})$ and an arc (a, b) in G^H , and its task is to 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 each agent has 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. These results are in stark contrast to the polynomial-time solvability of the problem of finding a stable matching with forced and forbidden edges in an instance of STABLE ROOMMATES [26].

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 $G_{X \prec}^H$ 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 graph D with vertex set V and arc set A , we construct a housing market H as follows (see Fig. 1 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 .

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




Fig. 1. Illustration of the housing market H constructed in the NP-hardness proof for ARC IN CORE. Here and everywhere else we depict markets through their acceptability graphs with all loops omitted. Preferences are indicated by numbers along the arcs; 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.

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 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 the set $\{c_i \mid v_i \in V_1\} \cup \{d_i \mid v_i \in V_2\}$ is allocated its first choice by X . Therefore 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.

Theorem 1 shows that there is a computational gap between the strict core and the core: even though all three problems considered in Theorem 1 are NP-complete even if agents' preferences are strict, the corresponding problems become computationally tractable for the strict core. This is trivial if the preferences are strict, since in that case the strict core contains a unique allocation [48]. If preferences are weak orders (that is, if each agent orders the houses she finds acceptable in a linear order allowing ties), then the set of houses an agent can obtain in a strict core allocation can be computed in polynomial time based on the characterization of the strict core by Quint and Wako [47]; see Appendix B for details. We remark that since the characterization of Quint and Wako crucially depends on weak orders, the same approach does not work when preferences are partial orders.

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) $a \prec'_q b$ if and only if $a \prec_q b$ for each $a, b \in N \setminus \{p\}$, and (ii) if $a \prec_q p$ (or $a \preceq_q p$), then $a \prec'_q p$ (or $a \preceq'_q p$, respectively) for each $a \in N$. We will also say that a housing market is a p -improvement of H , if it can be obtained by a sequence of (p, q_i) -improvements for a series q_1, \dots, q_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 $\Phi(H)$:

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

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 strongly satisfies the RI-best property. In fact, Theorem 2 (proved in Section 4.2) states a slightly stronger statement.

Theorem 2. *For any allocation X in the core of 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 $O(|H|)$ time.*

Corollary 3. *The core of housing markets strongly satisfies the RI-best property.*

By contrast, we show that the RI-worst property does not hold for the core.

Proposition 4. *The core of housing markets violates the RI-worst property, even if agents' preferences are strict orders.*

Proof. Let $N = \{a, b, c, p, q\}$ be the set of agents. The preferences indicated in Figure 2 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, we have $(b, a) \in X$. 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 2.




Fig. 2. The housing markets H and H' in the proof of Proposition 4. For both H and H' , the allocation represented by bold (and blue) arcs yields the worst possible outcome for p in any core allocation of the given market.

We remark that Corollary 3 and Proposition 4 illuminate both the similarities of and the contrast between the properties of the core and the strict core. Recall that for strict preferences, there is a unique allocation in the strict core, and in this case the results of Biró et al. [9] show that the strict core strongly satisfies both the RI-best and the RI-worst properties. On the one hand, Proposition 4 hence shows a sharp difference between the core and the strict core with respect to the RI-worst property. On the other hand, Corollary 3 is very close to the analogous findings by Biró et al [9]: both results establish that the given solution concept (the core or the strict core) satisfies the RI-best property. See also Table 2 in Section 6 for a comparison of the core and the strict core in relation to the RI-best and RI-worst properties. We remark that despite the proximity of Corollary 3 with the results by Biró et al. [9] for the strict core, they are independent in the sense that neither of them implies the other.

We describe our algorithm for proving Theorem 2 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 2, we need some notation.

Sub-allocations and their envy graphs. Given a housing market $H = (N, \{\prec_a\}_{a \in N})$ and two subsets U and V of agents in N with $|U| = |V|$, we say that a set Y of arcs in $G^H = (N, E)$ is a *sub-allocation* from U to V in H , if

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

Note that Y forms a collection of mutually vertex-disjoint cycles and paths P_1, \dots, P_k in G^H , with each path P_i leading from a vertex of U to a vertex of V . Moreover, the number of paths in this collection is $k = |U \Delta V|$, where Δ stands for the symmetric difference operation. We call U the *source set* of Y , and V its *sink set*. See Figure 3 for an illustration.

Given a sub-allocation Y from U to V in H , we say that an arc $(a, b) \in E$ is *Y -augmenting*, if either $a \in V$ or $Y(a) \prec_a b$. 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 $G_{Y, \prec}^H$. We say that the sub-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 2. Let q_1, \dots, q_k denote the agents for which H' can be obtained from H by a series of (p, q_i) -improvements, $i = 1, \dots, k$. Observe that we can assume w.l.o.g. that the agents q_1, \dots, q_k are all distinct.




Fig. 3. Illustration for the concept of sub-allocation. The arcs sets Y_1 and Y_2 , shown with bold, teal lines, are both sub-allocations from $\{p\}$ to $\{a\}$ in the depicted housing market (as usual, loops are omitted). Source and sink vertices of Y are depicted with a white black diamond, respectively. For each of Y_1 and Y_2 , we show the corresponding envy arcs (i.e., the arcs in the corresponding envy graphs) with dashed, red lines; as can be seen, Y_1 is stable while Y_2 is not.

Algorithm HM-Improve. For a pseudocode description, see Algorithm 1, and for an example demonstrating the algorithm see Example 5.

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' . Let Q denote that set of those agents among q_1, \dots, q_k that in H' prefer p 's house to the one they obtain in allocation X , that is,

$$Q = \{q_i : X(q_i) \prec'_{q_i} p, 1 \leq i \leq k\}.$$

Observe that if an arc is an X -augmenting arc in H' but not in H , then it must be an arc of the form (q, p) where $q \in Q$. Therefore any cycle that blocks X in H' must contain an arc from $\{(q, p) : q \in Q\}$, as otherwise it would block X in H as well.

HM-Improve proceeds by modifying the housing market: for each $q \in Q$, 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 (the preferences of p remain unchanged). 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) for each $q \in Q$ with a new vertex corresponding to agent \tilde{q} , i.e., replacing the arc (q, p) with arcs (q, \tilde{q}) and (\tilde{q}, p) . For an illustration of the construction, see Figure 4. Let $\tilde{Q} = \{\tilde{q} : q \in Q\}$, $\tilde{N} = N \cup \tilde{Q}$, and let us denote by \tilde{E} be the set of arcs in \tilde{G} .

Initialization. Let $Y = X \setminus \{(q, X(q)) : q \in Q\} \cup \{(q, \tilde{q}) : q \in Q\}$ in \tilde{G} . Observe that Y is a sub-allocation in \tilde{H} with source set $\{X(q) : q \in Q\}$ and sink set \tilde{Q} . 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 sub-allocation Y and the set R of irrelevant agents. It will maintain the property that Y is a sub-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 set of Y changes quite freely during the iteration, the sink set always remains a subset of \tilde{Q} .

At each iteration, HM-Improve performs the following steps (see Figure 5 for illustration):

- Step 1.** Let U be the source set of Y , and V its sink set. If $U = V$, then the iteration stops.
- Step 2.** Otherwise, if there exists a Y -augmenting arc (s, u) in $\tilde{G}_{Y \prec}$ entering some source vertex $u \in U$ (note that $s \in \tilde{N} \setminus R$), then proceed as follows.
- (a) If $s \notin V$, 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 sub-allocation from $U \setminus \{u\} \cup \{u'\}$ to V in $\tilde{H} - R$.




