Mechanism Design for Multi-Type Housing Markets

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

We study *multi-type housing markets*, where there are $p \geq 2$ types of items, each agent is initially endowed one item of each type, and the goal is to design mechanisms without monetary transfer to (re)allocate items to the agents based on their preferences over *bundles* of items, such that each agent gets one item of each type. In sharp contrast to classical housing markets, previous studies in multi-type housing markets have been hindered by the lack of natural solution concepts, because the strict core might be empty.

We break the barrier in the literature by leveraging AI techniques and making natural assumptions on agents' preferences. We show that when agents' preferences are lexicographic, even with different importance orders, the classical top-trading-cycles mechanism can be extended while preserving most of its nice properties. We also investigate computational complexity of checking whether an allocation is in the strict core and checking whether the strict core is empty. Our results convey an encouragingly positive message: *it is possible to design good mechanisms for multi-type housing markets under natural assumptions on preferences.*

Introduction

In this paper, we ask the following question: *is it possible at all to design good mechanisms for multi-type housing markets?* In multi-type housing markets (Moulin 1995; Konishi, Quint, and Wako 2001; Wako 2005; Klaus 2008), there are multiple types of items, for example houses and cars, each agent is initially endowed with one item of each type. The goal is to design mechanisms without monetary transfer to (re)allocate items to the agents based on their preferences over *bundles* of items, such that each agent gets one item of each type.

Multi-type housing markets are often described using examples of houses and cars as metaphors for indivisible items. However, the allocation problem is applicable to many other types of items and scarce resources. For example, students may want to exchange papers and dates for presentation (Mackin and Xia 2016); in cloud computing, agents may want to allocate multiple types of resources, including CPU, memory, and storage (Ghodsi et al. 2011; 2012); patients may want to allocate multiple types of medical resources, including surgeons, nurses, rooms, and equipments (Huh, Liu, and Truong 2013).

Mechanism design for single-type housing markets is a well-established field in economics, often referred to as *housing markets* (Shapley and Scarf 1974). In housing markets, the most sensible solution concept is the *strict core*, which is the set of allocations where no group of agents have incentive to deviate by exchanging their initial endowments within the group. Strict core is desirable because it is an intuitive stable solution, and when agents' preferences are linear orders, the strict core allocation always exists and is unique, which can be computed in polynomial time by Gale's celebrated *Top-Trading-Cycles (TTC)* algorithm (Shapley and Scarf 1974; Roth and Postlewaite 1977; Abdulkadiroğlu and Sönmez 1999). TTC enjoys many desirable axiomatic properties including *individual rationality*, *Pareto optimality*, and *group strategy-proofness* (Bird 1984). Many extensions of TTC to other single-type housing markets have been proposed and studied. See more details in Related Work.

In sharp contrast to the popularity of housing markets, there is little research on multi-type housing markets, despite their importance and generality. A potential reason for the absence of positive results is that the strict core can be empty or multi-valued in multi-type housing markets (Konishi, Quint, and Wako 2001). Therefore, as Sönmez and Ünver (2011) noted: *"Positive results of this section [on housing markets] no longer hold in an economy in which one agent can consume multiple houses or multiple types of houses"*. This is the problem we address in this paper and provide a number of positive results, a first in this field.

Our contributions

In this paper, we present novel algorithms building on AI techniques in preference representation and reasoning for allocation in multi-type housing markets. We assume that agents' preferences are represented by arbitrary acyclic *CP-nets* (Boutilier et al. 2004). Different agents may have arbitrarily different CP-net structure. We also assume that agents' preferences are *lexicographic*, meaning that agents have arbitrary importance orders over item types.

We propose the following natural extension of TTC, which we call *Multi-type TTC (MTTC)* for multi-type housing markets. MTTC builds a directed bipartite graph in each round, where agents and items are two separated groups of vertices. There is an edge (a, b_i), where a is an agent and b_i is agent b's initial type-i endowment, if (1) the most important type for agent a, from which she has not obtained an item, is type i, and (2) b_i is agent a's top-ranked type-i item among remaining type-i items. For any type i and any agent b, there is an edge (b_i, b). Then, MTTC finds and implements all cycles as TTC does, but only removes the items in the cycle. The agents always remain in the graph.

We note that in MTTC, a cycle may involve items of multiple types. Our main theorem is the following:

Theorem 1. *For lexicographic preferences, MTTC runs in polynomial-time and satisfies* strict-core-selection *(which implies* Pareto optimality *and* individual rationality*),* non-bossiness*, and* strong group strategy-proofness *when agents cannot lie about importance orders over types.*

We note that Theorem 1's assumption on lexicographic preferences are not only very mild given agents can have arbitrarily different CP-net structure and importance orders, but also an extension of naturalistic decision making structures studied in cognitive science literature for ordering items based on multiple criteria (Gigerenzer and Goldstein 1996). This positive result is especially surprising when it is put in comparison with similar research in *combinatorial voting*, where negative results emerge as soon as agents do not share similar CP-net structures (Lang and Xia 2016).

