# Top-Trading-Cycles Mechanisms with Acceptable Bundles

**Sujoy Sikdar**
Department of Computer Science
Rensselaer Polytechnic Institute
sikdas@rpi.edu

**Sibel Adalı**
Department of Computer Science
Rensselaer Polytechnic Institute
sibel@cs.rpi.edu

**Lirong Xia**
Department of Computer Science
Rensselaer Polytechnic Institute
xial@cs.rpi.edu

**Abstract**

In this paper, we extend the Top-Trading-Cycles (TTC) mechanism to housing markets with acceptable bundles. We propose *Conditionally-Most-Important trees (CMI-trees)* as a new language that generalizes commonly-studied languages such as generalized lexicographic preferences and LP-trees. Our main result is that for any instance of housing markets with acceptable bundles, and any CMI-tree profile, if TTC outputs an acceptable full allocation, then TTC is strict core selecting, non-bossy and indifferent to the order of implementing cycles as they become available.

We propose multi-type housing markets with size constraints as a generalization of multi-type housing markets (Moulin 1995; Sikdar, Adali, and Xia 2017) and the setting of Fujita et al. (2015), and show that our extension of TTC always outputs an acceptable full allocation. Finally, we show that TTC is computationally hard to manipulate in this setting.

## 1 Introduction

Shapley and Scarf (1974)'s *housing markets* are an important model of exchange economies. In a housing market, there are multiple agents, each initially endowed with some indivisible items and preferences over bundles of items. The goal is to design a mechanism without money to re-allocate the items. One prominent example is *kidney exchange* (Roth, Sonmez, and Unver 2004), where each recipient/donor pair is modeled as an agent, who initially owns a kidney (from the donor) and wants a new kidney (for the recipient).

When each agent is initially endowed with a single item and agents' preferences are linear orders over all items, Gale's *Top-Trading-Cycles (TTC)* mechanism (Abdulkadiroğlu and Sönmez 1999) satisfies many desirable properties. For example, it satisfies *core-selection* and *strategy-proofness*, and runs in polynomial time. Core-selection requires that the outcome of the mechanism must be in the *core*, which is the set of allocations where no group of agents has an incentive to deviate by reallocating their initial endowments. Strategy-proofness requires that no agent has incentive to misreport preferences to obtain a better outcome.

However when some agents initially own multiple items, the problem becomes much more challenging as no mechanism satisfies both core-selection and strategy-proofness in such cases as proved by Sönmez (1999). There have been some recent positive developments on core-selecting mechanisms under natural assumptions on agents' preferences with certain constraints on the final allocation. For example, Fujita et al. (2015) extend TTC under the assumption of lexicographic preferences, where each agent is allocated as many items as their initial endowment. Sikdar, Adali, and Xia (2017) extend TTC to multi-type housing markets (Moulin 1995) where each agents' initial endowment and final allocation consists of exactly one item of each type, under the assumption that agents' preferences are represented by lexicographic extensions of CP-nets.

We address the following questions: To what extent, and under what assumption on agents' preferences, can TTC be extended to housing markets with constraints on the final allocation? What properties does the extended TTC satisfy?

**Our Contributions.** We propose a general framework called *housing markets with acceptable bundles (HMAB)*, which is more general than previous works (Fujita et al. 2015; Moulin 1995; Sikdar, Adali, and Xia 2017). An HMAB is denoted by $(\mathcal{M}, \mathcal{D})$, where $\mathcal{M}$ is a housing market where each of $n$ agents may initially own multiple items and have preferences over *bundles*. Each agent $j \leq n$ has a set of acceptable bundles $D_j$, and an *acceptable allocation* is a member of $\mathcal{D} = D_1 \times \cdots \times D_n$.

We propose a new graphical language to represent preferences called *conditionally-most-important (CMI) trees* that extends several popular languages studied in previous work, as we show in Section 4. In a CMI-tree $V$, each node is labeled with an item, and each directed edge points to the next most-important item, conditioned on whether items along the path are allocated to the agent. Given any partial allocation of items, an agent either has a unique most important item that she desires, or is indifferent between all extensions of the partial allocation.

We extend the TTC mechanism to HMABs when agents' preferences are represented by CMI-trees with possibly different structures. At each round of the algorithm, each remaining agent points at her (unique) most important item conditioned on the partial allocation computed so far, and every item points at its initial owner. The algorithm then implements all cycles formed in the current round by assigning to agents involved in a cycle the items they were pointing at.

We show that for any HMAB $(\mathcal{M}, \mathcal{D})$, and any CMI-profile $P$, if $\text{TTC}(P)$ is an acceptable full allocation i.e. $\text{TTC}(P) \in \mathcal{D}$, and every item is either allocated to some agent, or cannot extend any agents' allocation, then (1) $\text{TTC}(P)$ is in the strict core (**Theorem 1**), (2) the cycles can be implemented one by one (instead of implementing all cycles in each step) in any order when they are available (**Theorem 2**), (3) TTC is non-bossy at $P$ (**Theorem 3**).

Here a strict core allocation means that no group has incentive to deviate and exchange their initial endowments to get a better *acceptable* allocation. Non-bossiness means that no agent can misreport her preferences to change the allocation of any other agent without changing her own allocation.

**How useful are these results?** Theorems 1–3 discussed above are quite general and positive. They can be applied to any housing market with any acceptable bundles under the relatively weak assumption of CMI-tree preferences.

First, if we can show that the output of TTC is always an acceptable full allocation, then it immediately implies that TTC is strict core-selecting, insensitive to the order of implementing cycles one by one, and non-bossy. To give an example of such settings, we introduce a new problem, *multi-type housing markets with size constraints*, that combines and generalizes the two settings from previous works (Fujita et al. 2015; Sikdar, Adali, and Xia 2017).

In this new setting, there are multiple types of items, agents may be endowed with multiple items of each type, and agents only accept bundles that contain as many items of each type as their initial endowment. We show that under CMI-profiles that are linear orders over acceptable bundles, *TTC always outputs an acceptable full allocation.* We also show that computing a beneficial misreport under TTC is NP-hard.

Second, even for problems for which the output of TTC is not always an acceptable full allocation, the three theorems still guarantee desirable properties of TTC for "good" instances.