Fig. 4. The housing market H of Example 5 and the modified housing market \tilde{H} constructed by Algorithm HM-Improve based on the p -improvement of H where q_1 and q_2 change their preferences so that q_1 comes to prefer p to a , and q_2 comes to prefer p to d . We depicted the core allocation X for H using blue lines, and we depicted the corresponding sub-allocation Y , as constructed by Algorithm HM-Improve in its initialization step, using teal lines. Sub-allocation Y has two sources, a and d , highlighted by diamonds, and two sinks, \tilde{q}_1 and \tilde{q}_2 . Envy arcs for both the original allocation X in H and the sub-allocation Y in \tilde{H} are shown using red, dashed lines.

- (b) If $s \in V$, then simply add the arc (s, u) to Y . In this case Y becomes a sub-allocation from $U \setminus \{u\}$ to $V \setminus \{s\}$ in $\tilde{H} - R$.

Step 3. Otherwise, let u be any vertex in $U \setminus V$ (not entered by any arc in $\tilde{G}_{Y \leftarrow}$), and 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 sub-allocation from $U \setminus \{u\} \cup \{u'\}$ to V in $\tilde{H} - R$.

Output. Let Y be the sub-allocation at the end of the above iteration, $U = V$ its source and sink set, and R the set of irrelevant agents. Note that $\tilde{Q} \setminus R \setminus U$ may contain at most one agent. Indeed, if $\tilde{q} \in \tilde{Q} \setminus R \setminus U$, then Y must contain the unique arc leaving \tilde{q} , namely (\tilde{q}, p) ; therefore, by $\delta_{\tilde{V}}(p) \leq 1$, at most one such agent \tilde{q} can exist.

To construct the desired allocation X' , the algorithm first applies the variant of the TTC algorithm that can deal with partial order preferences, described in Appendix A, to the sub-market $H'_{R \cap N}$ of H' when restricted to the set of irrelevant agents. This algorithm computes an allocation X_R in the core of $H'_{R \cap N}$.

HM-Improve next deletes all agents in \tilde{Q} . Since any agent in $\tilde{Q} \cap U = \tilde{Q} \cap V = V$ has zero in- and outdegree in Y , there is no need to modify our sub-allocation when deleting such agents; the same applies to agents in $\tilde{Q} \cap R$. By contrast, if there exists an agent $\tilde{q} \in \tilde{Q} \setminus R \setminus U$, then Y must contain the unique incoming and outgoing arcs of \tilde{q} , and therefore the algorithm replaces the arcs (q, \tilde{q}) and (\tilde{q}, p) with the arc (q, p) . This way we obtain an allocation on the submarket of H' on agents set $N \setminus R$.

Finally, HM-Improve outputs an allocation X' defined as

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

Let us now illustrate how Algorithm HM-Improve works on an example.






Fig. 5. Illustration of the possible steps performed during the iteration by HM-Improve. The edges of the current sub-allocation Y are depicted using bold, teal lines, while edges of the envy graph $\tilde{G}_{Y \prec}$ are shown by dashed, red lines. As in Figure 3, source and sink vertices of Y are depicted with a white black diamond, respectively. Vertices of R as well as all edges incident to them are shown in grey.

Example 5. Let us consider the housing market H shown in Figure 4, and let X denote the allocation in the core of H depicted, i.e., X consists of the cycles (p, j) , (a, b, c, q_1) , (d, q_2) , and (e, f, g, h, i) . Consider now the p -improvement H' of H where both q_1 and q_2 place p as their second favorite choice (instead of the third one).

The algorithm starts by checking whether X is in the core of H' , and finds that—since arcs (q_1, p) and (q_2, p) have become X -augmenting arcs—allocation X admits the blocking cycle (q_1, p, h) . Thus, the algorithm proceeds with modifying the housing market by subdividing the arcs (q_1, p) and (q_2, p) with newly added agents \tilde{q}_1 and \tilde{q}_2 .

In the initialization phase, algorithm HM-Improve constructs the sub-allocation Y based on X in the modified housing market \tilde{H} , as seen on Figure 4; we repeat the figure to the right here. Its source set is $U = \{a, d\}$ and its sink set is $V = \tilde{Q} = \{\tilde{q}_1, \tilde{q}_2\}$. The set of irrelevant agents is set to $R = \emptyset$. Then the algorithm starts iterating Steps 1–3 with the following results:

1. Considering the initial market \tilde{H} as shown in Figure 4, the algorithm finds that both sources, a and d , are entered by some Y -augmenting arc, namely by (q_1, d) and (i, a) . It may choose either one of these arcs to proceed with; we consider the course of the algorithm when it starts with the arc (q_1, d) : it replaces (q_1, \tilde{q}_1) with (q_1, d) in Y , so the source set becomes $U = \{a, \tilde{q}_1\}$, while the sink set remains $V = \{\tilde{q}_1, \tilde{q}_2\}$. The resulting sub-allocation is depicted to the right.




Algorithm 1 Algorithm HM-Improve

Input: housing market $H = (N, \prec)$, its p -improvement $H' = (N, \prec')$ for some agent p , and an allocation X in the core of H .


Output: an allocation X' in the core of H' such that $X(p) \prec_p X'(p)$ or $X(p) = X'(p)$.

- 1: **if** X is in the core of H' **then return** X
 - 2: Set $Q = \{a \in N : \prec_a \neq \prec'_a \text{ and } X(a) \prec'_a p\}$.
 - 3: Initialize housing market $\tilde{H} := H$.
 - 4: **for all** $q \in Q$ **do**
 - 5: Add new agent \tilde{q} to \tilde{H} , preferring only p to her own house.
 - 6: Replace p with \tilde{q} in the preferences of q in \tilde{H} .
 - 7: Set $\tilde{Q} = \{\tilde{q} : q \in Q\}$. $\triangleright \tilde{H}$ is now defined.
 - 8: Create sub-allocation $Y := X \setminus \{(q, X(q)) : q \in Q\} \cup \{(q, \tilde{q}) : q \in Q\}$.
 - 9: Set U and V as the source and sink set of Y , resp., and set $R := \emptyset$.
 - 10: **while** $U \neq V$ **do**
 - 11: **if** there exists an arc (s, u) in the envy graph $\tilde{G}_{Y, \prec}$ with $u \in U$ **then**
 - 12: **if** $s \notin V$ **then**
 - 13: Set $u' := Y(s)$, and update $Y \leftarrow Y \setminus \{(s, u')\} \cup \{(s, u)\}$ and $U \leftarrow U \setminus \{u\} \cup \{u'\}$.
 - 14: **else** \triangleright Case $s \in V$.
 - 15: Update $Y \leftarrow Y \cup \{(s, u)\}$, $U \leftarrow U \setminus \{u\}$ and $V \leftarrow V \setminus \{s\}$.
 - 16: **else** \triangleright No arc enters U in the envy graph $\tilde{G}_{Y, \prec}$.
 - 17: Pick any agent $u \in U \setminus V$, and set $u' := Y(u)$.
 - 18: Update $Y \leftarrow Y \setminus (u, u')$, $U \leftarrow U \setminus \{u\} \cup \{u'\}$ and $R \leftarrow R \cup \{u\}$.
 - 19: Compute a core allocation X_R in the submarket $H'_{R \cap N}$.
 - 20: **if** $\tilde{Q} \setminus R \setminus U = \emptyset$ **then set** $X' := X_R \cup Y$.
 - 21: **else set** $X' := X_R \cup Y \setminus \{(q, \tilde{q}), (\tilde{q}, p)\} \cup \{(q, p)\}$ where $\tilde{Q} \setminus R \setminus U = \{\tilde{q}\}$.
 - 22: **return** the allocation X' .
-


2. Next, the algorithm finds that only the source agent a (from among the source set $U = \{a, \tilde{q}_1\}$) is entered by some Y -augmenting arc, namely, by (i, a) . It replaces (i, e) with (i, a) in Y yielding the sub-allocation shown to the right; the source set becomes $U = \{e, \tilde{q}_1\}$.