We also prove an impossibility theorem to show that Theorem 1 cannot be improved by allowing agents to lie about the importance order. As for computing other strict cores, we prove that, it is coNP-complete to check whether a given allocation is in the strict core, even for two types of items and agents with separable preferences w.r.t. the same importance order. This hardness result is unexpected due to relatively small number of bundles (n^2), and the simplicity of the underlying preference structures. Furthermore, the same problem for single-type housing market can be easily computed in polynomial time. We also prove that it is NP-hard to check whether the strict core is non-empty.

We expect that the availability of MTTC will allow development of many applications for multi-type housing markets, especially those discussed in the beginning of the Introduction.

Related Work and Discussions

Housing markets are closely related to *house allocation*, where agents do not have initial endowments, and *matching*, where houses also have preferences over agents. See (Sönmez and Ünver 2011) for a recent survey. In the past decade, housing markets has been a popular topic under *multi-agent resource allocation* (Chevaleyre et al. 2006).

Many subsequent works in Economics and AI extend the standard, single-type housing markets. For example, agents may be indifferent between houses (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). Agents may desire multiple houses (Pápai 2007; Todo, Sun, and Yokoo 2014; Sonoda et al. 2014; Fujita et al. 2015; Sun et al. 2015). Some agents may not have initial endowments (Abdulkadiroğlu and Sönmez 1999; Chen and Sönmez 2002; Sönmez and Ünver 2010).

In our setting, agent's preferences are represented by the celebrated CP-nets (Boutilier et al. 2004). A CP-net allows agents to specify preferences over items within each type given other items allocated to her. In general, a CP-net represents a *partial order* over bundles of items. CP-nets have been heavily used in combinatorial voting, see for example (Rossi, Venable, and Walsh 2004; Lang 2007; Lang and Xia 2016), but negative results often emerge as soon as agents' CP-nets do not line up. *Lexicographic orders* (Booth et al. 2010) are special linear orders that extend some CP-nets, where agents can specify importance orders over types.

Previously, Bouveret, Endriss, and Lang (2009) proposed a CP-net-like language and Fujita et al. (2015) used lexicographic orders to model agents' preferences over bundles of items of a single type. Monte and Tumennasan (2015) characterized strategy-proof mechanisms for multi-type house allocation, when agents' preferences include lexicographic preferences. To the best of our knowledge, our paper is the first time that CP-nets and lexicographic preferences are investigated in multi-type housing markets.

Konishi, Quint, and Wako (2001) assumed that agents have separable and additive preferences. This assumption is in general incomparable to our assumptions on agents' preferences, because preferences represented by CP-nets are generally not separable, except when the CP-net has no edges. Therefore, our positive results are *not* obtained by putting further restrictions on agents' preferences.

There has been very little work on multi-type housing markets after Konishi, Quint, and Wako's negative results (Wako 2005; Klaus 2008). These papers focus on a different solution concept called *coordinate-wise core rule*, which are composed of type-wise strict-core allocations. We show that this naturally corresponds to the output of MTTC when agents' have separable and lexicographic preferences with a common importance order.

Preliminaries

We consider a market consisting of a set $N = \{1, ..., n\}$ of *agents* with $p \geq 2$ *types* of indivisible items. For any $i \leq p$, there are n items of type i, denoted by $T_i = \{i_1, \ldots, i_n\}$. For each item o, Type(o) is the type of o, that is, $o \in T_{\text{Type}(o)}$. Each agent $j \in N$ initially owns exactly one item of each type, and her endowment is denoted by a p-vector $O(j)$. W.l.o.g. in this paper we let $O(j) = (j_1, \ldots, j_p)$. Let $\mathcal{T} = T_1 \times \cdots \times T_p$ be the set of all *bundles*, each of which is represented by a p-vector. We will often use vectors such as $\vec{d}$ and $\vec{e}$ to represent bundles, and for any $i \leq p$, let $[\vec{d}]_i$ denote the type-i item in $\vec{d}$. A *multi-type housing market* M is given by the tuple $(N, \{T_1, \ldots, T_p\}, O)$.

Each agent desires to consume exactly one item of each type, and her preferences are represented by a linear order over $\mathcal{T}$. A *preference profile* $P = (R_1, \ldots, R_n)$ is a collection of agents' preferences. In any multi-type housing marketket M, an allocation A is a mapping from N to $\mathcal{T}$ such that for any $j \leq n$, $A(j)$ is the bundle allocated to j. Since no item is allocated twice, we have that for any $j \neq j'$ and any $i \leq p$, $[A(j)]_i \neq [A(j')]_i$. Given a market M, a *mechanism* f is a function that maps agents' profile P to an allocation in M.