**Related Work and Discussions.** We are not aware of previous work that explicitly formulates acceptable bundles as we do in this paper. The acceptable bundles can be seen as *soft* constraints on the allocation. We do not aim to design a mechanism that always outputs an acceptable full allocation. Instead, we prove that TTC has good theoretical guarantees *when the output is an acceptable full allocation*.

Our work is related to the literature on matching with constraints (see (Kojima 2015) for a recent survey), where there are two sets of agents, and the goal is to find a stable matching of agents from opposite sides where no pair of agents has an incentive to deviate. Constraints may be in the form of agents specifying acceptable matchings, or quotas on how many other agents a given agent can be matched with (Fragiadakis et al. 2016). In this paper, we are interested in the resource allocation problem where preferences are one sided, i.e. agents have preferences over items, and the solution concept is that of core stability.

A natural extension of the standard housing markets setting is that agents own and desire multiple items (Pápai 2007; Todo, Sun, and Yokoo 2014; Sonoda et al. 2014; Fujita et al. 2015; Sun et al. 2015; Sikdar, Adali, and Xia 2017). Konishi, Quint, and Wako (2001) showed that the core may be empty even when agents' preferences are additively separable. However, two lines of work provide core selecting mechanisms under certain restrictions:

(1) When there is a single type and agents may own multiple items, the ATTC mechanism by Fujita et al. (2015) is strict-core selecting under lexicographic preferences and allocates exactly as many items as agents' endowments.

(2) In multi-type housing markets (Moulin 1995), agents are endowed with, and must be allocated bundles containing exactly one item per type. Sun et al. (2015) assume separable lexicographic preferences with common importance orders. Under preferences that are represented by lexicographic extensions of CP-nets (Boutilier et al. 2004) with possibly different importance orders, the MTTC mechanism by Sikdar, Adali, and Xia (2017) is strict core-selecting and non-bossy.

Multi-type housing markets with size constraints combine and generalize these two settings. Therefore, Theorems 1–3 also apply to every instance of each of these problems under CMI-profiles that are linear orders over acceptable bundles, which is a weaker assumption than in previous works as we will show later. Aside from the standard TTC, the MTTC mechanism (Sikdar, Adali, and Xia 2017), and ATTC mechanisms (Fujita et al. 2015) are special cases of our extension of TTC. Theorems 2 and 3 are new to ATTC to the best of our knowledge. The NP-hardness of computing a beneficial misreport provides a computational barrier against agents' strategic behavior under TTC, similar to the idea of using computational complexity to protect elections (Bartholdi, Tovey, and Trick 1989; Conitzer and Walsh 2016).

Our setting is related to the work on housing markets with indifferences (Quint and Wako 2004; Yilmaz 2009; Alcalde-Unzu and Molis 2011; Jaramillo and Manjunath 2012; Aziz and de Keijzer 2012; Plaxton 2013; Saban and Sethuraman 2013) which assume that agents own and accept a single item and their preferences can be any weak order over the items. Therefore, their results do not directly apply to our setting where agents can be allocated multiple items. On the other hand, our results do not directly apply to the setting of these works since not all weak orders can be represented by CMI-trees.

**Domain Restriction on Preferences.** CMI-trees are more general, and impose a weaker restriction on agents' preferences than the assumptions of previous works such as lexicographic extensions of CP-nets (Sikdar, Adali, and Xia 2017), LP-trees (Booth et al. 2010), and GLPs (Monte and Tumennasan 2015). We look at the relationship with other languages in more detail in Section 4.

## 2 Preliminaries

A housing market with acceptable bundles (HMAB), denoted by $\mathcal{M}$, is given by a tuple $(\mathcal{N}, \mathcal{I}, \mathcal{O})$, where $\mathcal{N} = \{1, \ldots, n\}$ is a set of *agents*, $\mathcal{I}$ is a set of indivisible items, and for each $j \le n$, $\mathcal{O}(j) \subseteq \mathcal{I}$ denotes agents' initial endowments that are disjoint. For any $j \le n$, let $D_j \subseteq 2^{\mathcal{I}}$ denote agent $j$'s *acceptable bundles* and $\mathcal{D} = D_1 \times \cdots \times D_n$.

A *partial allocation* $A$ is a function that maps each agent $j$ to a pair of non-overlapping sets of items. For any $j \leq n$, $A(j) = (I_j, F_j)$, where $I_j \cap F_j = \emptyset$. $I_j$ and $F_j$ denote the sets of items assigned and denied to $j$ respectively. When $I_j \cup F_j = \mathcal{I}$ for every $j \leq n$, the allocation is called a *full allocation*. We often use $A = (I, F)$ to refer to partial allocations, where $I = (I_j)_{j \leq n}$ and $F = (F_j)_{j \leq n}$. A (partial) allocation $A'$ defined by pairs $(I'_j, F'_j)$ for every $j \leq n$, is an *extension* of partial allocation $A$, if for every $j \leq n$, $I_j \subseteq I'_j$ and $F_j \subseteq F'_j$. An allocation $A$ is acceptable if $A \in \mathcal{D}$.

Given an agent $j$'s acceptable bundles $D_j$, a partial allocation $(I_j, F_j)$ is *acceptable* to $j$, if it has an extension to an allocation $(I'_j, F'_j)$ such that $I'_j \in D_j$. An item $o$ is *allowable* for agent $j$ given $D_j$ and $(I_j, F_j)$, if $o \notin I_j \cup F_j$ and $(I_j \cup \{o\}, F_j)$ is acceptable. Otherwise $o$ is *denied*. Example 1 illustrates acceptable allocations and allowable items.

A *preference profile* is a collection of agents' preferences. Given an HMAB $(\mathcal{M}, \mathcal{D})$, a mechanism $f$ is a function that maps agents' profile $P$ to a (partial) allocation.

**Axiomatic Properties.** In this work, we are interested in the following desirable axiomatic properties. A mechanism $f$ is (a) *Individually Rational* if for any profile $P$, no agent prefers her initial endowment to her allocation in $f(P)$. (b) *Pareto optimal* if for any profile $P$, there does not exist an allocation $A$ such that (i) every agent weakly prefers her allocation in $A$ to her allocation in $f(P)$, and (ii) some agent strictly prefers her allocation in $A$ to her allocation in $f(P)$. (c) *Non-bossy* if for every profile $P$, no agent can change another agent's allocation by misreporting her preferences, without also changing her own allocation. (d) *Strict Core Selecting* if for every profile $P$, there is no coalition which *weakly blocks* $f(P)$. A coalition of agents $S \subseteq \mathcal{N}$ weakly blocks an allocation $A$, if there is a reallocation $B$ of items initially endowed to agents in $S$, such that (i) every agent in $S$ weakly prefers her allocation in $B$ to her allocation in $A$, and (ii) some agent in $S$ strictly prefers her allocation in $B$ to her allocation in $A$.