3. Next, the algorithm finds that neither of the sources e and \tilde{q}_1 is entered by an Y -augmenting arc; hence, it takes the unique source agent that is not a sink, namely e , and declares it irrelevant by setting $R = \{e\}$, and removing e from the market \tilde{H} . The arcs of Y incident to e are removed from Y , thus the source set becomes $U = \{f, \tilde{q}_1\}$.



4. Next, the algorithm finds that only the source agent f (from among the source set $U = \{f, \tilde{q}_1\}$) is entered by some Y -augmenting arc, namely, by (b, f) . It replaces (b, c) with (b, f) in Y , and the source set becomes $U = \{c, \tilde{q}_1\}$.



5. Next, the algorithm finds that only the source agent c (from among the source set $U = \{f, \tilde{q}_1\}$) is entered by some Y -augmenting arc, namely, the arcs (q_2, c) and (g, c) . Here we consider the course of the algorithm when it chooses the arc (q_2, c) : it replaces (q_2, \tilde{q}_2) with (q_2, c) in Y , and the source set becomes $U = \{\tilde{q}_2, \tilde{q}_1\}$.




At this point, the algorithm detects that the source set U equals the sink set $V = \tilde{Q}$, and stops the iteration. It computes a core allocation for the submarket $H'_{R \cap N}$ of irrelevant agents; since $R = \{e\}$, this allocation consists of the single arc (e, e) . Since $\tilde{Q} \setminus R \setminus U = \emptyset$, it outputs the allocation $Y \cup \{(e, e)\}$ in which agents trade along the cycles (p, j) , (q_1, c, q_2, d) , and (a, b, f, g, h, i) ; see Figure 6.


Consider now an alternative course for the algorithm when, after the fourth iteration (see the figure next to Step 4 above), in the fifth iteration step the arc (g, c) gets chosen instead of the arc (q_2, c) ; we omit the corresponding figures for Steps 5* and 6* in this alternative course:

- 5*. The algorithm replaces (g, h) with (g, c) in Y , and the source set becomes $U = \{h, \tilde{q}_1\}$.
6*. The algorithm replaces (p, j) with (p, h) in Y , and the source set becomes $U = \{j, \tilde{q}_1\}$.


- 7*. The algorithm finds that no Y -augmenting arc enters either of the sources, and thus removes agent j , the only agent in $U \setminus V$, together with the arc (j, p) . Hence, the set of irrelevant agents is set to $R = \{e, j\}$, and the source set becomes $U = \{p, \tilde{q}_1\}$.




- 8*. The algorithm finds that only the source p is entered by some Y -augmenting arc, namely by the arcs (\tilde{q}_1, p) and (\tilde{q}_2, p) . It chooses one of them, say (\tilde{q}_1, p) . Since \tilde{q}_1 is a sink, it adds (\tilde{q}_1, p) to Y , removes p from the source set, and removes \tilde{q}_1 from the sink set. This yields $U = \{\tilde{q}_1\}$ and $V = \{\tilde{q}_2\}$.



- 9*. The unique source \tilde{q}_1 is not entered by any Y -augmenting arc, therefore the algorithm declares it irrelevant by setting $R = \{e, j, \tilde{q}_1\}$, and removes it from the market. The arc (\tilde{q}_1, p) is removed from Y , and the source set becomes $U = \{p\}$.



- 10*. The algorithm finds that the (unique) source p is entered by a unique Y -augmenting arc, namely (\tilde{q}_2, p) . Since \tilde{q}_2 is a sink (recall that $V = \{\tilde{q}_2\}$ at this point), it adds (\tilde{q}_2, p) to Y , removes p from the source set, and removes \tilde{q}_2 from the sink set. This yields $U = V = \emptyset$.



At this point, the algorithm detects that the source set U equals the sink set V (both empty), and stops the iteration. It computes a core allocation for the submarket $H'_{R \cap N}$ of

irrelevant agents; since $R = \{e, j, \tilde{q}_1\}$, this allocation consists of the arcs (e, e) and (j, j) . Since $\tilde{Q} \setminus R \setminus U = \{\tilde{q}_2\}$, it outputs the allocation $Y \setminus \{(q_2, \tilde{q}_2), (\tilde{q}_2, p)\} \cup \{(q_2, p) \cup \{(e, e), (j, j)\}$ in which agents trade along the single cycle $(p, h, i, a, b, f, g, c, q_1, d, q_2)$; see Figure 6.

4.2 Correctness of algorithm HM-Improve.

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

Lemma 6. *At each iteration, sub-allocation Y is stable in $\tilde{H} - R$.*

Proof. The proof is by induction on the number n of iterations performed. For $n = 0$, suppose for the sake of contradiction that C is a cycle in $\tilde{G}_{Y \prec}$. First note that C cannot contain any agent in \tilde{Q} , since the unique arc entering \tilde{q} , that is, the arc (q, \tilde{q}) , is contained in Y by definition. Hence, C is also a cycle in H . Moreover, recall that initially $Y(a) = X(a)$ for each agent $a \in N \setminus Q$, and by the definition of Q , we also know $X(q) \prec_q \tilde{q} = Y(q)$ for each $q \in Q$. Therefore, any arc of C is an X -augmenting arc as well, and thus C is a blocking cycle for X in H . This contradicts our assumption that X is in the core of H . Hence, Y is stable in \tilde{H} at the beginning; note that $R = \emptyset$ initially.

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, and let Y' and R' be the sub-allocation and the set of irrelevant agents obtained after the modifications in this iteration. Let also U and V (U' and V') denote the source and sink set of Y (of Y' , respectively). By induction, we may 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 three cases:

- (a) The algorithm modifies Y in Step 2(a), by using a Y -augmenting arc (s, u) where $s \notin V$; then $R' = R$. Note that $s \in$ 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 2(b), by using a Y -augmenting arc (s, u) where $s \in V$; then $R' = R$. First observe that $V \subseteq \tilde{Q}$, as the only way the sink set of Y can change is when an agent ceases to be a sink of the current sub-allocation due to the application of Step 2(b). Thus, $s \in V$ implies $s \in \tilde{Q}$, which means that (s, u) must be the unique




Fig. 6. Two allocation in the core of housing market H' of Example 5, computed by Algorithm HM-Improve. The figure to the left depicts the allocation obtained by Steps 1–5, while the figure to the right depicts the allocation obtained when Steps 1–4 are followed by Steps 5*–10*.

arc (s, p) leaving s . Hence, adding (s, u) to Y amounts to deleting the arc (s, u) from the envy graph $\tilde{G}_{Y \prec}$.