Axiomatic Properties

A mechanism f satisfies *individually rationality* if for any profile P, no agent prefers her initial endowment to her allocation by f. f satisfies *Pareto optimality* if for any profile P, there does not exist an allocation A such that (1) every agent weakly prefers her allocation in A to her allocation in $f(P)$, and (2) some agent strictly prefers her allocation in A to that in $f(P)$. f is *non-bossy* if for any profile P, no agent can change any other agent's allocation without changing her own by reporting differently. f is *strategy-proof* if for each agent, falsely reporting her preferences is not beneficial. A mechanism satisfies *strong group strategy-proofness* if there is no group of agents S who can falsely report their preferences so that (1) every agent in S gets a weakly preferred bundle, and (2) at least one agent in S gets a strictly preferred bundle.

An allocation A is said to be *weakly blocked* by a *coalition* $S \subseteq N$, if the agents in S can find an allocation B of their initial endowments so that each agent weakly prefers allocation in B to that in A, and some agent is strictly better off in B than in A. The *strict core* of a market is the set of all allocations that are not weakly blocked by any coalition. A mechanism f is *strict-core-selecting*, if for any profile P, $f(P)$ is always in the strict core.

CP-nets and Lexicographic Preferences

A (directed) CP-net $\mathcal{N}$ over $\mathcal{T}$ is defined by (i) a directed graph $G = (\{T_1, ..., T_p\}, E)$, called the *dependency graph*, and (ii) for each T_i, there is a conditional preference table CPT_i that contains a linear order $\succ^i_{\vec{u}}$ over T_i for each valuation $\vec{u}$ of the parents of T_i in G, denoted $\text{Pa}(T_i)$. Each CPT-entry $\succ^i_{\vec{u}}$ carries the following meaning: my preferences over type i is $\succ^i_{\vec{u}}$ given that I get items $\vec{u}$, and these preferences are independent of other items I get. An agents' preferences are *separable* if there are no edges in the dependency graph.

Each CP-net $\mathcal{N}$ represents a partial order $\succ_{\mathcal{N}}$, which is the transitive closure of preference relations represented by all CPT entries, which are $\{(a_i, \vec{u}, \vec{z}) \succ_{\mathcal{N}} (b_i, \vec{u}, \vec{z}) : i \leq p; a_i, b_i \in T_i; \vec{u} \in \mathcal{T}_{Pa(T_i)}; \vec{z} \in \mathcal{T}_{-(Pa(T_i) \cup \{T_i\})}\}$.

For example, Figure 1 illustrates a separable CP-net. There are two types: houses (H) and cars (C), with two items per type. Since the preferences are separable, there is no edge in the dependency graph (in the left of the figure). The CPTs are shown in the middle of the Figure 1, and the partial order represented by the CP-net is shown in the right.

Let $\mathcal{O} = [T_1 \rhd \cdots \rhd T_p]$ be a linear order over the types. A CP-net is $\mathcal{O}$-legal, if there is no edge (T_k, T_l) with $k > l$ in its dependency graph. A *lexicographic extension* of an $\mathcal{O}$-legal CP-net $\mathcal{N}$ is a linear order V over $\mathcal{T}$, where for any $i \leq p$, any $\vec{x} \in T_1 \times \cdots \times T_{i-1}$, any $a_i, b_i \in T_i$,

H

Pref. over H
$1_H \succ 2_H$

C

Pref. over C
$1_C \succ 2_C$

$$\begin{tikzcd} (1_H, 1_C) \arrow[r] \arrow[d] & (1_H, 2_C) \arrow[d] \\ (2_H, 1_C) \arrow[r] & (2_H, 2_C) \end{tikzcd}$$

Figure 1: A CP-net.

and any $\vec{y}, \vec{z} \in T_{i+1} \times \cdots \times T_p$, if $a_i \succ^i_{\vec{x}} b_i$ in $\mathcal{N}$, then $(\vec{x}, a_i, \vec{y}) \succ_V (\vec{x}, b_i, \vec{z})$. In other words, the agent believes that type T_1 is most important type to her, T_2 is the second most important type, etc. In a lexicographic extension, $\mathcal{O}$ is called the *importance order*.

In this paper an agent's preferences are *lexicographic*, which means that each agent's ranking is a lexicographic extension of a CP-net. We note that the CP-net does not need to be separable and the importance order can be different.

Example 1. Suppose the agent's preferences is lexicographic w.r.t. the separable CP-net in Figure 1 and the importance order $H \rhd C$, then her preferences are $(1_H, 1_C) \succ (1_H, 2_C) \succ (2_H, 1_C) \succ (2_H, 2_C)$. If her importance order is $C \rhd H$, then her preferences are $(1_H, 1_C) \succ (2_H, 1_C) \succ (1_H, 2_C) \succ (2_H, 2_C)$. $\square$

The Multi-Type TTC Mechanism

We propose the *multi-type TTC (MTTC)* mechanism as Algorithm 1. MTTC assumes that agents' preferences are lexicographic (w.r.t. possibly different importance orders and possibly different CP-net structures).

