Multi-type Resource Allocation with Partial Preferences*

Haibin Wang[1] , **Sujoy Sikdar**[2] , **Xiaoxi Guo**[1] , **Lirong Xia**[2] , **Yongzhi Cao**[1] , **Hanpin Wang**[1]
[1]Key Laboratory of High Confidence Software Technologies (MOE), School of Electronics Engineering and Computer Science, Peking University, China
[2]Department of Computer Science Rensselaer Polytechnic Institute
beach@pku.edu.cn, sikdas@rpi.edu, guoxiaoxi@pku.edu.cn, xialirong@gmail.com, caoyz@pku.edu.cn, whpxhy@pku.edu.cn

Abstract

We propose multi-type probabilistic serial (MPS) and multi-type random priority (MRP) as extensions of the well known PS and RP mechanisms to the multi-type resource allocation problem (MTRA) with partial preferences. In our setting, there are multiple types of divisible items, and a group of agents who have partial order preferences over bundles consisting of one item of each type. We show that for the unrestricted domain of partial order preferences, no mechanism satisfies both sd-efficiency and sd-envy-freeness. Notwithstanding this impossibility result, our main message is positive: When agents' preferences are represented by acyclic CP-nets, MPS satisfies sd-efficiency, sd-envy-freeness, ordinal fairness, and upper-invariance, while MRP satisfies ex-post-efficiency, sd-strategyproofness, and upper-invariance, recovering the properties of PS and RP.

1 Introduction

Consider the example of rationing [Elster, 1992] two types of *divisible* resources: food (F) and drink (D) among two families who have heterogeneous preferences over combinations of food and drink they wish to consume. For example, a family may prefer water with rice, and milk with wheat. *How should we distribute available resources to the families fairly and efficiently?*

In this paper, we consider the problem of divisible *multi-type resource allocation (MTRA)* [Mackin and Xia, 2016] with partial preferences. Here, there are $p \geq 1$ types of n *divisible* items per type, with one unit of supply of each item, and a group of n agents with *partial preferences* over receiving *bundles* consisting of one unit of each type. Our goal is to design mechanisms to fairly and efficiently allocate one unit of items of each type to every agent given their partial preferences over bundles.

Examples of our setting include the division of land and water resources [Segal-Halevi, 2017]. In cloud computing, agents have preferences over their share of how long they use combinations of computational resources such as CPU, memory, and storage [Ghodsi *et al.*, 2011; 2012; Grandl *et al.*, 2015]. Partial preferences are natural in such problems since the number of bundles grows exponentially with the number of types, and it is often unreasonable to expect agents to form complete preferences over all bundles.

Unfortunately, it is well known that no mechanism which assigns each item fully to a single agent satisfies the basic fairness property of *equal treatment of equals*, meaning that everything else being equal, agents with the same preferences should receive the same share of the resources. For example, whenever two agents have equal and strict preferences over items, it is easy to see that no such mechanism satisfies equal treatment of equals.

Fractional mechanisms overcome this impossibility and possess more favorable properties. Indeed, the random priority (RP) [Abdulkadiroğlu and Sönmez, 1998] and probabilistic serial (PS) [Bogomolnaia and Moulin, 2001] mechanisms are two well-known fractional mechanisms for single-type resource allocations which satisfy equal treatment of equals as well as different notions of fairness and efficiency. RP satisfies ex-post-efficiency, weak-sd-envy-freeness, and sd-strategyproofness, PS is sd-efficient, sd-envy-free, and weak-sd-strategyproof [Bogomolnaia and Moulin, 2001]. PS is the only mechanism that simultaneously satisfies sd-efficiency, sd-envy-freeness, and bounded invariance [Bogomolnaia and Heo, 2012; Bogomolnaia, 2015]. Following in this vein, we focus on the class of *fractional* mechanisms adapted to MTRAs and partial preferences, whose output allocates fractional shares of bundles to the agents.