## 3 Multi-Type Housing Markets with Size Constraints

In a multi-type housing market with size constraints, there is a set $\mathcal{N}$ of $n$ agents and $p$ types of items. For each type $i \leq p$, $T_i$ denotes items of type $i$, and $\mathcal{I} = T_1 \cup \cdots \cup T_p$. Each agent $j \leq n$ is endowed with a non-empty subset of items $\mathcal{O}(j)$, and for every type $i$, let $[\mathcal{O}(j)]_i \subseteq T_i$ denote agent $j$'s possibly empty endowment of type $i$. Agents only accept bundles containing exactly as many items of each type as their endowments, i.e. $D_j = \{B : \text{for every type } i \leq p, |[B(j)]_i| = |[\mathcal{O}(j)]_i|\}$. Let $\mathcal{D}_{\text{MS}} = D_1 \times \cdots \times D_n$, and $\mathcal{M}_{\text{MS}} = (\mathcal{N}, \mathcal{I}, \mathcal{O})$. Multi-type housing markets with size constraints are a special case of HMABs, $(\mathcal{M}_{\text{MS}}, \mathcal{D}_{\text{MS}})$.

**Example 1.** *Consider the multi-type housing market with size constraints with* 2 *types* $H = \{1_H, 1'_H, 2_H\}$ *and* $C = \{1_C, 2_C, 2'_C\}$*, and* 2 *agents with initial endowments* $\mathcal{O}(1) = \{1_H, 1'_H, 1_C\}$ *and* $\mathcal{O}(2) = \{2_H, 2_C, 2'_C\}$*. Consider the partial allocation* $(I_1, F_1) = (\{1_H, 1_C\}, \{1'_H, 2'_C\})$ *to agent* 1*. The item* $2_H$ *is allowable and* $2_C$ *is denied.* $(I_1, F_1)$ *can be extended by adding* $2_H$ *to* $I_1$ *or* $2_C$ *to* $F_1$ *or both, but every other extension is unacceptable.* □

In a multi-type housing market (Moulin 1995), agents are initially endowed with and only accept bundles with exactly one item per type. In the setting of Fujita et al. (2015), there is a single type, and agents may own multiple items.

## 4 Conditionally-Most-Important (CMI) Preferences

*Conditionally-Most-Important (CMI)* preferences are represented by a tree defined below.

**Definition 1.** *Given* $\mathcal{I}$*, a CMI-tree* $V$ *is defined as follows.*

- *Every node* $d$ *is labeled with an item label*$(d) \in \mathcal{I}$*.*
- *Every item appears at most once on every branch.*
- *Every non-leaf node has either one outgoing edge labeled* $\{0, 1\}$ *or two outgoing edges labeled* 0 *or* 1*.*

*If each node in* $V$ *only has one outgoing edge, then* $V$ *is said to be* unconditional*. For any node* $d$*, let Anc*$(d)$ *denote the set of all ancestor nodes of* $d$*. Let Path*$_V(d)$ *denote the path from root to* $d$*. Let Absent*$(d)$ *(respectively, Present*$(d)$*) denote the set of all items labeling Anc*$(d)$ *with outgoing edges labeled* 0 *(respectively,* 1*) along Path*$_V(d)$*.*

**Semantics.** Given a CMI-tree, each node $d$ has the following meaning: given that the agent is allocated all items in Present$(d)$ and all of the items in Absent$(d)$ are denied in a given partial allocation, any bundle which contains the item label$(d)$ is more preferred than any bundle without label$(d)$. We therefore call item label$(d)$ the agent's *most important item* conditioned on items Present$(d)$ and Absent$(d)$ being present and absent respectively, in the partial allocation.

Comparing a pair of bundles $B_1, B_2 \subseteq \mathcal{I}$ involves traversing the CMI-tree from the root, following outgoing edges labeled 0 or 1, depending on whether the item label$(d)$ labeling the current node is absent from or present in both $B_1$ and $B_2$, until a node $d^*$ is encountered such that label$(d^*)$ is present in only one of the bundles. Such a node $d^*$ is said to *decide* the preference relation in favor of the bundle that includes the item label$(d^*)$. If no decision node is encountered then the agent is indifferent between $B_1$ and $B_2$.

A *CMI-profile* $P = (V_1, \ldots, V_n)$ is a collection of agents' preferences where each $V_j$ is agent $j$'s CMI-tree preferences over $2^{\mathcal{I}}$. A CMI-tree is said to be a strict CMI-tree over a set of bundles $B$ if it induces a linear order over $B$. Given an HMAB $(\mathcal{M}, \mathcal{D})$, a profile $P$ is a strict CMI-profile if for every agent $j$, $V_j$ is a strict CMI-tree over $D_j$. Given a partial allocation $A = (I, F)$, we use $V_j|_{(I_j, F_j)}$ to denote agent $j$'s most important item, conditioned on items in $I_j$ being present, and items in $F_j$ being absent in every bundle.

**Example 2.** Consider the CMI-profile in Figure 1, the multi-type housing market with size constraints from Example 1, and the partial allocation of $(I_2, F_2) = (\{2_H\}, \{1_H, 2_C\})$ to agent 2. Agent 2's conditionally most important item, $V_2|_{(I_2, F_2)} = 1_C$, can be computed by following the edges corresponding to $1_H$ being denied, $2_H$ being assigned, and $2_C$ being denied. Agent 2's preference over acceptable bundles is $[1_H 1_C 2_C \succ 1_H 1_C 2'_C \succ 1_H 2_C 2'_C \succ 2_H 2_C 1_C \succ 2_H 2_C 2'_C \succ 2_H 1_C 2'_C \succ 1'_H 1_C 2_C \succ 1'_H 1_C 2'_C \succ 1'_H 2_C 2'_C]$.