Example 2. Consider the market with 3 agents and 2 types in Figure 2. The two types are Houses (H) and Cars (C) with items $\{1_H, 2_H, 3_H\}$ and $\{1_C, 2_C, 3_C\}$, respectively. The initial endowments of each agent j is (j_H, j_C). Figure 2 shows agents' lexicographic preferences and the execution of MTTC. In particular, Figure 2 shows the graphs G_1, G_2, G_3 constructed in each round, and all the four cycles implemented in MTTC. The output is the allocation A where $A(1) = (2_H, 1_C), A(2) = (3_H, 2_C), A(3) = (1_H, 3_C)$. $\square$

It is not hard to see that when all agents have the same importance order, the output of MTTC is the same as outcome of the following p-step process. W.l.o.g. let the importance order be $T_1 \rhd T_2 \rhd \cdots \rhd T_p$. For each step i, we ask the agents to report their preferences over T_i, conditioned on items they got in previous steps; then we use TTC to allocate the items in T_i, and move on to the next step. In particular, when all agents' preference are separable, MTTC coincides with the *coordinate-wise core rule* (Konishi, Quint, and Wako 2001).

Theorem 1. *For lexicographic preferences, MTTC runs in polynomial-time and satisfies* strict-core-selection *(which implies* Pareto optimality *and* individual rationality*), non-bossiness, and* strong group strategy-proofneess *when agents cannot lie about importance orders over types.*

Agent 1	$H \rhd C$
H prefs.	$2_H \succ$ others
C prefs.	$3_C \succ 1_C \succ 2_C$

Agent 2	$H \rhd C$
H prefs.	$3_H \succ$ others
C prefs.	$2_C \succ$ others

Agent 3	$H \rhd C$
H prefs.	$1_H \succ$ others
C prefs.	$3_C \succ 1_C \succ 2_C$

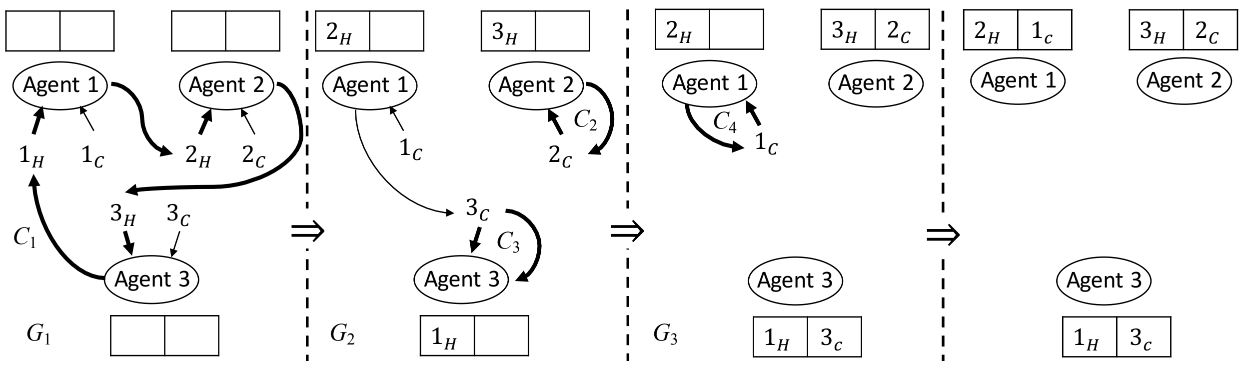

Figure 2: A lexicographic and separable profile P and an execution of MTTC on P.

Algorithm 1 MTTC

Input: A multi-type housing market M and a profile P of lexicographic preferences.
$t \leftarrow 1$. Let $L \leftarrow \cup_{i \leq p} T_i$ be the set of unallocated items. Let A be the empty assignment. For each $j \leq n$, let i_j^* is agent j's most desirable type.
while $L \neq \emptyset$ **do**
  Build a directed graph $G_t = (N \cup L, E)$. For every $j \in N, j_i \in L$, add edge (j_i, j) to E. For every $j \in N$, add edge $(j, \hat{j}_{i_j^*})$ to her most preferred item in L of type i_j^*, to E.
  Implement cycles in G_t**.** For each cycle C, for every $(j, \hat{j}_{i_j^*}) \in C$, assign $[A(j)]_{i_j^*} = \hat{j}_{i_j^*}$.
  Remove assigned items from L.
  For any agent j who is assigned an item, set i_j^* to be the next type according to j's importance order.
  $t \leftarrow t + 1$.
end while
Output: The allocation A.

Proof sketch: Our proof follows the following steps. All omitted proofs can be found in the ArXiv version.

Step 1: MTTC outputs an allocation in polynomial time. This follows after the observation that the number of rounds is no more than np.

