Multi-Type Resource Allocation with Partial Preferences

Haibin Wang,[1] Sujoy Sikdar,[2] Xiaoxi Guo,[1] Lirong Xia,[3] Yongzhi Cao,[1*] Hanpin Wang[4,1]
[1]Key Laboratory of High Confidence Software Technologies (MOE),
Department of Computer Science and Technology, Peking University, China
[2]Computer Science & Engineering, Washington University in St. Louis
[3]Department of Computer Science, Rensselaer Polytechnic Institute
[4]School of Computer Science and Cyber Engineering, Guangzhou University, China
beach@pku.edu.cn, sujoy@wustl.edu, guoxiaoxi@pku.edu.cn, xialirong@gmail.com,
caoyz@pku.edu.cn, whpxhy@pku.edu.cn, *Corresponding Author

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 problems (MTRAs) 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. Besides, we propose a hybrid mechanism, multi-type general dictatorship (MGD), combining the ideas of MPS and MRP, which satisfies sd-efficiency, equal treatment of equals and decomposability under the unrestricted domain of partial order preferences.

Introduction

Consider the example of rationing (Elster 1992) two types of *divisible* resources: food (F) and beverage (B) among two families who have heterogeneous preferences over combinations of food and beverage 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 problems of divisible *multi-type resource allocation problems (MTRAs)* (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. Such mechanisms are called *fractional mechanisms*, because an agent may receive fractions of items.

When there is one type ($p = 1$), fractional mechanisms broadly fall in two classes. The random priority (RP) mechanism (Abdulkadiroğlu and Sönmez 1998) exemplifies the first class: agents consume their favorite remaining item one-by-one according to an ordering drawn from the uniform distribution. Each agent is allocated a fraction of each item equal to the probability that they consume the item. It is easy to see that RP satisfies ex-post-efficiency and equal treatment of equals. Additionally, RP satisfies notions of envy-freeness and strategyproofness through the idea of *stochastic dominance (sd)*. Given strict preferences, a fractional allocation p dominates another q, if at every item o, the total share of o and items preferred to o under p is at least the total share under q. RP satisfies weak-sd-envy-freeness and sd-strategyproofness (Bogomolnaia and Moulin 2001).

The probabilistic serial (PS) mechanism belongs to the class of "simultaneous eating" mechanisms (Bogomolnaia and Moulin 2001). PS proceeds in multiple rounds. In each round, all agents simultaneously "eat" their favorite remaining item at a constant, uniform rate, until one of the items being consumed is exhausted. This terminates when all items are fully consumed, and the output allocates each agent with a fraction of each item that they would consume by this procedure. PS satisfies sd-efficiency, sd-envy-freeness, and weak-sd-strategyproofness (Bogomolnaia and Moulin 2001). Besides, PS is the only mechanism that simultaneously satisfies sd-efficiency, sd-envy-freeness, and bounded invariance (Bogomolnaia and Heo 2012; Bogomolnaia 2015).

Our work is the first to consider the design of fair and efficient mechanisms for MTRAs with partial preferences, and the first to extend fractional mechanisms to MTRAs with partial preferences, to the best of our knowledge. Katta and Sethuraman (2006) mention that PS can be extended to partial orders but we are not aware of a (formal or informal) work that explicitly defines such an extension and studies its properties. Monte and Tumennasan (2015) and Mackin and Xia (2016) consider the problem of MTRAs 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 resources. However, the problem of finding fair and efficient assignments for MTRAs with partial preferences remains open.

Our mechanisms output *fractional assignments*, where each agent receives a fractional share of bundles consisting

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Table 1: Properties of MRP, MPS and MGD 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). Other results are proved in this paper.

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

of an item of each type, which together amount to one unit per type. 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* (Brams and Taylor 1996; Brams, Jones, and Klamler 2006; 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 one 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 multi-type probabilistic serial (MPS) and multi-type random priority (MRP) as the 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. 2004a), MRP and MPS satisfy all of the fairness and efficiency properties of their counterparts for single types and complete preferences.