Katta and Sethuraman [2006] mention that PS can be extended to partial orders but we are not aware of a followup work. Monte and Tumennasan [2015] and Mackin and Xia [2016] consider the problem of MTRA under linear preferences, but do not fully address the issue of fairness. Ghodsi *et al.* [2011] consider the problem of allocating multiple types of resources, when the resources of each type are indistinguishable, and agents have different demands for each type of resource. However, the problem of finding fair and efficient assignments for MTRA with partial preferences remains open.

Our mechanisms output *fractional assignments*, where each agent receives fractional shares of bundles consisting of an item of each type, which together amount to one unit per

*The full version is at: https://gofile.io/?c=qe21wn

Table 1: Properties of MRP and MPS under different domain restrictions on partial preferences. A "Y" indicates that the row mechanism satisfies the column property, and an "N" indicates that it does not. Results annotated with † are from [Bogomolnaia and Moulin, 2001], ‡ are from [Hashimoto *et al.*, 2014]. The rest are results proved in this paper.

Mechanism and Preference Domain		**SE**	**EPE**	**OF**	**SEF**	**WSEF**	**UI**	**SS**	**WSS**
MRP	General partial preferences	N†	Y	N‡	N†	Y	N	N	Y
	CP-nets	N†	Y	N‡	N†	Y	Y	Y	Y
	CP-nets with common dependency graph	N†	Y	N‡	N†	Y	Y	Y	Y
MPS	General partial preferences	Y	N	N	N	Y	N	N†	N
	CP-nets	Y	N	Y	Y	Y	Y	N†	N
	CP-nets with common dependency graph	Y	N	Y	Y	Y	Y	N†	Y

type. In settings such as cloud computing, agents' consumption at any point in time must consist of a bundle composed of every type of item simultaneously at the same rate. The fractional assignments output by our mechanisms also specify for each agent how to form bundles for consumption from the assigned fractional shares of items. Our setting may be interpreted as a special case of *cake cutting* [Procaccia, 2013], where the cake is divided into parts of unit size of p types, and n parts per type, and agents have complex combinatorial preferences over being assigned combinations of parts of the cake which amount to a unit of each type.

Our Contributions. Our work is the first to provide fair and efficient mechanisms for MTRAs, and the first to extend PS and RP both to MTRAs and to partial preferences, to the best of our knowledge. We propose MPS and MRP as extensions of PS and RP to MTRAs respectively. Our main message is positive: Under the well-known and natural domain restriction of CP-net preferences [Boutilier *et al.*, 2004], MRP and MPS satisfy all of the fairness and efficiency properties of their counterparts for single types and complete preferences.

Our technical results are summarized in Table 1. For the unrestricted domain of general partial preferences, unfortunately, no mechanism satisfies both sd-efficiency (SE) and sd-envy-freeness (SEF) as we prove in Proposition 1. Despite this impossibility result, MRP and MPS retain several of the properties of their counterparts RP and PS: We show in Theorem 1 that MRP satisfies ex-post-efficiency (EPE), weak-sd-envy-freeness (WSEF), weak-sd-strategyproofness (WSS), and in Theorem 2 that MPS is sd-efficient (SE) and weak-sd-envy-free (WSEF).

Remarkably, we recover the fairness and efficiency properties of MPS and the truthfulness and invariance properties for MRP and MPS under the well-known and natural domain restriction of CP-net preferences [Boutilier *et al.*, 2004]. We show in:

- Theorem 3, that MRP is sd-strategyproof (SS),

- Theorem 4, that MPS satisfies sd-envy-freeness (SEF) and ordinal fairness (OF), and

- Proposition 3 that MPS is upper-invariant (UI), and in Proposition 4 that MPS satisfies weak-sd-strategyproofness (WSS) under the special case where agents' CP-nets share a common dependency structure.