Step 2: Strict-core-selecting. Suppose for the sake of contradiction, there exists a coalition S and an allocation B that blocks $A = \text{MTTC}(P)$. Let t^* be the first round in which MTTC's assignments restricted to S at the end of the round differs from the assignments in B. We then prove that there is some $j \in S$ and a type i such that j is assigned a different item of type i in B than in A. Let $j_i^* = [B(j)]_i \neq [A(j)]_i = \hat{j}_i$. Because j's preferences are lexicographic and j get the same items in more important types in B as in A, we have $j_i^* \succ_j \hat{j}_i$. This means that j_i^* is unavailable at round t^*. Therefore, j_i^* must already have been assigned to another agent in a strictly earlier round $t' < t^*$, which is a contradiction to the choice of t^*.

Step 3: MTTC*, the single-cycle-elimination variant of MTTC. MTTC* is similar to MTTC except that in each round a *single* trading cycle is implemented.

Definition 1. *Given a housing market M and any profile P, let MTTC*(P) denote the set of algorithms, each of which is a modification of MTTC (Algorithm 1), where instead of implementing all cycles in each round, the algorithm implements exactly one available cycle in each round.*

We note that MTTC*(P) depends on P because for different profiles the cycles in each round might be different. For each $\mathcal{A} \in$ MTTC*(P), we let Order$(\mathcal{A})$ denote the linear order over the cycles that $\mathcal{A}$ implements. That is, if Order$(\mathcal{A}) = C_1 \succ \cdots \succ C_k$, then it means that for any $t \leq k$, C_t is the cycle implemented by $\mathcal{A}$ in round t.

Example 3. Let M, P be defined as in Figure 2. Let C_1, C_2, C_3, C_4 be the same cycles as in Figure 2. An $\mathcal{A}$ in MTTC* can be defined by following $C_1 \succ C_2 \succ C_3 \succ C_4$ to implement the cycles. □

Definition 2. *For any multi-type housing market M, let Cycles(P) denote the set of cycles implemented in the execution of MTTC on P. We define a partial order PO(P) over Cycles(P) as follows. For every pair of cycles C_k, C_l, $C_k \succ C_l$ in PO(P) if (1) there is an agent who gets an item of a more important type in C_k than in C_l, or (2) there is an agent in C_l who prefers an item in C_k over the item she is pointing to in C_l of the same type, conditioned on the items she got in previous rounds.*

Example 4. Continuing Example 3, PO$(P) = \{C_1 \succ C_2, C_1 \succ C_3, C_1 \succ C_4, C_3 \succ C_4\}$. For all $2 \leq i \leq 4$, $C_1 \succ C_i$ because houses are more important than cars to all agents. $C_3 \succ C_4$ because agent 1 has a more preferred car in C_3 than in C_4. □

We now present two key lemmas.

Agent 1	$H \rhd C$
H prefs.	$2_H \succ 1_H \succ 3_H$
C prefs.	$1_C \succ 3_C \succ 2_C$

Agent 2	$C \rhd H$
H prefs.	$3_H \succ 2_H \succ 1_H$
C prefs.	$2_H \succ$ others

Agent 3	$H \rhd C$
H prefs.	$1_H \succ 3_H \succ 2_H$
C prefs.	$1_C \succ 3_C \succ 2_C$

Figure 3: An example of 3 agents' separable lexicographic preferences over 2 types: Houses (H) and Cars (C).

Lemma 1. *For any multi-type housing market M and any profile P, we have Order(MTTC*(P)) = Ext(P).*

Lemma 2. *For any multi-type housing market M, any profile P, any agent j, and any P' obtained from P by changing agent j's top-ranked bundle to be MTTC(P)(j), then MTTC(P') = MTTC(P).*

Step 4: Non-bossiness. Suppose for the sake of contradiction MTTC does not satisfy non-bossiness. W.l.o.g. let P and P' denote profiles where only agent 1's preferences are different, MTTC$(P)(1)$ = MTTC$(P')(1)$, yet MTTC$(P) \neq$ MTTC(P'). Let $\hat{P}$ denote the market obtained from P by letting agent 1's top-ranked bundle to be MTTC$(P)(1)$. By Lemma 2, MTTC$(\hat{P})$ = MTTC(P) and MTTC$(\hat{P})$ = MTTC(P'), which is a contradiction.

Step 5: Strong group strategy-proofness when the agents cannot lie about the importance orders. Suppose for the sake of contradiction, there exist truthful profile P, $S \subseteq N$, and a beneficial false profile P'. Let $\hat{P}$ denote the profile obtained from P' by letting the top-ranked bundle of all agents $j \in S$ be MTTC$(P')(j)$. By sequentially applying Lemma 2 to agents in S, we have that MTTC$(\hat{P})$ = MTTC(P').