- (c) The algorithm modifies Y in Step 3, by adding an agent $u \in U \setminus V$ 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.

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 7. *Let Y_1 and Y_2 be two sub-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 .*

In the next two lemmas, we prove that HM-Improve produces a core allocation. We start by explaining why irrelevant agents may not become the cause of instability in the housing market.

Lemma 8. *At the end of algorithm HM-Improve, there does not exist an arc $(a, b) \in \tilde{E}$ such that $a \notin R$, $b \in R$ 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. Suppose that HM-Improve adds b to R during the n -th iteration, and let Y_n be the sub-allocation at the beginning of the n -th iteration. By Proposition 7, 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.

Lemma 9. *The output of HM-Improve is an allocation in the core of H' .*

Proof. Let Y and R be the sub-allocation and the set of irrelevant agents, respectively, at the end of algorithm HM-Improve, and let U be the source set of Y . To begin, we prove it formally that the output X' of HM-Improve is an allocation for H' .

Since HM-Improve stops only when $U = V$, the arc set Y forms a collection of mutually vertex-disjoint cycles in $\tilde{H} - R$ that covers each agent in $\tilde{N} \setminus R \setminus U$; agents of U have neither incoming nor outgoing arcs in Y . As no agent outside \tilde{Q} can become a sink of Y , we know that $U = V \subseteq \tilde{Q}$.

First, assume $\tilde{Q} \setminus R \setminus U = \emptyset$, that is, $\tilde{Q} \setminus R = U = V$. In this case, Y is the union of cycles covering each agent in $N \setminus R$ exactly once. Hence, Y is an allocation in the submarket of H' restricted to agent set $N \setminus R$, i.e., $H'_{N \setminus R}$.

Second, assume $\tilde{Q} \setminus R \setminus U \neq \emptyset$. In this case, Y is the union of cycles covering each agent in $\tilde{N} \setminus R \setminus V$ exactly once. Let \tilde{q} be an agent in $\tilde{Q} \setminus R \setminus V$. Since \tilde{q} is not a sink of Y , is not irrelevant, and has a unique outgoing arc to p , we know $(\tilde{q}, p) \in Y$. As Y cannot contain two arcs entering p , this proves that $\tilde{Q} \setminus R \setminus V = \tilde{Q} \setminus R \setminus U = \{\tilde{q}\}$. Moreover, since the unique arc entering \tilde{q} is from q , we get $(q, \tilde{q}) \in Y$. Therefore, the arc set $Y \setminus \{(q, \tilde{q}), (\tilde{q}, p)\} \cup \{(q, p)\}$ is an allocation in $H'_{N \setminus R}$.

Consequently, as X_R is an allocation on $H'_{R \cap N}$, we obtain that X' is indeed an allocation in H' in both cases.

Now we prove that 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 $H'_{R \cap N}$ and hence is acyclic.

Claim 1 (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. Let us suppose first that $(a, b) \notin \{(q, p) : q \in Q\}$: then (a, b) is an arc in $G^{\tilde{H}}$. If $a \notin Q$ or $Y(a) \notin \tilde{Q}$, then $Y(a) = X'(a)$ and thus the claim follows immediately. If $a \in Q$ and $Y(a) = \tilde{a} \in \tilde{Q}$, then $X'(a) = p \prec'_a b$ implies that a prefers b to $Y(a) = \tilde{a}$ in \tilde{H} as well, that is, (a, b) is Y -augmenting.

Suppose now that $(a, b) = (q, p)$ for some $q \in Q$. We finish the proof of the claim by showing that (q, p) is not X' -augmenting if $q \notin R$ (recall that we assumed $q = a \notin R$).

First, if $\tilde{q} \notin U$, then necessarily $\{(q, \tilde{q}), (\tilde{q}, p)\} \subseteq Y$, and so $(q, p) \in X'$, which means that (q, p) is not X' -augmenting.

Second, if $\tilde{q} \in U$, then consider the iteration in which \tilde{q} became a source for our sub-allocation, and let Y_n denote the sub-allocation at the end of this iteration. Agent \tilde{q} can become a source either in Step 2(a) or in Step 3, since Step 2(b) always results in one agent being deleted from the source set without a replacement. Recall that the only arc entering \tilde{q} is (q, \tilde{q}) . If \tilde{q} became the source of Y_n in Step 2(a), then we know $\tilde{q} \prec'_q Y_n(q)$. By Proposition 7, this implies $\tilde{q} \prec'_q Y(q)$. By the construction of \tilde{H} , we obtain that q prefers $Y(q) = X'(q)$ to p in H' , so (q, p) is not X' -augmenting. Finally, if agent \tilde{q} became the source of Y_n in Step 3, then this implies $q \in R$, which contradicts our assumption $a = q \notin R$.

Our claim implies that $G_{X', \prec}^{H'}[N \setminus R]$ is a subgraph of $\tilde{G}_{Y, \prec}$ and therefore it is acyclic by Lemma 6. 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 8. Thus $G_{X', \prec}^{H'}$ is acyclic and X' is in the core of H' .

The following lemma, the last one necessary to prove Theorem 2, 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. *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 Appendix A can be computed in $O(|H|)$ time. Hence, the overall running time of our algorithm is $O(|H|) + O(|E|) = O(|H|)$.

We are now ready to prove Theorem 2.

Proof (of Theorem 2). Lemma 10 shows that algorithm HM-Improve runs in linear time, and by Lemma 9 its output is an allocation X' in the core of H' . It remains to prove that either $X'(p) = X(p)$ or p prefers X' to X . Observe that it suffices to show $p \notin R$, by Proposition 7.

For the sake of contradiction, assume that HM-Improve puts p into the set of irrelevant vertices at some point, during an execution of Step 3. Let Y denote the sub-allocation at the beginning of this step, and let V be its sink set. Clearly, $V \neq \emptyset$ (as in that case the source and the sink set of Y would coincide). Recall also that $V \subseteq \tilde{Q}$. Thus, there exists some $\tilde{q} \in V \subseteq \tilde{Q}$. However, then (\tilde{q}, p) is an Y -augmenting arc by definition, entering p , which contradicts our assumption that the algorithm put p into the set of irrelevant agents in Step 3 of this iteration.

4.3 Strict improvement.

Looking at Theorem 2 and Corollary 3, 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. *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 (a^*, b^*) . Therefore, (\hat{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).

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 of them is either unmatched, or 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, illustrated in Examples 13 and 14.

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.*

Example 13. Let $N = \{a, b, c, d, e, p, q\}$ be the set of agents. The preferences indicated in Figure 7 define two housing markets H and H' such that H' is a (p, q) -improvement of H . Note that agent 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 14. Let $N = \{a, b, p, q\}$ be the set of agents. The preferences indicated in Figure 8 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.

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 if preferences are strict orders; hence, stable matchings do not strongly satisfy the RI-best property.