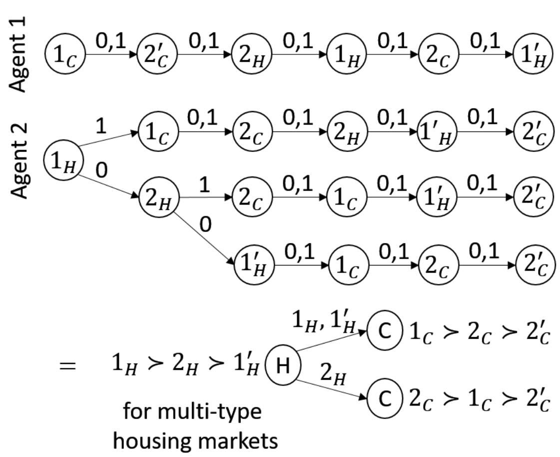


Figure 1: A CMI-profile with two agents and two types $H = \{1_H, 1'_H, 2_H\}$ and $C = \{1_C, 2_C, 2'_C\}$. Agent 2's preferences can be represented by an LP-tree with type labels for multi-type housing markets (Moulin 1995).

Comparing bundles $1_H 1'_H 2_C$ and $1_H 1'_H 2'_C$ according to agent 2's preferences is performed by following edges corresponding to $1_H$ being present, and $1_C$ being absent in both the bundles to reach the node labeled $2_C$. Now, $2_C$ is the most important item and decides the pairwise relationship in favor of $1_H 1'_H 2_C$. □

**CMI-Trees Generalize Other Languages in Multi-Type Housing Markets** (Moulin 1995). We start by showing that CMI-trees are a strict generalization of two important preference languages:

(1) LP-trees with type labels (Booth et al. 2010) which are given by a tree $V$ where each node $v$ is labeled with a type $type(v)$, and a conditional preference table $CPT(v)$, and outgoing directed edges from each non-leaf node are labeled by the items of $type(v)$, such that each item in $type(v)$ labels exactly one outgoing edge of $v$. Each type appears once, and only once on every branch of $V$. $CPT(v)$ is composed of local preferences over items of $type(v)$, conditioned on the allocation of each of the items of the type corresponding to the parent node of $v$ in $V$.

(2) Generalized lexicographic preferences (GLPs) (Monte and Tumennasan 2015) which are represented by a linear order $\eta$ over $\mathcal{I}$. Given any pair $\vec{a}, \vec{b} \in 2^{\mathcal{I}}$, $\vec{a} \succ \vec{b}$ if and only if there is an item $o$ in $\vec{a}$ but not in $\vec{b}$, and all items preferred to $o$ in $\eta$ are either in both $\vec{a}$ and $\vec{b}$, or in neither.

**Example 3.** For multi-type housing markets, agent 2's preferences in Figure 1 can be represented as an LP-tree as shown. Agent 1's preferences in Figure 1 realizes the GLP with $\eta = [1_C \succ 2'_C \succ 2_H \succ 1_H \succ 2_C \succ 1'_H]$. □

**Proposition 1.** *LP-trees with type labels and GLPs are strict subsets of strict CMI-trees.*

**Proof sketch.** It is easy to check that we can model any GLP or LP-tree with type labels as a CMI-tree. We provide examples to prove that LP-trees with type labels and GLPs are a strict subset of CMI-trees. (1) and (2) show that not every GLP can be represented by an LP-tree with type labels and vice versa, and (3) Not every CMI-tree can be represented by either an LP-tree with type labels or a GLP.

(1) Consider the GLP given by $\eta = [1_H \succ 1_C \succ 2_C \succ 2'_C \succ 2_H \succ 1'_H]$. An LP-tree would need to have $H$ as most important type with preferences $[1_H \succ 2_H \succ 1'_H]$ which cannot represent the relation $[1'_H 1_C \succ 2_H 2_C]$ by the GLP.

(2) It is easy to see that no GLP can represent agent 2's LP-tree preferences over $H \times C$ in Figure 1 which induces the order $[1_H 1_C \succ 1_H 2_C \succ \cdots \succ 2_H 2_C \succ 2_H 1_C \succ \dots]$.

(3) By a similar reasoning to the two examples above, we can verify that the CMI-tree in Figure 2 with bundle preferences $[1_H 1_C \succ 1_H 2_C \succ 1_H 2'_C \succ 2_H 2_C \succ 1'_H 2_C \succ 2_H 1_C \succ 1'_H 1_C \succ 2_H 2'_C \succ 1'_H 2'_C]$ over $H \times C$ can neither be represented by an LP-tree with type labels, nor a GLP. □

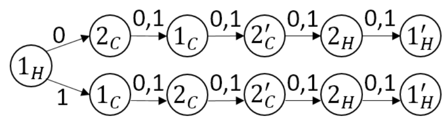

Figure 2: A CMI-tree over 2 types $H$ and $C$.

**Relationship with LP-trees with item labels.**An LP-tree with item labels is represented by a tree, where every node $v$ is labeled by an item $item(v)$, each item is treated as a binary variable, and the local preferences specify whether it is preferred that the item labeling the node is present $(1 \succ 0)$ or absent $(0 \succ 1)$ from a bundle. Since each item appears once, and only once on every branch, this induces a linear order over $2^{\mathcal{I}}$, unlike LP-trees with type labels where the preferences over bundles with multiple items of a type may not be well defined. We say that an LP-tree with item labels is *monotonic* if for every node $v$, $CPT(v) = 1 \succ 0$, i.e. being assigned an item is always preferred. Otherwise, it is *non-monotonic*.

**Proposition 2.** *(1) Strict CMI-trees over $2^{\mathcal{I}}$ are equivalent to monotonic LP-trees with item labels, and (2) CMI-trees cannot represent the preferences of any non-monotonic LP-tree with item labels over $2^{\mathcal{I}}$.*

**Proof sketch.** (1) Constructing a monotonic LP-tree with item labels that represents a strict CMI-tree is as follows. We start with an LP-tree with the same nodes and edges (as well as labels) as the CMI-tree, and populate the local CPTs with $1 \succ 0$. For every leaf node, if the path from root to node does not involve all items, we iteratively add outgoing edges labeled $0, 1$ and nodes labeled by the remaining items as children one by one in arbitrary order. Constructing a CMI-tree to represent a monotonic LP-tree is trivial.