2 Related Work and Discussion

We are not aware of any previous works which extend RP and PS to MTRAs. MTRAs belong to a long line of research on mechanism design for multi-agent resource allocation (see [Chevaleyre *et al.*, 2006] for a survey), where the literature focuses on the settings with a single type of items. The exchange economy of multi-type housing markets [Moulin, 1995] is considered in [Sikdar *et al.*, 2018; 2017] under lexicographic preferences, while Fujita *et al.* [2015] consider the exchange economy where agents may consume multiple units of a single type of items under lexicographic preferences.

Our work is the first to extend RP and PS under partial preferences, to the best of our knowledge despite the vast literature on fractional assignments. The remarkable properties of PS has encouraged extensions to several settings. Hashimoto *et al.* [2014] provide two characterizations of PS: (1) by sd-efficiency, sd-envy-freeness, and upper-invariance, and (2) by ordinal fairness and non-wastefulness. In [Heo, 2014; Hatfield, 2009] there is a single type of items, and agents have multi-unit demands. In [Saban and Sethuraman, 2014], the supply items may be different, while agents have unit demand and are assumed to have lexicographic preferences. Other works extend RP and PS to settings where indifferences are allowed [Katta and Sethuraman, 2006; Heo and Yılmaz, 2015; Aziz *et al.*, 2015; Hosseini and Larson, 2019]. Yilmaz [2009] and Athanassoglou and Sethuraman [2011] extend PS to the housing markets problem [Shapley and Scarf, 1974]. Bouveret *et al.* [2010] study the complexity of computing fair and efficient allocations under partial preferences represented by SCI-nets for allocation problems with a single type of indivisible items.

3 Preliminaries

A *multi-type resource allocation problem (MTRA)* [Mackin and Xia, 2016], is given by a tuple (N, M, R). Here, (1) $N = \{1, \dots, n\}$ is a set of agents. (2) $M = D_1 \cup \cdots \cup D_p$ is a set of items of p types, where for each $i \le p$, D_i is a set of n items of type i, and there is one unit of *supply* of each item in M. We use $\mathcal{D} = D_1 \times \cdots \times D_p$ to denote the set of *bundles*. (3) $R = (\succ_j)_{j \le n}$ is a *preference profile*, where for each $j \le n$, $\succ_j$ represents the preference of agent j.

Bundles. For any type $i \le p$, we use k_i to refer to the k-th item of type i. Each bundle $\mathbf{x} \in \mathcal{D}$ is a p-tuple, and we use $o \in \mathbf{x}$ to indicate that bundle $\mathbf{x}$ contains item o. We define

$T = \{D_1, \ldots, D_p\}$, and for any $S \subseteq T$, we define $\prod_S = \times_{D \in S} D$, and $-S = T \setminus S$. For any $S \subseteq T$, $\hat{S} \subseteq T \setminus S$, and any $\mathbf{x} \in \prod_S, \mathbf{y} \in \prod_{\hat{S}}$, $(\mathbf{x}, \mathbf{y})$ denotes the bundles consisting of all items in $\mathbf{x}$ and $\mathbf{y}$.

Partial Preferences and Profiles. A partial preference $\succ$ is a partial order over $\mathcal{D}$ which is an irreflexive, anti-symmetric, and transitive binary relation. Given a partial preference $\succ$ over $\mathcal{D}$, we define the corresponding *preference graph* denoted $G_{\succ}$ to be a directed graph whose nodes correspond to the bundles in $\mathcal{D}$, and for every $\mathbf{x}, \mathbf{y} \in \mathcal{D}$, there is an directed edge $(\mathbf{x}, \mathbf{y})$ if and only if $\mathbf{x} \succ \mathbf{y}$ and there exists no $\mathbf{z} \in \mathcal{D}$ such that $\mathbf{x} \succ \mathbf{z}$ and $\mathbf{z} \succ \mathbf{y}$. We use $\mathcal{R}$ to denote the set of all possible preference profiles. We use $\succ_{-j}$ to denote the preferences of agents in $N \setminus \{j\}$. Given a partial order $\succ$ over $\mathcal{D}$, we define the *upper contour set* of $\succ$ at a bundle $\mathbf{x} \in \mathcal{D}$ as $U(\succ, \mathbf{x}) = \{\hat{\mathbf{x}} : \hat{\mathbf{x}} \succ \mathbf{x} \text{ or } \hat{\mathbf{x}} = \mathbf{x}\}$.