Fig. 7. The housing markets H and H' as instances of STABLE ROOMMATES with ties, in the Example 13. 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.




Fig. 8. The housing markets H and H' in Example 14. 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.

Proposition 15. *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 9 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.

Contrasting Propositions 12 and 15, 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.

Theorem 16. *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 $M(p) \preceq_p M'(p)$. Moreover, given H , H' and M it is possible to find such a matching M' in polynomial time, or conclude correctly that H' admits no stable matchings.*

Corollary 17. *In the STABLE ROOMMATES model with strict preferences stable matchings satisfy the RI-best property.*

Comparing our results for the core of housing markets and for stable matchings in the STABLE ROOMMATES model, we find that these solution concepts exhibit similarities as well as disparities in connection to the notion of respecting improvement. First, neither solution concept satisfies the RI-worst property (see Propositions 4 and 12). Second, both satisfy the RI-best property, as established by our main algorithmic results, Theorems 2 and 16; however, the core of housing markets *strongly* satisfies the RI-best property even if preferences are partial orders, as opposed to stable matchings in the STABLE ROOMMATES model, where it may happen




Fig. 9. Housing markets H and H' illustrating the proof of Proposition 15. 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.

that an improvement yields a market without any stable matchings (see Proposition 15), and the RI-best property holds only if preferences are strict, but fails if preferences can contain ties (see Proposition 12). Table 2 in Section 6 offers a comparison of these two models from the viewpoint of the property of respecting improvement.

We describe our algorithm for Theorem 16 in Section 5.1, and prove its correctness in Section 5.2.

5.1 Description of algorithm SR-Improve.

To prove Theorem 16 we are going to rely on the concept of proposal-rejection alternating sequences introduced by Tan and Hsueh [55], 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 18 as observed and used by Tan and Hsueh. Since the first claim of Lemma 18 is only implicit in the paper by Tan and Hsueh [55], we prove it for the sake of completeness.

Lemma 18 ([55]). *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 $M_{i-1}(\beta_i)$.

Proof (of the first statement of Lemma 18). We prove the statement 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 $M_i \triangle M_{i-1} = \{\{\alpha_i, \beta_i\}, \{\alpha_{i-1}, \beta_i\}\}$ 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.

We are now ready to describe algorithm SR-Improve; see Algorithm 2 for its pseudocode.

Algorithm SR-Improve.

Let $H = (N, \{\prec_a\}_{a \in N})$ be a housing market containing a stable matching M , 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.

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 [31]; 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, algorithm 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 $M' = M_k$.

5.2 Correctness of algorithm SR-Improve.

To show that algorithm SR-Improve is correct, we first state the following two lemmas.

Lemma 19. *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

Algorithm 2 Algorithm SR-Improve

Input: housing market $H = (N, \prec)$, its (p, q) -improvement $H' = (N, \prec')$ for two agents p and q , and a stable matching M in H .

Output: a stable matching M' in H' such that $M(p) \preceq_p M'(p)$ or $M(p) = M'(p)$, if H' admits some stable matching.

- 1: **if** M is stable in H' **then return** M
 - 2: **if** H' admits a stable matching **then** let M^* be any stable matching in H' .
▷ Use Irving's algorithm [31]
 - 3: **else return** "No stable matching exists for H' ."
 - 4: **if** $M(p) \preceq_p M^*(p)$ **then return** $M' := M^*$
 - 5: Create housing market \tilde{H} by deleting the agents $\{a \in N : a \preceq'_q p\}$ from $A(q)$ and vice versa.
 - 6: Set $i := 0$, $\alpha_0 := M(q)$, and $M_0 := M \setminus \{\alpha_0, q\}$
 - 7: **repeat** ▷ Computing a proposal-rejection sequence S .
 - 8: Set $i \leftarrow i + 1$.
 - 9: Set $B_i := \{b : \alpha_{i-1} \in A(b), b \text{ is unmatched in } M_{i-1} \text{ or } M_{i-1}(b) \prec_b \alpha_{i-1}\}$.
 - 10: **if** $B_i = \emptyset$ **then return** $M' := M_{i-1}$ ▷ S stops at $i - 1$.
 - 11: Set β_i as the agent most preferred by α_{i-1} in B_i .
 - 12: **if** β_i is unmatched in M_{i-1} **then return** $M' := M_{i-1} \cup \{\{\alpha_{i-1}, \beta_i\}\}$ ▷ S stops at i .
 - 13: Set $\alpha_i := M_{i-1}(\beta_i)$ and $M_i := M_{i-1} \cup \{\{\alpha_{i-1}, \beta_i\}\} \setminus \{\{\alpha_i, \beta_i\}\}$.
 - 14: **until** $\alpha_i = p$ **return** $M' := M_i \cup \{\{p, q\}\}$
-

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 18, 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 [28, Theorem 4.5.2]. Now, if S stops with the last induced matching M_k , then by the third statement of Lemma 18 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 .

Lemma 20. *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 18, 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 19, we know that p is not among the agents $\alpha_0, \beta_1, \alpha_1, \dots, \beta_{k-1}$. Thus, $M_{k-1}(p) = M_0(p) = M(p)$. 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 is not in $\{\alpha_{k-1}, \beta_k, p\}$), 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 18, 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 18 the partner that p receives in the matchings M_0, M_1, \dots, M_k can only get better for p , and thus $M(p) = M_0(p) \preceq'_p M_k(p) = M'(p)$.

We can now piece together the proof of Theorem 16.

Proof (of Theorem 16). From the description of SR-Improve and Lemma 20 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 19 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 18, we also know that the length of S is at most $2|E|$. Thus, the algorithm finishes in $O(|E|)$ time.

5.3 A note on strongly stable matchings in Stable Roommates.

Given an instance of STABLE ROOMMATES where preferences are not strict, strong stability is an alternative notion of stability based on the notion of weakly blocking pairs. Given a matching M in a housing market $H = (N, \{\prec_a\}_{a \in N})$, an edge $\{a, b\}$ in the acceptability graph G^H is *weakly blocking*, if (i) a is either unmatched or weakly prefers b to $M(a)$, and (ii) b is either unmatched or weakly prefers a to $M(b)$, and (iii) if a and b are both matched in M , then a prefers b to $M(a)$, or b prefers a to $M(b)$. If there is no weakly blocking pair for M , then M is *strongly stable*.

Note that a strongly stable matching for H can be thought of as an allocation that (i) contains only cycles of length at most 2, and (ii) does not admit a *weakly blocking* cycle of length at most 2. Recall that stable matchings correspond to the concept of core if we restrict allocations to pairwise exchanges; analogously, strongly stable matchings correspond to the concept of strict core for pairwise exchanges. Observe also that if agents' preferences are strict orders, then strong stability is equivalent with stability, or in other words, the strict core and the core coincide.

In view of Corollary 17, it is natural to ask whether the set of strongly stable matchings satisfies the RI-best property in the case when preferences may not be strict. The following statement answers this question in the negative. Interestingly, the result holds even in the STABLE MARRIAGE model, the special case of STABLE ROOMMATES where the acceptability graph is bipartite.