Unlike single-type resources allocations, in MTRAs, not all fractional assignments are decomposable, where assignments can be represented as a probability distribution over assignments where each agent receives a bundle consisting of whole items. Unfortunately, the output of MPS may be indecomposable. In response to this, we propose a new mechanism, multi-type general dictatorship (MGD), which is decomposable and matches the efficiency of MPS, and satisfies equal treatment of equals.

Our technical results are summarized in Table 1. We extend *stochastic dominance* to compare two fractional allocations under partial preferences. Here, a fractional allocation p is said to stochastically dominate another allocation q w.r.t. an agent's partial preference, if at any bundle, the fractional share of weakly dominating bundles in p is larger than or equal to the fractional shares of the bundles in q according to her preference. Formal definitions of stochastic dominance and properties in Table 1 can be found in Preliminaries.

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 3. Despite this impossibility result, MRP, MPS and MGD satisfy several desirable properties: We show in

- Theorem 1 that MRP satisfies ex-post-efficiency (EPE), weak-sd-envy-freeness (WSEF), equal treatment of equals (ETE), weak-sd-strategyproofness (WSS), and decomposibility (DC);
- Theorem 2 that MPS satisfies sd-efficiency (SE), equal treatment of equals (ETE), weak-sd-envy-freeness (WSEF);
- Theorem 3 that MGD satisfies sd-efficiency (SE), ex-post-efficiency (EPE), equal treatment of equals (ETE), and decomposibility (DC).

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

- Theorem 4, that MRP is sd-strategyproof (SS);
- Theorem 5, that MPS is sd-envy-free (SEF) and ordinally fair (OF);
- Proposition 5 that MPS is upper invariant (UI); and
- Proposition 6 that MPS is weak-sd-strategyproof (WSS) under the special case where agents' CP-nets share a common dependency structure.

Discussions. MGD may be viewed as a hybrid between MRP and MPS: on the one hand as a modification of MRP where the priorities depend on the preference profile, instead of being drawn from the uniform distribution; and on the other hand, as a modification of MPS where only a subset of the agents are allowed to eat simultaneously in each round. The design of MGD is motivated by the natural desire for decomposability (not satisfied by MPS), while maintaining the stronger efficiency notion of sd-efficiency (not satisfied by MRP), and the basic fairness property of equal treatment of equals. Decomposable assignments are desired when sharing of items imposes overhead costs, as noted by (Sandomirskiy and Segal-Halevi 2019). However, (i) for MTRAs, fractional assignments are not guaranteed to be decomposable as we show in assignment (3) of Example 2, and indeed, MPS does not guarantee decomposable assignments; (ii) fractional mechanisms are inevitable when certain fairness guarantees are desired. Even when there is

one type ($p = 1$), no mechanism which assigns each item fully to a single agent, can satisfy the basic fairness property of equal treatment of equals, whereby, everything else being equal, agents with the same preferences should receive the same share of the resources (e.g. two agents having identical strict preferences). On the flip side, MGD does not satisfy weak-sd-envy-freeness, while MRP and MPS both do.

Related Work

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. Mackin and Xia (2016) characterize serial dictatorships for MTRAs by strategyproofness, neutrality, and non-bossiness. The exchange economy of multi-type housing markets (Moulin 1995) is considered in (Sikdar, Adali, and Xia 2017; 2018) 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. Hosseini and Larson (2019) consider RP under lexicographic preferences. 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 of items may be different, while agents have unit demand and are assumed to have lexicographic preferences. Other works extend PS to settings where indifference relationships are allowed (Katta and Sethuraman 2006; Heo and Yılmaz 2015). Aziz et al. (2015) consider fair assignments when indifference allowed in preferences (but not incomparabilities). Yilmaz (2009), Athanassoglou and Sethuraman (2011) extend PS to the housing markets problem (Shapley and Scarf 1974). Bouveret, Endriss, and Lang (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.