Acyclic CP-nets. A CP-net [Boutilier *et al.*, 2004] $\succ$ over the set of variables $\mathcal{D}$ is given by two components (i) a directed graph $G = (T, E)$ called the *dependency* graph, and (ii) for each $i \leq p$, there is a *conditional preference table* $CPT(D_i)$ that contains a linear order $\succ^{\mathbf{x}}$ over D_i for each $\mathbf{x} \in \prod_{Pa(D_i)}$, where $Pa(D_i)$ is the set of types corresponding to the parents of D_i in G. When G is (a)cyclic we say that $\succ$ is a (a)cyclic CP-net. The partial order induced by an acyclic CP-net $\succ$ over $\mathcal{D}$ is the transitive closure of $\{(o, \mathbf{x}, \mathbf{z}) \succ (\hat{o}, \mathbf{x}, \mathbf{z})\} : i \leq p; o, \hat{o} \in D_i; \mathbf{x} \in \prod_{Pa(D_i)}; o \succ^{\mathbf{x}} \hat{o}; \mathbf{z} \in \prod_{-(Pa(D_i) \cup D_i)}$.

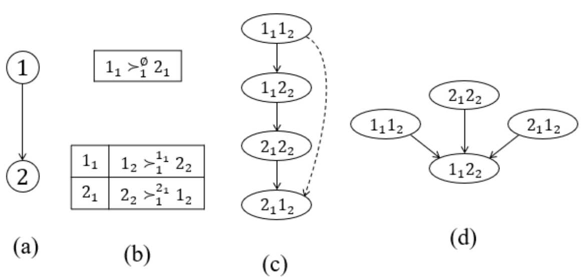

Figure 1: an acyclic CP-net and a general partial preference

Example 1. Consider MTRA (N, M, R) with $p \geq 2$ types where $N = \{1, 2\}, M = \{1_1, 2_1, 1_2, 2_2\}$, where 1_2 is item 1 of type 2 and so on. Let agent 1's preferences $\succ_1$ be represented by the acyclic CP-net in Figure 1, where the dependency graph (Figure 1 (a)) shows that her preference on type 2 depends on her assignment of type 1. The corresponding conditional preference tables (Figure 1 (b)) show that agent 1 prefers 1_2 with 1_1, and she prefers 2_2 with 2_1. This induces the preference graph in Figure 1 (c) which happens to be a linear order. Let agent 2's preferences $\succ_2$ be represented by the preference graph in Figure 1 (d) which represents a partial order, where $1_1 2_2$ is the least preferred bundle. □

Assignments. A *deterministic assignment* $A : N \to \mathcal{D}$ is a one to one mapping from agents to bundles such that no item is assigned to more than one agent. A *fractional allocation* shows the fractional shares an agent acquires over $\mathcal{D}$, represented by a vector $p = [p_{\mathbf{x}}]_{\mathbf{x} \in \mathcal{D}}$, $p \in [0, 1]^{1 \times |\mathcal{D}|}$ such that $\sum_{\mathbf{x} \in \mathcal{D}} p_{\mathbf{x}} = 1$. We use Π to denote the set of all possible fractional allocations on an agent. A *fractional assignment* is a combination of all agents' fractional allocations, represented by a matrix $P = [p_{j,\mathbf{x}}]_{j \leq n, \mathbf{x} \in \mathcal{D}}$, $P \in [0, 1]^{|N| \times |\mathcal{D}|}$, such that (i) for every $j \leq n$, $\sum_{\mathbf{x} \in \mathcal{D}} p_{j,\mathbf{x}} = 1$, (ii) for every $o \in M$, $S_o = \{\mathbf{x} : \mathbf{x} \in \mathcal{D} \text{ and } o \in \mathbf{x}\}$, $\sum_{j \leq n, \mathbf{x} \in S_o} p_{j,\mathbf{x}} = 1$. The j-th row of P represents agent j's fractional allocation under P, denoted $P(j)$. We use $\mathcal{P}$ to denote the set of all possible fractional assignment matrices. If a fractional assignment matrix can be represented as a probability distribution over deterministic assignments, we say the matrix is *realizable*.