(2) For every non-monotonic LP-tree $V$ with item labels, there must be a pair of bundles $B_1, B_2$ such that there is a single item $o$ which is absent in $B_1$ and present in $B_2$, and every other item is either present or absent in both $B_1$ and $B_2$, and $B_1 \succ B_2$ according to $V$. However, comparing $B_1$ and $B_2$ according to any strict CMI-tree $\hat{V}$ corresponds to a path in $\hat{V}$ involving a node labeled with $o$ which decides the pairwise preference relation as $B_2 \succ B_1$. □

Figure 3 illustrates the relationship between CMI-trees, GLPs, LP-trees with type labels, and LP-trees with item labels.

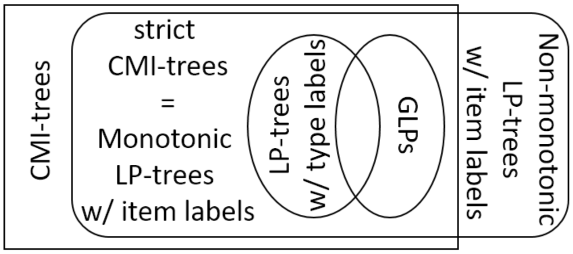


Figure 3: Relationship between popular preference languages (Propositions 1 and 2).

**Relationship with CI-nets** (Bouveret, Endriss, and Lang 2009). In a CI-net, preferences are expressed as a collection of CI-statements where conditioned on some items being present and some other items being absent, one set of items is specified to be more important than another. Given a CI-statement $(S^+, S^-, S_1, S_2)$, comparisons between pairs of bundles where one contains all items in $S_1$, and the other contains all items in $S_2$, and both bundles contain all items in $S^+$ and the same items from $\mathcal{I} \setminus (S_1 \cup S_2)$, and none of the items in $S^-$, are decided in favor of the bundle containing items of $S_1$. This induces a partial order over bundles. Every CMI-tree may be expressed by CI-statements. However such a representation may not be compact. For example, the CMI-tree over items $\{1_H, \dots, m_H\}$ with one node labeled with $1_H$ as the most important item needs an exponential number of CI-statements to represent.

**Properties of CMI-trees.** CMI-trees are compact in the same way that LP-trees are compact. The size of the representation and time taken to compute most important items may depend on the amount of branching in the CMI-tree. Importantly, they can be compact for special cases such as when they represent the same preferences as an LP-tree or GLP. We note that computing the most important item and deciding pairwise comparisons can be performed with polynomial number of queries for the next conditionally most important item. This can be done efficiently for the special cases where CMI-trees are compact. Also, note that TTC only requires knowledge of agents' most important items at each round, instead of requiring agents to formulate and communicate full preferences a priori. Our next proposition is that CMI-trees induce a weak order over bundles, with no incomparabilities.

**Proposition 3.** *Given any CMI-tree $V$ and any $D \subseteq 2^{\mathcal{I}}$, the restriction of $V$ to $D$ is a weak order over $D$.*

**Proof sketch.** The idea is that each decision node in the tree partitions the set of all bundles that are indistinguishable so far into two sets: a more preferred set (following the edge labeled with 1), and a less preferred set (following the edge labeled with 0). Any two bundles that cannot be separated by any node must belong to the same equivalence class. □

Not all linear orders can be represented by CMI-trees. For example, no CMI-tree can represent $1_H \succ 1_H 2_H \succ 2_H$ over $D = \{1_H, 2_H, 1_H 2_H\}$.

## 5 Housing Markets with Acceptable Bundles

Given a housing market with acceptable bundles $(\mathcal{M}, \mathcal{D})$, for any CMI-profile $P = (V_1, \dots, V_n)$, where each $V_j$ is a CMI-tree, we propose an extension of TTC in Algorithm 1.

**Algorithm 1** TTC with output constraints for CMI-profiles.

**Input:** $(\mathcal{M}, \mathcal{D})$ and a CMI-profile $P$.
$t \leftarrow 1$. $\mathcal{N}^t \leftarrow \mathcal{N}, \mathcal{I}^t \leftarrow \mathcal{I}$. Let $G_t = (\mathcal{N}^t \cup \mathcal{I}^t, \emptyset)$. For each $j \leq n$, let $I_j^t = F_j^t = \emptyset$.
**while** at least one agent remains in $\mathcal{N}^t$ **do**
  **Identify most important item** $\hat{o}_j$ in polynomial time. For every agent $j$ in $\mathcal{N}^t$, do the following:
  4.1 Let $o = V_j|_{(I_j^t, F_j^t)}$.
  4.2 If $o$ is denied or removed, add $o$ to $F_j^t$ and update $o = V_j|_{(I_j^t, F_j^t)}$, until $o$ is allowable or no item is left.
  4.3 If an allowable item $o$ is found, then we let $\hat{o}_j = o$, otherwise we let $\hat{o}_j = null$.
  **Build the graph** $G_t = (\mathcal{N}^t \cup \mathcal{I}^t, E)$**.** For every agent $j$ in $\mathcal{N}^t$, do the following:
  5.1 For every item $o \in \mathcal{O}(j) \cap \mathcal{I}^t$, add edge $(o, j)$ to $E$.
  5.2 If $\hat{o}_j \neq null$, add edge $(j, \hat{o}_j)$ to $E$.
  **Implement cycles.** For each cycle $C$, for every $(j, \hat{o}_j)$ in $C$, add $\hat{o}_j$ to $I_j^t$ and remove $\hat{o}_j$ from $\mathcal{I}^t$.
  For all agents $j$ with no outgoing edge in $G_t$, remove $j$ from $\mathcal{N}^t$ and remove items in $\mathcal{O}(j) \cap \mathcal{I}^t$ from $\mathcal{I}^t$.
  $\mathcal{N}^{t+1} \leftarrow \mathcal{N}^t$. $\mathcal{I}^{t+1} \leftarrow \mathcal{I}^t$. For every agent $j \in \mathcal{N}$, $(I_j^{t+1}, F_j^{t+1}) \leftarrow (I_j^t, F_j^t)$. Set $t \leftarrow t + 1$.
**end while**
**Output:** The allocation $A = ((I_j^t, F_j^t)_{j \in \mathcal{N}})$.

At each round of the algorithm, the remaining agents point to their most important allowable item that remains, and each remaining item points to its initial owner. Every node in the corresponding graph has an out-degree of exactly 1, so a cycle is guaranteed to exist. TTC proceeds by implementing all available cycles at each round, by assigning to agents involved in each cycle, the item they were pointing at, updating the partial allocations, and removing the assigned items. At any round $t$ of the algorithm, for any agent $j$, the partial allocation $(I_j^t, F_j^t) = (\text{Present}(d), \text{Absent}(d))$ is used to represent $\text{Path}_{V_j}(d)$, where $\text{label}(d) = V_j|_{(I_j^t, F_j^t)}$ represents the most important item that is under consideration. If an agent does not have any allowable item that remains, she is removed.