We then compare side by side two parallel runs of two MTTC* algorithms: $\mathcal{A} \in$ MTTC*(P) and $\hat{\mathcal{A}} \in$ MTTC*$(\hat{P})$. We will define $\mathcal{A}$ and $\hat{\mathcal{A}}$ dynamically. Starting with $t = 0$, let G_t and $\hat{G}_t$ denote the graphs of MTTC* at the beginning of round t for input P and input $\hat{P}$, respectively. If there is a common cycle C in G_t and $\hat{G}_t$, then we let both $\mathcal{A}$ and $\hat{\mathcal{A}}$ implement C, and move on to the next round. Let t^* be the earliest round where G_{t^*} and $\hat{G}_{t^*}$ do not have a common cycle. We can prove that there exists $j \in S$ such that $(j, o_i) \in C$ and $(j, s_i) \in \hat{G}_{t^*}$, where $o_i \neq s_i$. Agent j must strictly prefers o_i to s_i, given her allocations in previous rounds. Also agent j must get s_i in MTTC$(\hat{P})$ because this is how we define $\hat{P}$. This contradicts the assumption that none of agents in S is strictly worse off in MTTC$(\hat{P})$. □

Proposition 1. *MTTC is not strategy-proof w.r.t. only misreporting the importance order (i.e. without misreporting local preferences over types).*

Proof. Consider the preferences in Figure 2. We recall that when agents are truthful, the output of MTTC is $((2_H, 1_C), (3_H, 2_C), (1_H, 3_C))$ to agents $1, 2, 3$, respectively (Example 2). If agent 1 misreports the importance order as $C \rhd H$ without misreporting any preferences over types, then the output of MTTC is $((2_H, 3_C), (3_H, 2_C), (1_H, 1_C))$. We note that agent 1 prefers $(2_H, 3_C)$ to $(2_H, 1_C)$. □

Proposition 2. *The strict core of a multi-type housing market can be multi-valued, even when agents' preferences are separable and lexicographic w.r.t. the same order.*

Proof. Consider the preferences in Figure 2. Let B denote the allocation such that $B(1) = (2_H, 3_C), B(2) = (3_H, 2_C), B(3) = (1_H, 1_C)$. For the sake of contradiction, suppose S be a blocking coalition to B. We can observe that agents $1, 2$ each receive their top bundles in B and have no incentive to participate in a coalition. Therefore $S = N$. However, it can be verified that B is Pareto optimal. Further, agent 3 cannot benefit by not participating since $(1_H, 1_C) \succ (3_H, 3_C)$. This means that there is no coalition that blocks B, which is a contradiction. □

Theorem 2. *For any multi-type housing market M with $n \geq 3$ and $p \geq 2$, no mechanism satisfies individually rationality, Pareto optimality, and strategy-proofness, even when agents' preferences are lexicographic and separable.*

Proof. For the sake of contradiction, let f be such a mechanism. Consider the agents' lexicographic and separable preferences P as in Figure 3. Explicitly, P is the following
1: $(2_H, 1_C) \succ (2_H, 3_C) \succ (2_H, 2_C) \succ (\mathbf{1_H}, \mathbf{1_C}) \succ$ others.
2: $(3_H, 2_C) \succ (\mathbf{2_H}, \mathbf{2_C}) \succ$ others.
3: $(1_H, 1_C) \succ (1_H, 3_C) \succ (1_H, 2_C) \succ (3_H, 1_C) \succ (\mathbf{3_H}, \mathbf{3_C}) \succ$ others.

We can verify that the only allocations that satisfy individual rationality and Pareto optimality are $((2_H, 3_C), (3_H, 2_C), (1_H, 1_C))$ and $((2_H, 1_C), (3_H, 2_C), (1_H, 3_C))$. If $f(P) = ((2_H, 3_C), (3_H, 2_C), (1_H, 1_C))$, then we show that agent 1 has incentive misreports her lexicographic order as $2 \rhd 1$, leading to the only individually rational and Pareto optimal allocation $((2_H, 1_C), (3_H, 2_C), (1_H, 3_C))$. If $f(P) = ((2_H, 1_C), (3_H, 2_C), (1_H, 3_C))$, then we show that agent 3 has incentive misreports her lexicographic order as $2 \rhd 1$ and her local preferences over type 1 as $3_H \succ 1_H \succ 2_H$, leading to the only individually rational and Pareto optimal allocation $((2_H, 3_C), (3_H, 2_C), (1_H, 1_C))$. Either case leads to a contradiction with strategy-proofness. □

Computing the Membership of the Strict Core

Definition 3 (InStrictCore)**.** *Given a multi-type housing market M, agents' preferences P, and an allocation A, we are asked whether A is a strict core allocation w.r.t. M.*

Theorem 3. *InStrictCore is co-NPC even when agents have separable lexicographic preferences over $p = 2$ types.*

Proof sketch: Membership in co-NPC is easy to verify. For co-NP hardness, we sketch a reduction from 3-SAT. An instance of 3-SAT is given by a formula F in 3-CNF, consisting of clauses $c_1, ..., c_n$ involving Boolean variables $x_1, ..., x_m$ that can take on values in {0,1}, and we are asked whether there is a valuation of the variables that satisfies F. We will use $F = (x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3)$ to illustrate the proof.