CP-net Preferences

Compact preference representations are a common approach to deal with the preference formation and elicitation bottleneck faced in MTRAs, where the number of bundles grows exponentially with the number of types. CP-nets (Boutilier et al. 2004b) are perhaps the most well-studied and natural compact preference representation language allowing agents to express conditional (in)dependence of their preferences over combinations of different types (see Example 1). CP-nets are an important restriction on the domain of partial preferences, and induce a partial ordering on the set of all bundles. Sikdar, Adali, and Xia (2017) design mechanisms for multi-type housing markets under lexicographic extensions of CP-nets. Several works in the combinatorial voting literature assume CP-net preferences (Rossi, Venable, and Walsh 2004; Lang 2007), and that agents' CP-nets have a common dependence structure (see (Lang and Xia 2016) for a recent survey).

Preliminaries

A *multi-type resource allocation problem (MTRA)* (Mackin and Xia 2016), is given by a tuple (N, M, R). Here, (1) $N = \{1, \ldots, n\}$ 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 $i \leq p$, 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) $R = (\succ_j)_{j \leq n}$ is a *preference profile*, where for each $j \leq n$, $\succ_j$ represents the preference of agent j, and R_{-j} represents the preferences of agents in $N \setminus \{j\}$. We use $\mathcal{R}$ to denote the set of all possible preference profiles.

Bundles. For any type $i \leq p$, we use k_i or k_t to refer to the k-th item of type i where t represents the name 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 $\Pi_S = \bigtimes_{D \in S} D$, and $-S = T \setminus S$. For any $S \subseteq T$, $\hat{S} \subseteq T \setminus S$, and any $\mathbf{x} \in \Pi_S, \mathbf{y} \in \Pi_{\hat{S}}$, $(\mathbf{x}, \mathbf{y})$ denotes the bundles consisting of all items in $\mathbf{x}$ and $\mathbf{y}$. For any $S \subseteq T$, $D \in T \setminus S$, and any $\mathbf{x} \in S, o \in D$, $(o, \mathbf{x})$ denotes the bundles consisting of o and the items in $\mathbf{x}$.

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 by $G_\succ$, to be the directed graph whose nodes are 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 $\mathbf{z} \succ \mathbf{y}$. 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. 2004a) $\succ$ over the set of variables $\mathcal{D}$ has two parts: (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 $\mathbf{x} \in \Pi_{Pa(D_i)}$, 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 $\{(o, \mathbf{x}, \mathbf{z}) \succ (\hat{o}, \mathbf{x}, \mathbf{z}) : i \leq p; o, \hat{o} \in D_i; o \succ^{\mathbf{x}} \hat{o}; \mathbf{x} \in \Pi_{Pa(D_i)}; \mathbf{z} \in \Pi_{-(Pa(D_i) \cup \{D_i\})}\}$. A CP-profile is a profile of agents' preferences, each of which is represented by an acyclic CP-net.

Example 1. *Consider MTRA (N, M, R) with $p = 2$ types, food (F) and beverage (B), where $N = \{1, 2\}, M = \{1_F, 2_F, 1_B, 2_B\}$, where 1_B is item 1 of type B and so on. Let agent 1's preference $\succ_1$ be represented by the acyclic CP-net in Figure 1, where the dependency graph (Figure 1 (a)) shows that her preference on type B depends on her assignment in type F. The corresponding conditional preference tables (Figure 1 (b)) show that agent 1 prefers 1_B with 1_F, and 2_B with 2_F. This induces the preference graph in Figure 1 (c) which happens to be a linear order. Let agent 2's preference $\succ_2$ be represented by the preference graph in*