**Example 4.** Figure 4 is an example of a run with the CMI-profile in Figure 1 as input for a multi-type housing market. Only $I_j^t$'s are shown. Computation of most important items (darker circle), by following the dashed nodes and edges, are shown for agent 1 at $t_2$ and agent 2 at $t_3$. □

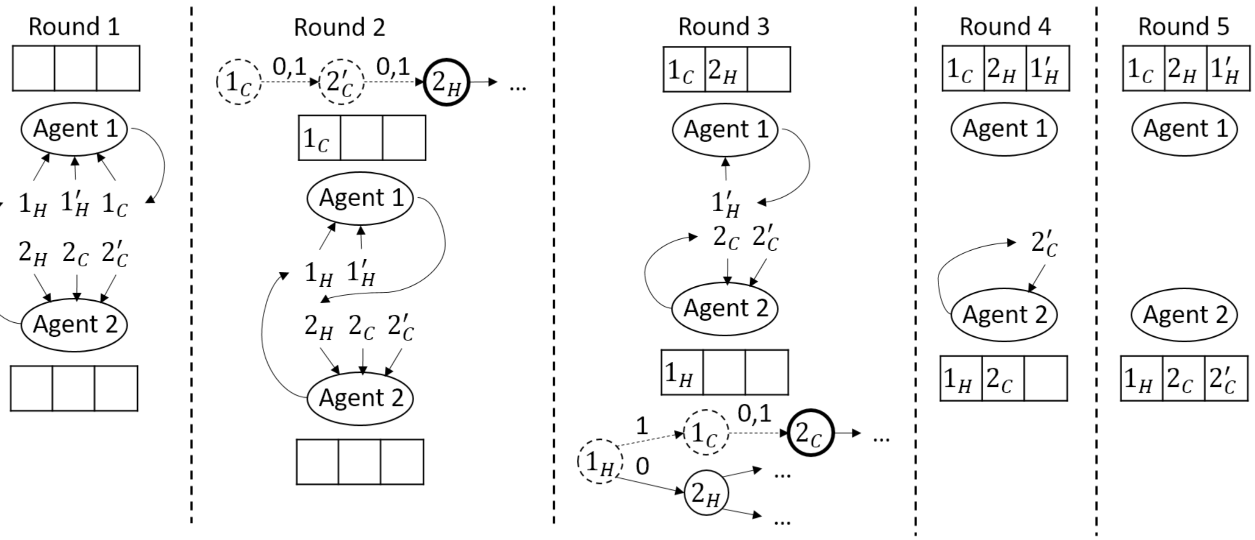


Figure 4: A run of TTC for the multi-type housing market with size constraints from Example 1 and preferences in Figure 1.

**Theorem 1.** *For any HMAB $(\mathcal{M}, \mathcal{D})$ and any CMI-profile $P$, if TTC$(P)$ is an acceptable full allocation then it is in the strict core.*

*Proof.* For the sake of contradiction, suppose that the allocation $A = \text{TTC}(P)$ is an acceptable full allocation and does not belong to the strict core. Then, there exists a coalition $S$ and an acceptable allocation $B = (I^B, F^B)$ on $S$ that blocks $A$, i.e. that every agent in $S$ weakly prefers $B$ to $A$, and at least one agent in $S$ strictly prefers $B$ to $A$. Let $t^*$ be the first round in Algorithm 1 when the partial assignments at the end of the round for agents in $S$ are not compatible with $B$, by which we mean that there exists an agent $j \in S$ such that either $I_j^{t^*} \not\subseteq I_j^B$, or $F_j^{t^*} \cup I_j^B \neq \emptyset$. We can ignore the case where $A$ is a partial allocation because an agent was removed in Step 7. We note that $I_j^{t^*}$ and $F_j^{t^*}$ are updated only at Steps 4 and 6. Therefore, there exists $j \in S$ such that at least one of the following cases hold:

(1) **Denied item in Step 4.** $j$ skips a denied item that is assigned to her in $B$.

(2) **Conflict in $I_j^{t^*}$ in Step 6.** $j$ is assigned an item absent in $B(j)$.

(3) **Removed item in Step 4.** $j$'s item is allocated to a different agent in $A$ than in $B$ in the previous round.

Suppose Case (1) holds and $j$ is denied from getting item $o$ in Step 4.2 of Algorithm 1, but $o \in B(j)$. Immediately before $o$ is added to $F_j^{t^*}$, we have that $(B(j), \mathcal{I} - B(j))$ is an extension of $(I_j^{t^*}, F_j^{t^*})$. Because $o$ is denied, we must have that $B(j)$ is unacceptable, which is a contradiction.

Suppose Case (2) holds and $j$ is assigned $o_A$ in Step 6. Therefore, $o_A$ is $j$'s most important item given $(I_j^{t^*}, F_j^{t^*})$. We note that by assumption, Step 6 in round $t^*$ is the first moment when $(I_j^{t^*}, F_j^{t^*})$ is incompatible with $B(j)$. Therefore, right before $o_A$ is assigned to $j$, $(B(j), \mathcal{I} - B(j))$ is an extension of $(I_j^{t^*}, F_j^{t^*})$. This means that all items in $(B(j) - I_j^{t^*})$ must be allowable at this point. By the definition of CMI-tree, we have $A(j) \succ_j B(j)$, a contradiction.

Case (3) can be reduced to Case (2) by considering the cycle involving $j$ in the previous round. There must be an agent $\hat{j}$ in $S$ who gets a different item in the cycle from $B(\hat{h})$.

This means that since $A$ is a full allocation, it is impossible that an agent in $S$ gets a different allocation in $B$. □

Our next theorem states that if TTC$(P)$ is acceptable, then the order of implementing cycles in the execution of TTC does not matter. We begin by defining a class of algorithms called TTC$^*$ for CMI-profiles similar to (Sikdar, Adali, and Xia 2017), each of which implements one trading cycle per round. There can be multiple TTC$^*$ algorithms.