Proposition 21. *Strongly stable matchings in the STABLE MARRIAGE model do not satisfy the RI-best property, even if agents' preferences are weak orders.*

Proof. Consider the housing markets H and H' depicted in Figure 10; note that H' is a (p, q) -improvement of H . Note that the preferences in H are strict, but in H' agent q is indifferent between p and b .

First observe that the matching M shown in bold in the first part of Figure 10 is stable in H , so it is possible for p to be matched with its second choice, namely a , in a (strongly) stable matching in H . We claim that the best possible partner p can obtain in any strongly stable matching in H' is its third choice. To see this, first note that any matching containing $\{p, q\}$ is weakly blocked by $\{q, b\}$ in H , so p cannot be matched to its first choice, agent q , in any




Fig. 10. The housing markets H and H' in the proof of Proposition 21. For both H and H' , the allocation represented by bold arcs yields the best possible strongly stable matchings.

strongly stable matching in H' . Second, note that any matching M' containing $\{p, a\}$ must match q to its first choice (otherwise the pair $\{p, q\}$ weakly blocks M') and hence M' must match b to its third choice (so as not to form a blocking pair with it); however, then $\{a, b\}$ is a blocking pair for M' . Thus, p cannot be matched in any strongly stable matching of H' to its second choice, agent a , either.

By contrast, it is easy to verify that the matching shown in bold in the second part of Figure 10, matching p to its third choice, is strongly stable in H' . This proves our proposition.

6 Summary and further research.

		Housing market		STABLE ROOMMATES	
		Core	Strict core	Core	Strict core
Forced edge/arc:	strict pref.	NP-c (Thm. 1)	P [48]	P [26]	P [26]
	weak pref.	NP-c (Thm. 1)	P (Thm. 22)	NP-c [44]	P [39, 40]
	partial order pref.	NP-c (Thm. 1)	open	NP-c [44]	NP-c [32]
Forbidden edge/arc:	strict pref.	NP-c (Thm. 1)	P [48]	P [26]	P [26]
	weak pref.	NP-c (Thm. 1)	P (Thm. 22)	NP-c [21]	P [39, 40]
	partial order pref.	NP-c (Thm. 1)	open	NP-c [21]	NP-c [32]
Agent trading:	strict pref.	NP-c (Thm. 1)	P [48]	P [28]	P [28]
	weak pref.	NP-c (Thm. 1)	P (Thm. 22)	NP-c [44]	P [39, 42]
	partial order pref.	NP-c (Thm. 1)	open	NP-c [44]	NP-c [32]

Table 1. Summary of known results on the problems of finding an allocation in the core or strict core of a housing market or a STABLE ROOMMATES instance that additionally (i) contains a given arc (or edge), (ii) avoids a given arc (or edge), or (iii) includes a trading cycle containing a given agent. The table classifies each of these problems either as polynomial-time solvable (P) or NP-complete (NP-c for short), except for the cases whose computational complexity remains open. Recall that in an instance of STABLE ROOMMATES, we interpret the core as the set of stable matchings, and the strict core as the set of strongly stable matchings; these two notions coincide for strict preferences.

We have investigated questions about the notion of improvement in connection to the core of housing markets. Table 1 puts into context our algorithmic results on finding allocations with various restrictions in the core of housing markets; the table presents analogous results for the strict core of housing markets, as well as for the core and strict core of STABLE ROOMMATES

	Housing market		STABLE ROOMMATES	
	Core	Strict core	Core	Strict core
RI-best: strict pref.	✓(Cor. 3)	✓[9]	✓ [∅] (Cor. 17, Prop. 15)	✓ [∅] (Cor. 17, Prop. 15)
weak pref.	✓(Cor. 3)	✓ [∅] [9]	✗[9]	✗(Prop. 21)
partial order pref.	✓(Cor. 3)	open	✗[9]	✗(Prop. 21)
RI-worst: strict pref.	✗(Prop. 4)	✓[9]	✗[9]	✗[9]
weak pref.	✗(Prop. 4)	✓ [∅] [9]	✗[9]	✗[9]
partial order pref.	✗(Prop. 4)	open	✗[9]	✗[9]

Table 2. Summary of known results on the property of respecting improvement for the core and the strict core in housing markets and in the STABLE ROOMMATES model. Symbol ✓ signifies that the given solution concept strongly satisfies the given property (namely, RI-best or RI-worst), while symbol ✓[∅] means that the given property is satisfied, but not strongly satisfied. Symbol ✗ means that the given property fails to hold.

instances (when interpreted as the set of stable and strongly stable matchings, respectively). Table 2 summarizes our results regarding the property of respecting improvement both for housing markets and also for the STABLE ROOMMATES model; we also included the known facts about the strict core, mostly by Biró et al. [9]. We remark that several questions remain open in the case when agents’ preferences are partial orders; in fact, we are not aware of any result concerning the strict core of housing markets under such preferences.

Even though the property of respecting improvement is deeply connected to agents’ incentives in exchange markets, many solution concepts have not yet been studied from this aspect. A solution concept that seems interesting from this point of view is the set of stable half-matchings (or equivalently, stable partitions) in instances of STABLE ROOMMATES without a stable matching. Although Figure 11 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.




Fig. 11. An example where an agent’s improvement has a detrimental effect on the agent’s situation in a model where allocations are defined as half-matchings (see also [54]). Given a STABLE ROOMMATES instance with underlying graph (V, E) , a *half-matching* is a function $f : E \rightarrow \{0, \frac{1}{2}, 1\}$ that satisfies $\sum_{e=\{u,v\} \in E} f(e) \leq 1$ for each agent $v \in V$. The figure contains housing market H and its (p, q) -improvement H' , and a unique stable half-matching for each market; see [43] for the definition of stable half-matchings. 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' .

Acknowledgement

The authors thank the area editor, the associate editor and two anonymous referees for their insightful comments and suggestions.

Funding

Ildikó Schlotter is supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2021-2) and its János Bolyai Research Scholarship. The research reported in this paper and carried out by Tamás Fleiner at the Budapest University of Technology and Economics was supported by the “TKP2020, National Challenges Program” of the National Research Development and Innovation Office (BME NC TKP2020 and OTKA K143858) and by the Higher Education Excellence Program of the Ministry of Human Capacities in the frame of the Artificial Intelligence research area of the Budapest University of Technology and Economics (BME FIKP-MI/SC). Péter Biró gratefully acknowledges financial support from the Hungarian Scientific Research Fund, OTKA, Grant No. K143858, and the Hungarian Academy of Sciences, Momentum Grant No. LP2021-2.

Appendix A Top Trading Cycles for partial order preferences.

Here we present an adaptation of the Top Trading Cycles algorithm for the case when agents' preferences are represented as partial orders; this algorithm always finds an allocation in the core of the given housing market in linear time. We start by recalling how TTC works for strict preferences, propose a method to deal with partial orders, and finally discuss how the obtained algorithm can be implemented in linear time.

Strict preferences. If agents' preferences are represented by strict orders, then the TTC algorithm [52] 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 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 in D . 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 the set $A(a)$ of agents whose house is acceptable for 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 deleting a and adding those in-neighbors of a in H_b which have no out-neighbor still present in the market. 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 $O(|N| + \sum_{a \in N} |H_a|) = O(|H|)$.

Appendix B Arc restrictions for the strict core.