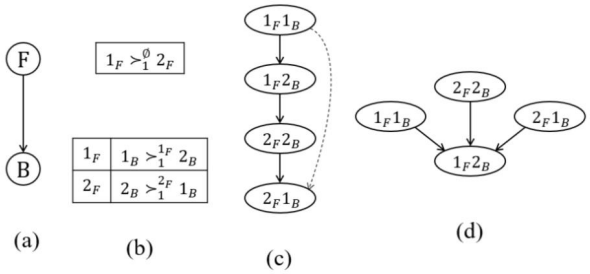

Figure 1: Agent 1's preferences are represented by an acyclic CP-net with dependence graph (a), and CP-tables (b), whose preference graph is in (c). Agent 2's preferences are represented by the preference graph in (d).

Figure 1 (d) which represents a partial order, where 1_F2_B is the least preferred bundle, and providing no information on the relative ordering of other bundles. □

Assignments. A *discrete 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* P is a combination of all agents' fractional allocations, and can be represented by a matrix $P = [p_{j,\mathbf{x}}]_{j\le n,\mathbf{x}\in\mathcal{D}}$, $P \in [0,1]^{|N|\times|\mathcal{D}|}$, such that (i) for every $j \le n$, $\sum_{\mathbf{x}\in\mathcal{D}} p_{j,\mathbf{x}} = 1$, (ii) for every $o \in M$, $S_o = \{\mathbf{x} : \mathbf{x} \in \mathcal{D} \text{ and } o \in \mathbf{x}\}$, $\sum_{j\le n,\mathbf{x}\in S_o} p_{j,\mathbf{x}} = 1$. 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 assignments.

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 use $f(R)$ to refer to the fractional assignment output by f, and for any agent $j \le n$ and any bundle $\mathbf{x} \in \mathcal{D}$, we use $f(R)_{j,\mathbf{x}}$ to refer to the value of the element of the matrix $f(R)$ indexed by j and $\mathbf{x}$.

Stochastic Dominance. In this paper we will use the natural extension of *stochastic dominance* (Bogomolnaia and Moulin 2001) to compare fractional allocations.

Definition 1. (stochastic dominance for fractional allocations) *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 $\mathbf{x} \in \mathcal{D}$, $\sum_{\hat{\mathbf{x}}\in U(\succ,\mathbf{x})} p_{\hat{\mathbf{x}}} \ge \sum_{\hat{\mathbf{x}}\in U(\succ,\mathbf{x})} q_{\hat{\mathbf{x}}}$.*

We use $P \succ_j^{sd} Q$ to denote $P(j) \succ_j^{sd} Q(j)$. We write $P \succ^{sd} Q$ if for every $j \le n$, we have $P \succ_j^{sd} Q$.

Example 2. *Consider the situation in Example 1 with the three fractional assignments (1), (2), and (3) shown below. The four upper contour sets for agent 1 are $\{1_F1_B\}$, $\{1_F1_B, 1_F2_B\}$, $\{1_F1_B, 1_F2_B, 2_F2_B\}$, and $\{1_F1_B, 1_F2_B, 2_F1_B, 2_F2_B\}$. The allocations of the four upper contour sets for agent 1 in assignment (2) are 0.5, 0.5, 1, 1, respectively and in assignment (3) are 0, 0.5, 0.5, 1. The former is greater than or equal to the latter respectively, and the same can be concluded for agent 2 by considering the upper contour sets $\{1_F1_B\}$, $\{2_F2_B\}$, $\{2_F1_B\}$, and $\{1_F1_B, 1_F2_B, 2_F1_B, 2_F2_B\}$. Hence assignment (2) stochastically dominates (3). We can check that (2) does not dominate (1), and that the reverse is true by considering the sets of $\{1_F1_B\}$ and $\{2_F1_B\}$ for agent 2.*

	1_F1_B	1_F2_B	2_F1_B	2_F2_B
Agent 1	*0.5*	*0.5*	*0*	*0*
Agent 2	*0*	*0*	*0.5*	*0.5*