Mechanisms. A *mechanism* $f : \mathcal{R} \to \mathcal{P}$ is a mapping from profiles to fractional assignments. For any profile $R \in \mathcal{R}$, we will use $f(R)$ to refer to the fractional assignment matrix output by f.

Stochastic Dominance. We extend the notion of stochastic dominance [Bogomolnaia and Moulin, 2001] to compare fractional assignments for MTRAs under partial preferences.

Definition 1. (stochastic dominance) *Given a partial preference $\succ$ over $\mathcal{D}$, the* stochastic dominance *relation associated with $\succ$, denoted $\succ^{sd}$ is a weak ordering over Π such that for any pair of fractional allocations $p, q \in \Pi$, p* stochastically dominates *q, denoted $p \succ^{sd} q$, if and only if for every $\mathbf{x} \in \mathcal{D}$, $\sum_{\hat{\mathbf{x}} \in U(\succ, \mathbf{x})} p_{\hat{\mathbf{x}}} \geq \sum_{\hat{\mathbf{x}} \in U(\succ, \mathbf{x})} q_{\hat{\mathbf{x}}}$.*

We write $P \succ^{sd}_j Q$ to denote $P(j) \succ^{sd}_j Q(j)$. We use $P \succ^{sd} Q$ to indicate that $P \succ^{sd}_j Q$ for every $j \leq n$ and $P \neq Q$.

Desirable Properties. A fractional assignment P satisfies: (i) **sd-efficiency**, if there is no fractional assignment $Q \neq P$ such that $Q \succ^{sd} P$, (ii) **ex-post-efficiency**, if P can be represented as a probability distribution over *sd-efficient* deterministic assignments, (iii) **sd-envy-freeness**, if for every pair of agents $j, \hat{j} \leq n$, $P(j) \succ^{sd}_j P(\hat{j})$, (iv) **weak-sd-envy-freeness**, if for every pair of agents $j, \hat{j} \leq n$, $P(\hat{j}) \succ^{sd}_j P(j) \implies P(j) = P(\hat{j})$, and (v) **ordinal fairness**, if for every bundle $\mathbf{x} \in \mathcal{D}$ and every pair of agents $j, \hat{j} \leq n$ with $P_{j,\mathbf{x}} > 0$, $\sum_{\hat{\mathbf{x}} \in U(\succ_j, \mathbf{x})} P_{j,\hat{\mathbf{x}}} \leq \sum_{\hat{\mathbf{x}} \in U(\succ_{\hat{j}}, \mathbf{x})} P_{\hat{j},\hat{\mathbf{x}}}$

A mechanism f satisfies $X \in$ {sd-efficiency, ex-post-efficiency, sd-envy-freeness, weak-sd-envy-freeness, ordinal fairness}, if for every $R \in \mathcal{R}$, $f(R)$ satisfies X. A mechanism f satisfies: (i) **sd-strategyproofness** if for every profile $R \in \mathcal{R}$, every agent $j \leq n$, every $R' \in \mathcal{R}$ such that $R' = (\succ'_j, \succ_{-j})$, it holds that $f(R) \succ^{sd}_j f(R')$, and (ii) **weak-sd-strategyproofness** if for every profile $R \in \mathcal{R}$, every agent $j \leq n$, every $R' \in \mathcal{R}$ such that $R' = (\succ'_j, \succ_{-j})$, it holds that $f(R') \succ^{sd}_j f(R) \implies f(R')(j) = f(R)(j)$.