In this section we investigate the variants of our questions Q1, Q2, and Q3 for the strict core of housing markets. Recall that an allocation X in a housing market $H = (N, \{\prec_a\}_{a \in N})$ is in the *strict core* of H if there is no coalition S of agents with an allocation X' on S such that (i) $X(a) \preceq_a X'(a)$ for each agent $a \in S$, and (ii) $X(a) \prec_a X'(a)$ for at least one agent $a \in S$.

Recall that if agents' preferences are strict, then the strict core contains a unique allocation, and this allocation can be efficiently computed by the TTC algorithm [48]. Thus, it is trivial to decide whether an agent can obtain a given house in the unique allocation in the strict core.

When agents' preferences are weak orders, then the strict core can be empty [52]. However, there is a polynomial-time algorithm by Quint and Wako [47] that decides whether the strict core is empty when preferences are weak orders. We generalize this result by giving a polynomial-time algorithm for the following problem:

The ARC RESTRICTIONS FOR STRICT CORE problem in housing markets: given a housing market H , a set F^+ of *forced* arcs and a set F^- of *forbidden* arcs in the underlying graph G^H , find an allocation in the strict core of H that contains F^+ and is disjoint from F^- .

Note that the ARC RESTRICTIONS FOR STRICT CORE problem is a generalization of the problems underlying questions Q1–Q3 when interpreted in relation for the strict core: for a given agent $a \in N$ and an arc (a, b) in G^H , we can use an algorithm for the ARC RESTRICTIONS FOR STRICT CORE problem in order to decide whether

- some allocation in the strict core contains (a, b) , by setting $F^+ = \{(a, b)\}$ and $F^- = \emptyset$;
- some allocation in the strict core avoids (a, b) , by setting $F^+ = \emptyset$ and $F^- = \{(a, b)\}$;
- agent a is trading in some allocation in the strict core, by setting $F^+ = \emptyset$ and $F^- = \{(a, a)\}$.

The following theorem is obtained through a straightforward modification of the algorithm by Quint and Wako for deciding the emptiness of the strict core.

Theorem 22. *If agents' preferences are weak orders, then the ARC RESTRICTIONS FOR STRICT CORE problem can be solved in polynomial time.*

Proof. Let $H = (N, \{\prec_a\}_{a \in N})$ be the housing market with underlying graph $G^H = (N, E)$ given as our input, together with arc sets $F^+ \subseteq E$ and $F^- \subseteq E$.

We need the following concept from graph theory: we say that a set V of vertices in a directed graph is an *absorbing set*, if (i) no arc leaves V and (ii) V is strongly connected, meaning that for each vertices $v_1, v_2 \in V$ there are paths from v_1 to v_2 and from v_2 to v_1 in the graph. It is easy to see that two absorbing sets in a directed graph are either identical or vertex-disjoint. Recall that an arc $(a, b) \in E$ is *undominated*, if there exists no agent b' such that $b \prec_a b'$. Let $U(a)$ denote the set of all undominated arcs in G^H leaving some agent a , and let $U = \bigcup_{a \in N} U(a)$. Let T denote the union of all absorbing sets in the subgraph (N, U) of undominated arcs within G^H , and let G_T denote the subgraph of (N, U) induced by vertices of T , i.e., $G_T = (T, \bigcup_{t \in T} U(t))$. Informally speaking, G_T contains the “top” agents and their most-preferred choices.

Quint and Wako [47] proved that if agents' preferences are weak orders, then an allocation X in H is in the strict core of H if and only if

- for each agent $t \in T$, the arc of X leaving t is undominated, and
- $X[N \setminus T]$ is an allocation in the strict core of $H_{N \setminus T}$.

Using this characterization, it is straightforward to see that the following algorithm solves the ARC RESTRICTIONS FOR STRICT CORE problem:

- Step 1.** Compute the set T and the subgraph G_T .
- Step 2.** If F^+ contains an arc running between T and $N \setminus T$, then return ‘No’.
- Step 3.** Find a set $C \subseteq \bigcup_{t \in T} U(t)$ of arcs in G_T that is an allocation in the submarket H_T , and (i) contains all arcs of $F^+[T]$ and (ii) is disjoint from F^- . Return ‘No’ if no such set C exists.
- Step 4.** Use a recursive call to compute an allocation X' in the strict core of the submarket $H_{N \setminus T}$ that (i) contains all edges of $F^+[N \setminus T]$, and (ii) is disjoint from F^- . Return ‘No’ if no such allocation X' exists.
- Step 5.** Return the allocation $X = C \cup X'$.

From the above characterization of the strict core by Quint and Wako, the correctness of the above algorithm follows immediately. Hence, it remains to check its running time.

Step 1 can be performed in linear time using standard algorithms on directed graphs, see e.g., [38]. Step 2 takes linear time as well. The task in Step 3 can be performed by computing a maximum-weight matching in the following bipartite graph \widehat{G}_T : the vertex set of \widehat{G}_T is $\{t_1, t_2 : t \in T\}$, and for each arc (t, t') in G_T with $(t, t') \notin F^-$ we add an edge (t_1, t'_2) to \widehat{G}_T ; the weight of this edge is 2 if $(t, t') \in F^+[T]$, otherwise it is 1. Then allocations⁸ in G_T disjoint from F^- correspond bijectively to perfect matchings in \widehat{G}_T , and moreover, allocations in G_T disjoint from F^- and containing all arcs of $F^+[T]$ correspond bijectively to (perfect) matchings in \widehat{G}_T of weight $|T| + |F^+[T]|$. Since we can compute a maximum-weight matching in \widehat{G}_T in $|T|^3$ time using the Hungarian algorithm (see e.g., [38]), the recursion in Step 4 yields an overall running time of $|N|^3$.

Note that Theorem 22 is in sharp contrast with our results for the core: by Theorem 1, given a housing market H , it is NP-hard to decide whether a given arc of G^H can be contained in some core allocation of H , even if agents’ preferences are strict orders. Hence, in this aspect, we can perceive a computational gap between the core the strict core in housing markets with strictly or weakly ordered preferences.

We remark that we are not aware of any result that would settle the computational complexity of the ARC RESTRICTIONS FOR STRICT CORE problem in the case when agents’ preferences are partial orders. In fact, even deciding the emptiness of the strict core seems to be a problem whose computational complexity is open.

Appendix C 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].

⁸ By an allocation in G_T , we mean an edge set in G_T that is an allocation in the submarket H_T .

In Theorem 23 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 note that Biró and Cechlárová [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 23. Instead, our reduction relies on ideas we use to prove Theorem 1.

Theorem 23. *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

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

that is, we replace the arc (a^*, b^*) with the path $(a^*, p_1, p_2, \dots, p_K, b^*)$; see Figure 12 for an illustration. Let $N' = N \cup \{p_1, \dots, p_K\}$. Formally, we define preferences \prec'_a for each agent $a \in N'$ as follows: first, \prec'_a is identical to \prec_a for each $a \in N \setminus \{a^*\}$; second, a^* only prefers the house of agent p_1 to her own house; third, 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 the arc (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 X' must contain at least one vertex from $\{p_1, p_2, \dots, p_K\}$ by $|N| = 4n + 4$, which by the previous paragraph implies that there exists an allocation in the core of H that contains (a^*, b^*) , and thus D must be a yes-instance of ACYCLIC PARTITION. 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




Fig. 12. Illustration for the proof of Theorem 23, for constructing the housing market H' from H by subdividing the arc (a^*, b^*) . The figure omits vertices of $N \setminus \{a^*, b^*\}$; arcs between b^* and $N \setminus \{a^*, b^*\}$ are shown in gray.