**Definition 2.** *Given an HMAB $(\mathcal{M}, \mathcal{D})$ and a CMI-profile $P$, let TTC$^*_P$ denote the set of algorithms, each of which is a modification of TTC (Algorithm 1), where instead of implementing all cycles in each round, the algorithm implements exactly one available cycle in each round.*

We note that TTC$^*_P$ depends on $P$. Each algorithm $\mathcal{A} \in \text{TTC}^*_P$ is represented by a linear order $\text{Order}(\mathcal{A}) = C_1 \succ \cdots \succ C_k$, where for any $t \leq k$, $\mathcal{A}$ implements $C_t$ in round $t$. $\text{Order}(\text{TTC}^*_P) = \{\text{Order}(\mathcal{A}) : \mathcal{A} \in \text{TTC}^*_P\}$.

**Definition 3.** *For any HMAB $(\mathcal{M}, \mathcal{D})$ and any CMI-profile $P$, let Cycles$(P)$ denote the set of cycles implemented in the execution of TTC on $P$. We define a partial order PO$(P)$ over Cycles$(P)$ as follows. For every pair of cycles $C_k, C_l$, we let $C_k \succ C_l$ in PO$(P)$ if the following conditions hold:* **Conditions.** *There is an agent $j$ in $C_l$ and an item $o$ in $C_k$ such that $o$ is earlier than $C_l(j)$ in Path$_{V_j}$(TTC$(P)(j)$), and $o$ is allowable given the subpath from root to $o$ in Path$_{V_j}$(TTC$(P)(j)$).*

*PO$(P)$ is the transitive closure of the binary relations mentioned above. We use Ext$(P)$ to denote the set of linear orders that extend PO$(P)$.*

The condition is used to guarantee that when all cycles ahead of $C_l$ in Cycles$(P)$ are removed, the graph by any $\text{TTC}^*_P$ algorithm contains $C_l$. We are now ready to present the key lemma that establishes the equivalence between $\text{TTC}^*_P$ and $Ext(P)$, which immediately leads to Theorem 2.

**Lemma 1.** *For any HMAB $(\mathcal{M}, \mathcal{D})$ and any CMI-profile $P$, if TTC$(P)$ is an acceptable full allocation, then Order$(TTC^*_P) =$ Ext$(P)$.*

**Proof sketch.** The full proof is provided in the appendix.
$\supseteq$: We prove that for every $W \in \text{Ext}(P)$, there exists a $\text{TTC}^*_P$ algorithm $\mathcal{A}$ that implements cycles exactly as in $W$ by contradiction. Suppose that $W = C_1 \succ \cdots \succ C_k$ and that there exists $\mathcal{A} \in \text{TTC}^*_P$ where the first $h-1$ cycles are exactly $C_1, \ldots, C_{h-1}$, but there is no mechanism in $\text{TTC}^*_P$ where the first $h$ cycles are $C_1, \ldots, C_h$. The proof proceeds by first showing that $C_h$ is a cycle in the graph $G_h$ in Step 5 of round $h$. Next, we show that there must exist a mechanism in $\text{TTC}^*_P$ that implements $C_h$ in round $h$ by contradiction.
$\subseteq$: Suppose for the sake of contradiction that there exists $\mathcal{A} \in \text{TTC}^*_P$, such that Order$(\mathcal{A}) \notin \text{Ext}(P)$. W.l.o.g. let Order$(\mathcal{A}) = [C_1 \succ \cdots \succ C_{h-1} \succ C^*_h \succ \cdots]$, and there is some $L \in \text{Ext}(P)$ that agrees with Order$(\mathcal{A})$ with the order of the top $h-1$ elements (cycles), but no $\hat{L} \in \text{Ext}(P)$ that agrees with the order of the top $h$ cycles. W.l.o.g. let $L = [C_1 \succ \cdots \succ C_k]$, where $C_h \neq C^*_h$. By the $\supseteq$ direction, there exists $\mathcal{A}_L \in \text{TTC}^*_P$ such that Order$(\mathcal{A}_L) = L$. The proof relies on establishing that $C^*_h \in G_h$ and $C^*_h \in \{C_{h+1}, \ldots, C_k\}$, because $G_h$ is the graph at the beginning of round $h$ in both $\mathcal{A}$ and $\mathcal{A}_L$. □

**Theorem 2.** *For any HMAB $(\mathcal{M}, \mathcal{D})$ and any CMI-profile $P$, if TTC$(P)$ is an acceptable full allocation, then the output of every $TTC^*_P$ algorithm is the same and equals TTC$(P)$.*

The last theorem in this section proves that TTC is non-bossy when TTC$(P)$ is full and acceptable which follows from Lemma 2 whose proof relies on Theorem 2. The full proofs are provided in the appendix.

**Lemma 2.** *For any HMAB $(\mathcal{M}, \mathcal{D})$, any CMI-profile $P$, and any agent $j$, let $\hat{V}_j$ denote an unconditional CMI-tree obtained from $V_j$ by letting TTC$(P)(j)$ to be the most important items w.r.t. the order they are allocated to $j$. Then if TTC$(P)$ is an acceptable full allocation, we have that TTC$(P) =$ TTC$(P_{-j}, \hat{V}_j)$.*

**Proof sketch.** The proof is done by examining the execution of an algorithm in $\text{TTC}^*_P$. □

**Theorem 3.** *For any HMAB $(\mathcal{M}, \mathcal{D})$, any CMI-profile $P$, if TTC$(P)$ is an acceptable full allocation, then TTC is non-bossy at $P$.*

### TTC$(P)$ is Always Acceptable For Multi-Type Housing Markets with Size Constraints.

**Proposition 4.** *For any strict CMI-profile $P$ and any multi-type housing market with size constraints, $(\mathcal{M}_{MS}, \mathcal{D}_{MS})$, we have that TTC$(P)$ is a full acceptable allocation.*

**Proof sketch.** A proof by induction shows that the partial allocation at the end of every round is acceptable. Either an agent pointed at and was assigned a denied item which is impossible in TTC, or there is no allowable item that meets the agents' demand for some type. The latter case is impossible because agents only point to allowable items, and each agent only demands as many items of each type as they are endowed with. Then, at the end of each round, the number of items of each type equals the demand from remaining agents, so an allowable item is always available. It is easy to see that the output is a full allocation. Example 4 illustrates TTC for multi-type housing markets with size constraints. □

**Manipulation of TTC in Multi-Type Housing Markets with Size Constraints.** We first show by an example, that TTC is not necessarily strategy-proof for multi-type housing markets under CMI-tree preferences. However, as we show subsequently, computing a beneficial misreport is NP-complete, for CMI-trees, GLPs, and LP-trees.