Given any partial preferences $\succ$, we denote $\succ|_{\mathcal{B}}$ by the restriction of $\succ$ to $\mathcal{B} \subseteq \mathcal{D}$, i.e., $\succ|_{\mathcal{B}}$ is a preference relation over $\mathcal{B}$ such that for all $\mathbf{x}, \mathbf{y} \in \mathcal{B}$, $\mathbf{x} \succ|_{\mathcal{B}} \mathbf{y} \Leftrightarrow \mathbf{x} \succ \mathbf{y}$. Then for any $j \leq n$, $\succ'_j$ is an **upper invariant transformation** of $\succ_j$ at $\mathbf{x} \in \mathcal{D}$ under a fractional assignment P if for some $\mathcal{Z} \subseteq \{\mathbf{y} \in \mathcal{D} \mid P_{j,\mathbf{y}} = 0\}$, $U(\succ'_j, \mathbf{x}) = U(\succ_j, \mathbf{x}) \setminus \mathcal{Z}$ and $\succ'_j|_{U(\succ'_j, \mathbf{x})} = \succ_j|_{U(\succ'_j, \mathbf{x})}$. A mechanism f satisfies **upper**

invariance if it holds that $f(R)_{\hat{j},\mathbf{x}} = f(R')_{\hat{j},\mathbf{x}}$ for every $\hat{j} \le n$, $j \le n$, $R \in \mathcal{R}$, $R' \in \mathcal{R}$, and $\mathbf{x} \in \mathcal{D}$, such that $R' = (\succ'_j, \succ_{-j})$ and $\succ'_j$ is an upper invariant transformation of $\succ_j$ at $\mathbf{x}$ under $f(R)$.

4 Extensions of RP and PS Mechanisms to MTRAs with Partial Preferences

We propose MRP (Algorithm 1) and MPS (Algorithm 2) as extensions of the RP [Abdulkadiroğlu and Sönmez, 1998] and PS [Bogomolnaia and Moulin, 2001] mechanisms to MTRAs with partial preferences. Given an instance of MTRA with agents' partial preferences, MRP and MPS operate on a modified preference profile of strict preferences, where for every agent with partial preferences $\succ$, an arbitrary deterministic topological sorting is applied to obtain a strict ordering $\succ'$ over $\mathcal{D}$, such that for any pair of bundles $\mathbf{x}, \mathbf{y} \in \mathcal{D}$, $\mathbf{x} \succ \mathbf{y} \implies \mathbf{x} \succ' \mathbf{y}$. Given a strict order $\succ'$ obtained in this way, and remaining M', we use $Ext(\succ', M')$ to denote the first available bundle in $\succ'$, which we refer to as the agents' favorite bundle.

Algorithm 1 MRP

Input: An MTRA (N, M, R)
Output: Assignment matrix P

For each $j \le n$, compute a linear ordering $\succ'_j$ corresponding to a deterministic topological sort of $G_{\succ_j}$.
$P \leftarrow \mathbf{0}$ and $M' \leftarrow M$.
Pick a random priority order $\rhd$ over agents.
Successively pick a highest priority agent j^* according to $\rhd$. $\mathbf{x}^* \leftarrow Ext(\succ'_j, M')$ and set $P_{j^*,\mathbf{x}^*} \leftarrow 1$. Remove j^*, and remove all items contained by $\mathbf{x}^*$ in M'.
return P

Given an instance of MTRA with agents' partial preferences, MRP fixes an arbitrary deterministic topological sorting $\succ'$ of agents' preferences, and sorts the agents uniformly at random. Then agents get one unit of their favorite available bundle $Ext(\succ', M')$ M' turn by turn.