Given an arbitrary instance I of 3-SAT, we construct an instance J of InStrictCore for $p = 2$ as follows. There are four classes of **agents** (1) $\{c_1, \ldots, c_n\}$; (2) for every $i \leq m$ and $j \leq n$, we have agents $1_i^j, 0_i^j$, corresponding to the values that might be used to satisfy the clauses; (3) for every $i \leq m$, we have three agents $x_i, \bar{x}_i, d_i$; (4) two extra agents e_1, e_2. For each agent a, her initial **endowments** are $([a]_1, [a]_2)$. Agents' **preferences** are separable and lexicographic w.r.t. $1 \rhd 2$. Preferences and the **allocation** A are illustrated in Figure 4 (graph G_1) and Figure 5 (graph G_2), which correspond to type 1 and type 2, respectively.

For each $k \in \{1, 2\}$, a *dashed* edge (a, b) in G_k represents (i) the preference $[b]_k \succ [a]_k$, and (ii) the allocation $[A(a)]_k = [b]_k$. A *solid* edge (a, b) indicates a strict preference $[b]_k \succ [A(a)]_k$. The absence of an edge (a, b) indicates $[b]_k \prec [a]_k$.

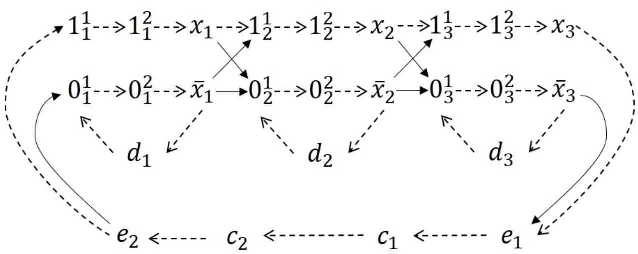

Figure 4: G_1: preferences and allocations for type 1.

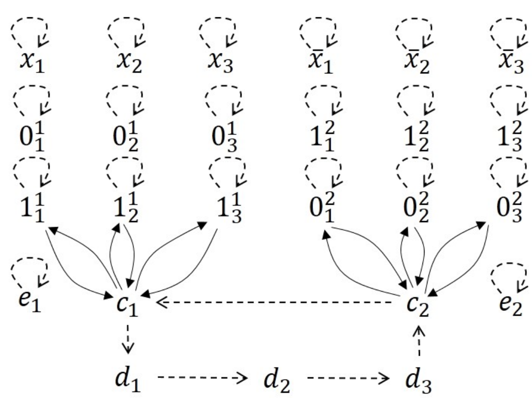

Figure 5: G_2: preferences and allocations for type 2.

G_1 is used to ensure that the valuation of variables is consistent, and G_2 is used to ensure that all clauses are satisfied. If the 3-SAT instance is satisfiable, the blocking coalition S includes all c_i's, e_1, e_2, and the agents correspond to the valuation. The allocation B corresponds to cycles in G_1 and G_2, respectively. On the other hand, if A is not a strict core, we can prove that any blocking coalition corresponds to a satisfying assignment to the 3-SAT instance. □

Definition 4 (StrictCoreNonEmpty). *Given a multi-type housing market M, agents' preferences P, and an allocation A, we are asked whether the strict core of M is non-empty.*

Theorem 4. *StrictCoreNonEmpty is NP-hard in general for even $p = 2$ types.*

Proof sketch: The hardness is proved by a reduction from 3-SAT. Given a 3-SAT instance F, we construct a StrictCoreNonEmpty instance with the following agents: for every clause c_j, there are three agents c_j^1, c_j^2, c_j^3; for every variable x_i there is an agent x_i and two agents $1_i^j, 0_i^j$ for every clause c_j. Agents' preferences will be defined by leverage the delicate example with 3 agents and 2 types, where the strict core is empty (Konishi, Quint, and Wako 2001).

Example 5. [Example 2.2 in (Konishi, Quint, and Wako 2001)] The strict core is empty for the following multi-type housing market with 3 agents and 2 types.
Agent 1: $13 \succ 33 \succ 12 \succ$ others.
Agent 2: $23 \succ 21 \succ 33 \succ 31 \succ 22 \succ$ others.
Agent 3: $21 \succ 22 \succ 31 \succ 11 \succ 32 \succ 12 \succ 23 \succ 33 \succ 13 \succ$ others. □