**Example 5.** Consider the multi-type housing market with 4 agents and 2 types $\{H, C\}$ with the following GLP preferences. Agent 1: $3_H \succ 2_C \succ 1_H \succ 1_C \succ$ others. Agent 2: $1_H \succ 2_C \succ 2_H \succ$ others. Agent 3: $1_H \succ 1_C \succ 3_H \succ 3_C \succ$ others. Agent 4: $4_H \succ 4_C \succ$ others. The TTC allocation for the true preferences is: $(3_H, 1_C), (2_H, 2_C), (1_H, 3_C), (4_H, 4_C)$.

However, if agent 1 misreports her preferences as: $4_H \succ 2_C \succ 3_H \succ 1_C \succ$ others, the resulting allocation by TTC is: $(3_H, 2_C), (1_H, 3_C), (2_H, 1_C), (4_H, 4_C)$, which is a beneficial manipulation for agent 1. □

**Definition 4.** *(*TTC-BENEFICIAL-MISREPORT*) Given an HMAB $(\mathcal{M}, \mathcal{D})$, a profile $P$, and an agent $j^*$, we are asked if agent $j^*$ has a misreport $\hat{V}_{j^*}$ so that $TTC(\hat{P} = (\hat{V}_{j^*}, V_{-j^*}))(j^*) \succ_{V_{j^*}} TTC(P)(j^*)$.*

**Theorem 4.** TTC-BENEFICIAL-MISREPORT *in $(\mathcal{M}_{MT}, \mathcal{D}_{MT})$ is NP-complete to compute under GLP, LP-tree, or CMI preferences.*

**Proof sketch.** We show in the appendix a reduction from ATTC-BENEFICIAL-MISREPORT for ATTC (Fujita et al. 2015) which is known to be NP-complete.

□

## 6 Applications and Future Directions

We have extended TTC to HMABs when agents' preferences are represented by CMI-trees. We proved that TTC satisfies some desirable properties when the output is an acceptable full allocation. We further showed that for multi-type housing markets with size constraints which combine the settings of previous works, TTC always computes acceptable allocations for CMI-profiles. Although TTC may not be strategy-proof, computing a beneficial manipulation is NP-hard. Open questions include characterizations of other properties of the extended TTC under other types of preferences. How to choose an acceptable allocation when the output of TTC is unacceptable is also a promising direction for future research in mechanism design.

## 7 Acknowledgements

This work is partly supported by NSF #1453542 and #1716333.

## References


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

Alcalde-Unzu, J., and Molis, E. 2011. Exchange of indivisible goods and indifferences: The Top Trading Absorbing Sets mechanisms. *Games and Economic Behavior* 73:1–16.

Aziz, H., and de Keijzer, B. 2012. Housing Markets with Indifferences: A Tale of Two Mechanisms. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 1249–1255.

Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.

Booth, R.; Chevaleyre, Y.; Lang, J.; Mengin, J.; and Sombattheera, C. 2010. Learning conditionally lexicographic preference relations. In *Proceeding of the 2010 conference on ECAI 2010: 19th European Conference on Artificial Intelligence*, 269–274.

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research* 21:135–191.

Bouveret, S.; Endriss, U.; and Lang, J. 2009. Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In *Proceedings of the 21st international jont conference on Artifical intelligence*, IJCAI'09, 67–72.

Conitzer, V., and Walsh, T. 2016. Barriers to Manipulation in Voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 6.

Fragiadakis, D.; Iwasaki, A.; Troyan, P.; Ueda, S.; and Yokoo, M. 2016. Strategyproof matching with minimum quotas. *ACM Transactions on Economics and Computation* 4(1):6.

Fujita, E.; Lesca, J.; Sonoda, A.; Todo, T.; and Yokoo, M. 2015. A Complexity Approach for Core-Selecting Exchange with Multiple Indivisible Goods under Lexicographic Preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 907–913.

Jaramillo, P., and Manjunath, V. 2012. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory* 147(5):1913–1946.

Kojima, F. 2015. Recent Developments in Matching Theory and its Practical Applications. Working paper.

Konishi, H.; Quint, T.; and Wako, J. 2001. On the Shapley–Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics* 35(1):1—15.

Monte, D., and Tumennasan, N. 2015. Centralized allocation in multiple markets. *Journal of Mathematical Economics* 61:74—85.

Moulin, H. 1995. *Cooperative Microeconomics: A Game-Theoretic Introduction*. Prentice Hall.

Pápai, S. 2007. Exchange in a general market with indivisible goods. *Journal of Economic Theory* 132(1):208—235.

Plaxton, C. G. 2013. A Simple Family of Top Trading Cycles Mechanisms for Housing Markets with Indifferences. In *Proceedings of the 24th International Conference on Game Theory*.

Quint, T., and Wako, J. 2004. On Houseswapping, the Strict Core, Segmentation, and Linear Programming. *Mathematics of Operations Research* 29(4):861–877.

Roth, A. E.; Sonmez, T.; and Unver, M. U. 2004. Kidney exchange. *Quarterly Journal of Economics* 119(2):457–488.

Saban, D., and Sethuraman, J. 2013. House allocation with indifferences: a generalization and a unified view. In *Proceedings of the fourteenth ACM conference on Electronic Commerce*, 803–820.

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

Sikdar, S.; Adali, S.; and Xia, L. 2017. Mechanism Design for Multi-Type Housing Markets. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.

Sönmez, T. 1999. Strategy-proofness and essentially single-valued cores. *Econometrica* 67(3):677–689.

Sonoda, A.; Fujita, E.; Todo, T.; and Yokoo, M. 2014. Two Case Studies for Trading Multiple Indivisible Goods with Indifferences. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 791–797.

Sun, Z.; Hata, H.; Todo, T.; and Yokoo, M. 2015. Exchange of Indivisible Objects with Asymmetry. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, 97–103.

Todo, T.; Sun, H.; and Yokoo, M. 2014. Strategyproof Exchange with Multiple Private Endowments. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 805–811.

Yilmaz, Ö. 2009. Random assignment under weak preferences. *Games and Economic Behavior* 66(1):546–558.