$\text{OPT}(H') = |N'|$, 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.

We contrast Theorem 23 with the observation that an algorithm that outputs *any* allocation in the core yields an approximation for MAX CORE with ratio $|N|$.

Proposition 24. 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 Appendix A.

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 *Proc. of the 8th ACM Conference on Electronic Commerce (EC '07)*, pages 295–304, 2007.
- [3] D. J. Abraham, K. Cechlárová, D. F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In *Proc. of the 15th International Conference on Algorithms and Computation (ISAAC 2004)*, volume 3341 of *Lecture Notes in Computer Science*, pages 3–15, 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 *Proc. of the 26th AAAI Conference on Artificial Intelligence (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. Baltessová, I. Boletis, C. Bolotinha, G. Bond, G. Böhmig, L. Burnapp, 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. *Math. Oper. Res.*, 2023.
- [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, W. Pettersson, T. Andersson, L. Burnapp, P. Chromy, P. Delgado, P. Dworzak, B. Haase, 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 *Proc. of the 8th International Symposium on Operational Research in Slovenia (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] K.-M. Chow, U. Maggiore, and F. J. Dor. Ethical issues in kidney transplant and donation during COVID-19 pandemic. *Semin. Nephrol.*, 42:151272, 2022.

- [21] Á. Cseh and K. Heeger. The stable marriage problem with ties and restricted edges. *Discrete Optim.*, 36:100571, 2020.
- [22] Á. Cseh and A. Juhos. Pairwise preferences in the stable marriage problem. *ACM Trans. Econ. Comput.*, 9(1), 2021.
- [23] Á. Cseh and D. F. Manlove. Stable marriage and roommates problems with restricted edges: Complexity and approximability. *Discrete Optim.*, 20:62–89, 2016.
- [24] 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.
- [25] J. Drummond and C. Boutilier. Elicitation and approximately stable matching with partial preferences. In *Proc. of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013)*, page 97–105, 2013.
- [26] 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.
- [27] M. Gelain, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Male optimal and unique stable marriages with partially ordered preferences. In C. Guttmann, F. Dignum, and M. Georgeff, editors, *Collaborative Agents - Research and Development*, volume 6066 of *Lecture Notes in Computer Science*, pages 44–55. Springer Berlin Heidelberg, 2011.
- [28] D. Gusfield and R. W. Irving. *The Stable Marriage problem: Structure and Algorithms*. MIT press, 1989.
- [29] 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.
- [30] C.-C. Huang. Circular stable matching and 3-way kidney transplant. *Algorithmica*, 58(1):137–150, 2010.
- [31] R. W. Irving. An efficient algorithm for the “stable roommates” problem. *J. Algorithms*, 6(4):577–595, 1985.
- [32] R. W. Irving, D. F. Manlove, and S. Scott. Strong stability in the Hospitals/Residents problem. In H. Alt and M. Habib, editors, *Proc. of the 20th International Symposium on Theoretical Aspects of Computer Science (STACS 2003)*, volume 2607 of *Lecture Notes in Computer Science*, pages 439–450, 2003.
- [33] P. Jaramillo and V. Manjunath. The difference indifference makes in strategy-proof allocation of objects. *J. Econ. Theory*, 147(5):1913–1946, 2012.
- [34] Y. Kamijo and R. Kawasaki. Dynamics, stability, and foresight in the Shapley–Scarf housing market. *J. Math. Econ.*, 46(2):214–222, 2010.
- [35] R. Kawasaki. Roth–Postlewaite stability and von Neumann–Morgenstern stability. *J. Math. Econ.*, 58:1–6, 2015.
- [36] B. Klaus and F. Klijn. Minimal-access rights in school choice and the deferred acceptance mechanism. *Math. Oper. Res.*, 2023.
- [37] D. E. Knuth. *Mariages stables et leurs relations avec d’autres problèmes combinatoires*. Les Presses de l’Université de Montréal, 1976.
- [38] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 5th edition, 2012.
- [39] A. Kunysz. The strongly stable roommates problem. In *Proc. of the 24th Annual European Symposium on Algorithms (ESA 2016)*, volume 57 of *LIPICs*, pages 60:1–60:15, 2016.
- [40] A. Kunysz. An algorithm for the maximum weight strongly stable matching problem. In *Proc. of the 29th International Symposium on Algorithms and Computation (ISAAC 2018)*, volume 123 of *LIPICs*, pages 42:1–42:13, 2018.
- [41] M. Kurino. House allocation with overlapping generations. *Am. Econ. J. Microecon.*, 6(1):258–289, 2014.
- [42] D. F. Manlove. Stable marriage with ties and unacceptable partners. Technical Report TR-1999-29, University of Glasgow, Department of Computing Science, 1999.

- [43] D. F. Manlove. *Algorithmics of matching under preferences*, volume 2 of *Series on Theoretical Computer Science*. World Scientific, 2013.
- [44] D. F. Manlove, R. W. Irving, K. Iwama, S. Miyazaki, and Y. Morita. Hard variants of stable marriage. *Theor. Comput. Sci.*, 276(1-2):261–279, 2002.
- [45] B. Pittel. On random stable matchings: Cyclic ones with strict preferences and two-sided ones with partially ordered preferences. *Adv. Appl. Math.*, 120:102061, 2020.
- [46] C. G. Plaxton. A simple family of Top Trading Cycles mechanisms for housing markets with indifferences. In *Proc. of the 24th International Conference on Game Theory (ICGT 2013)*, 2013.
- [47] T. Quint and J. Wako. On houseswapping, the strict core, segmentation, and linear programming. *Math. Oper. Res.*, 29(4):861–877, 2004.
- [48] A. E. Roth and A. Postlewaite. Weak versus strong domination in a market with indivisible goods. *J. Math. Econ.*, 4:131–137, 1977.
- [49] A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *Quarterly J. of Econ.*, 119:457–488, 2004.
- [50] A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *J. Econ. Theory*, 125(2):151–188, 2005.
- [51] D. Saban and J. Sethuraman. House allocation with indifferences: A generalization and a unified view. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC '13)*, pages 803–820, 2013.
- [52] L. Shapley and H. Scarf. On cores and indivisibility. *J. Math. Econ.*, 1:23–37, 1974.
- [53] T. Sönmez and T. Switzer. Matching with (Branch-of-Choice) contracts at the United States Military Academy. *Econometrica*, 81:451–488, 2013.
- [54] J. J. M. Tan. Stable matchings and stable partitions. *Int. J. Comput. Math.*, 39(1-2):11–20, 1991.
- [55] J. J. M. Tan and Y.-C. Hsueh. A generalization of the stable matching problem. *Discrete Appl. Math.*, 59(1):87–102, 1995.
- [56] M. U. Ünver. Dynamic kidney exchange. *Rev. Econ. Stud.*, 77(1):372–414, 2010.