Given an instance of MTRA with agents' partial preferences, MPS involves applying the PS mechanism to a modified profile $\succ'$ over $\mathcal{D}$ using an arbitrary deterministic topological sorting in multiple rounds as follows. In each round, each agent consumes their favorite *available* bundle by consuming each item in the bundle at a uniform rate of one unit of an item per type per unit of time, until one of the bundles being consumed becomes unavailable because the supply for one of the items in it is exhausted. Although MPS (or MRP) always form the same topological sorting given the same partial preferences, the output of MPS (or MRP) may be different for different topological sortings as Example **??** shows.

5 Properties of MRP and MPS under General Partial Preferences

Theorem 1. *Given any partial preference profile R, MRP(R) is ex-post-efficient, weak-sd-envy-free, and weak-sd-strategyproof.*

Algorithm 2 MPS

Input: An MTRA (N, M, R)
Output: Assignment matrix P

For each $j \le n$, compute a linear ordering $\succ'_j$ corresponding to a deterministic topological sort of $G_{\succ_j}$.
$P \leftarrow \mathbf{0}$ and $M' \leftarrow M$. For every $o \in M$, supply$(o) \leftarrow 1$, $B \leftarrow \emptyset$, progress $\leftarrow 0$.
while $M' \neq \emptyset$ **do**
  $top(j) \leftarrow Ext(\succ'_j, M')$ for every agent $j \le n$.
  Consume.
    5.1: For each $o \in M'$, consumers$(o) \leftarrow |\{j \in N : o \text{ is in } top(j)\}|$.
    5.2: progress $\leftarrow \min_{o \in M'} \frac{\text{supply}(o)}{\text{consumers}(o)}$.
    5.3: For each $j \le n$, $P_{j,top(j)} \leftarrow P_{j,top(j)} + \text{progress}$.
    5.4: For each $o \in M'$, supply$(o) \leftarrow$ supply$(o) -$ progress $\times$ consumers(o).
  $B \leftarrow \arg\min_{o \in M'} \frac{\text{supply}(o)}{\text{consumers}(o)}$, $M' \leftarrow M' \setminus B$
return P

Remark 1. *MRP is not upper invariant and sd-strategyproof under general partial preferences.*

Theorem 2. *Given any partial preference profile R, MPS(R) is sd-efficient and weak-sd-envy-free.*

Remark 2. *MPS is not ex-post-efficient since its output may not be realizable when coming to multi-type resources. MPS is not ordinally fair, sd-envy-free and upper invariant under general partial preferences.*

Proposition 1. *No mechanism can satisfy both sd-efficiency and sd-envy-freeness under general partial preferences.*

6 Properties of MRP and MPS under Acyclic CP-net Preferences

For convenience, we refer to a profile of acyclic CP-net preferences as a *CP-profile*.

Theorem 3. *Given any CP-profile R, MRP(R) is sd-strategyproof.*

Proposition 2. *Given any CP-profile R, MRP(R) is upper invariant for any other CP-profile R'.*

Theorem 4. *Given any CP-profile R, MPS(R) is sd-envy-free and ordinally fair.*

Proposition 3. *Given any CP-profile R, MPS(R) is upper invariant for any other CP-profile R'.*

Remark 3. *MPS is not weak-sd-strategyproof under CP-net preferences.*

Proposition 4. *Given any CP-profile R with a identical dependency graph, MPS(R) is weak-sd-strategyproof for any other CP-profile R' with the same dependency graph.*

References

[Abdulkadiroğlu and Sönmez, 1998] Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–702, 1998.

[Athanassoglou and Sethuraman, 2011] Stergios Athanassoglou and Jay Sethuraman. House allocation with fractional endowments. *International Journal of Game Theory*, 40(3):481–513, 2011.

[Aziz *et al.*, 2015] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.