The preferences of each triple of c_j^1, c_j^2, c_j^3 is similar to Example 5, except that c_j^1's preferences are: $(1_{j_1}^j, c_{j+1}^1) \succ (1_{j_2}^j, c_{j+1}^1) \succ (1_{j_3}^j, c_{j+1}^1) \succ (c_j^1, c_j^3) \succ (c_j^3, c_j^3) \succ (c_j^1, c_j^2) \succ (c_j^1, c_j^1) \succ$ others. For $k \in \{1, 2, 3\}$, if $l_{j_k} = \bar{x}_{j_k}$, then we replace $1_{j_k}^j$ with $0_{j_k}^j$. This modification is to ensure that if the strict core is non-empty, agent c_j^1 must get the type-2 endowment of agent c_{j+1}^1 and the type-1 endowment of a satisfying valuation; otherwise there is no strict core even restrict to agents c_j^1, c_j^2, c_j^3. Based on this we can prove the NP-hardness of StrictCoreNonEmpty. □

Summary and Future Work

We propose MTTC for multi-type housing markets with lexicographic preferences, and prove that it satisfies many desirable axiomatic properties. There are many future directions in mechanism design for multi-type housing markets. Are there good mechanisms when agents demand more than one item of some type? Can we design strategy-proof mechanisms under other assumptions about agents' preferences, such as LP-trees (Booth et al. 2010)? What is the computational complexity of manipulation under MTTC? What if agents' preferences are partial orders such as CP-nets only?

Acknowledgments

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-09-2-0053 (the ARL Network Science CTA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Lirong Xia acknowledges National Science Foundation under grant IIS-1453542 for support.

References

Abdulkadiroğlu, A., and Sönmez, T. 1999. House Allocation with Existing Tenants. *JET* 88:233–260.

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

Aziz, H., and de Keijzer, B. 2012. Housing Markets with Indifferences: A Tale of Two Mechanisms. In *Proc. AAAI*, 1249–1255.

Bird, C. G. 1984. Group incentive compatibility in a market with indivisible goods. *Economics Letters* 14(4):309–313.

Booth, R.; Chevaleyre, Y.; Lang, J.; Mengin, J.; and Sombattheera, C. 2010. Learning conditionally lexicographic preference relations. In *Proc. ECAI*, 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. *JAIR* 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 *Proc. IJCAI*, 67–72.

Chen, Y., and Sönmez, T. 2002. Improving Efficiency of On-Campus Housing: An Experimental Study. *AER* 92(5):1669–1686.

Chevaleyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaitre, M.; Maudet, N.; Padget, J.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in multiagent resource allocation. *Informatica* 30:3–31.

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 *Proc. AAAI*, 907–913.

Ghodsi, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proc. NSDI*, 323–336.

Ghodsi, A.; Sekar, V.; Zaharia, M.; and Stoica, I. 2012. Multi-resource Fair Queueing for Packet Processing. In *Proc. SIGCOMM*, volume 42, 1–12.

Gigerenzer, G., and Goldstein, D. G. 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review* 103(4):650.

Huh, W. T.; Liu, N.; and Truong, V.-A. 2013. Multiresource Allocation Scheduling in Dynamic Environments. *Manufacturing and Service Operations Management* 15(2):280–291.

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

Klaus, B. 2008. The coordinate-wise core for multiple-type housing markets is second-best incentive compatible. *Journal of Mathematical Economics* 44(9–10):919–924.

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.

Lang, J., and Xia, L. 2016. Voting in Combinatorial Domains. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 9.

Lang, J. 2007. Vote and aggregation in combinatorial domains with structured preferences. In *Proc. IJCAI*, 1366–1371.

Mackin, E., and Xia, L. 2016. Allocating Indivisible Items in Categorized Domains. In *Proc. IJCAI*, 359–365.

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. *JET* 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. *MOR* 29(4):861–877.

Rossi, F.; Venable, K. B.; and Walsh, T. 2004. mCP nets: Representing and reasoning with preferences of multiple agents. In *Proc. AAAI*, 729–734.

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

Saban, D., and Sethuraman, J. 2013. House allocation with indifferences: a generalization and a unified view. In *Proc. EC*, 803–820.

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

Sönmez, T., and Ünver, M. U. 2010. House Allocation with Existing Tenants: A Characterization. *GEB* 69(2):425–445.

Sönmez, T., and Ünver, M. U. 2011. Matching, Allocation, and Exchange of Discrete Resources. In Benhabib, J.; Bisin, A.; and Jackson, M. O., eds., *Handbook of Social Economics*. North-Holland. chapter 17, 781–852.

Sonoda, A.; Fujita, E.; Todo, T.; and Yokoo, M. 2014. Two Case Studies for Trading Multiple Indivisible Goods with Indifferences. In *Proc. AAAI*, 791–797.

Sun, Z.; Hata, H.; Todo, T.; and Yokoo, M. 2015. Exchange of Indivisible Objects with Asymmetry. In *Proc. IJCAI*, 97–103.

Todo, T.; Sun, H.; and Yokoo, M. 2014. Strategyproof Exchange with Multiple Private Endowments. In *Proc. AAAI*, 805–811.

Wako, J. 2005. Coalition-proof Nash allocation in a barter game with multiple indivisible goods. *MSS* 49:179—199.

Yilmaz, Ö. 2009. Random assignment under weak preferences. *GEB* 66(1):546–558.