[Bogomolnaia and Heo, 2012] Anna Bogomolnaia and Eun Jeong Heo. Probabilistic assignment of objects: Characterizing the serial rule. *Journal of Economic Theory*, 147(5):2072–2082, 2012.

[Bogomolnaia and Moulin, 2001] Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[Bogomolnaia, 2015] A. Bogomolnaia. Random assignment: Redefining the serial rule. *Journal of Economic Theory*, 158:308–318, 2015.

[Boutilier *et al.*, 2004] Craig Boutilier, Ronen Brafman, Carmel Domshlak, Holger Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[Bouveret *et al.*, 2010] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*, 2010.

[Chevaleyre *et al.*, 2006] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[Elster, 1992] Jon Elster. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation, 1992.

[Fujita *et al.*, 2015] Etsushi Fujita, Julien Lesca, Akihisa Sonoda, Taiki Todo, and Makoto Yokoo. 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*, pages 907–913, 2015.

[Ghodsi *et al.*, 2011] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, pages 323–336, Boston, MA, USA, 2011.

[Ghodsi *et al.*, 2012] Ali Ghodsi, Vyas Sekar, Matei Zaharia, and Ion Stoica. Multi-resource fair queueing for packet processing. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, volume 42, pages 1–12, Helsinki, Finland, 2012.

[Grandl *et al.*, 2015] Robert Grandl, Ganesh Ananthanarayanan, Srikanth Kandula, Sriram Rao, and Aditya Akella. Multi-resource packing for cluster schedulers. *ACM SIGCOMM Computer Communication Review*, 44(4):455–466, 2015.

[Hashimoto *et al.*, 2014] Tadashi Hashimoto, Daisuke Hirata, Onur Kesten, Morimitsu Kurino, and M Utku Ünver. Two axiomatic approaches to the probabilistic serial mechanism. *Theoretical Economics*, 9(1):253–277, 2014.

[Hatfield, 2009] John William Hatfield. Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, 33(3):505–515, 2009.

[Heo and Yılmaz, 2015] Eun Jeong Heo and Özgür Yılmaz. A characterization of the extended serial correspondence. *Journal of Mathematical Economics*, 59:102–110, 2015.

[Heo, 2014] Eun Jeong Heo. Probabilistic assignment problem with multi-unit demands: A generalization of the serial rule and its characterization. *Journal of Mathematical Economics*, 54:40–47, 2014.

[Hosseini and Larson, 2019] Hadi Hosseini and Kate Larson. Multiple assignment problems under lexicographic preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-19)*, 2019.

[Katta and Sethuraman, 2006] Akshay-Kumar Katta and Jay Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic theory*, 131(1):231–250, 2006.

[Mackin and Xia, 2016] Erika Mackin and Lirong Xia. Allocating indivisible items in categorized domains. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 359–365, 2016.

[Monte and Tumennasan, 2015] Daniel Monte and Norovsambuu Tumennasan. Centralized allocation in multiple markets. *Journal of Mathematical Economics*, 61:74—85, 2015.

[Moulin, 1995] Hervé Moulin. *Cooperative Microeconomics: A Game-Theoretic Introduction*. Prentice Hall, 1995.

[Procaccia, 2013] Ariel D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[Saban and Sethuraman, 2014] Daniela Saban and Jay Sethuraman. A note on object allocation under lexicographic preferences. *Journal of Mathematical Economics*, 50:283–289, 2014.

[Segal-Halevi, 2017] Erel Segal-Halevi. *Fair Division of Land*. PhD thesis, 2017.

[Shapley and Scarf, 1974] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.

[Sikdar *et al.*, 2017] Sujoy Sikdar, Sibel Adali, and Lirong Xia. Mechanism design for multi-type housing markets. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017.

[Sikdar *et al.*, 2018] Sujoy Sikdar, S Adali, and Lirong Xia. Top-trading-cycles mechanisms with acceptable bundles. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI*, volume 18, 2018.

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