(1)

	1_F1_B	1_F2_B	2_F1_B	2_F2_B
Agent 1	*0.5*	*0*	*0*	*0.5*
Agent 2	*0.5*	*0*	*0*	*0.5*

(2)

	1_F1_B	1_F2_B	2_F1_B	2_F2_B
Agent 1	*0*	*0.5*	*0.5*	*0*
Agent 2	*0.5*	*0*	*0*	*0.5*

(3)

□

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* discrete assignments, (iii) **sd-envy-freeness**, if for every pair of agents $j, \hat{j} \le n$, $P(j) \succ_j^{sd} P(\hat{j})$, (iv) **weak-sd-envy-freeness**, if for every pair of agents $j, \hat{j} \le n$, $P(\hat{j}) \succ_j^{sd} P(j) \implies P(j) = P(\hat{j})$, (v) **equal treatment of equals**, if for every pair of agents $j, \hat{j} \le n$ such that agents j and $\hat{j}$ have the same preference, $P(j) = P(\hat{j})$, (vi) **ordinal fairness**, if for every bundle $\mathbf{x} \in \mathcal{D}$ and every pair of agents $j, \hat{j} \le n$ with $P_{j,\mathbf{x}} > 0$, $\sum_{\hat{\mathbf{x}}\in U(\succ_j,\mathbf{x})} P_{j,\hat{\mathbf{x}}} \le \sum_{\hat{\mathbf{x}}\in U(\succ_{\hat{j}},\mathbf{x})} P_{\hat{j},\hat{\mathbf{x}}}$, and (vii) **decomposability**, if P can be represented as a probability distribution over discrete assignments.

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

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 $j \le n$, $\succ'_j$ 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 $\hat{j} \leq n$, $j \leq n$, $R \in \mathcal{R}$, $R' \in \mathcal{R}$, 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 $f(R)$.

Mechanisms for MTRAs with Partial Preferences

In this section, we propose MRP (Algorithm 1 as extension of RP), MPS (Algorithm 2, as extension of PS), and MGD (Algorithm 3), which can be seen as not only an eating algorithm but a special random priority algorithm.

The three mechanisms operate on a modified preference profile of strict preferences, where for every agent with partial preference $\succ$, 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. It is easy to see that no available bundle is preferred over $Ext(\succ', M')$ according to $\succ$.

Our results apply to arbitrary deterministic topological sortings (induced by a fixed ordering over items), even though different topological sortings may lead to different outputs.

Algorithm 1 MRP

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

1: For each $j \leq n$, compute a linear ordering $\succ'_j$ corresponding to a deterministic topological sort of $G_{\succ_j}$.
2: $P \leftarrow 0^{|N|\times|\mathcal{D}|}$ and $M' \leftarrow M$.
3: Pick a random priority order $\rhd$ over agents.
4: 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'.
5: **return** P

Given an instance of MTRA with agents' partial preferences, MRP fixes an arbitrary deterministic topological sorting $\succ'$ of agents' preferences, and sorts agents uniformly at random. Then agents get one unit of their favorite available bundle from the remaining M' in turns as in RP.

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 an 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.

Given an instance of MTRA with agents' partial preferences, MGD proceeds by operating on an arbitrary deterministic topological sorting $\succ'$. Let $Group(j, \succ')$ denote the set of all agents who have the same order with agent j in $\succ'$. MGD proceeds in n rounds as follows. In each round

Algorithm 2 MPS

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

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

Algorithm 3 MGD

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

1: For each $j \leq n$, compute a linear ordering $\succ'_j$ corresponding to a deterministic topological sort of $G_{\succ_j}$.
2: $P \leftarrow 0^{|N|\times|\mathcal{D}|}$ and $M' \leftarrow M$.
3: **for** $j = 1$ to n **do**
4: $top(j) \leftarrow Ext(\succ'_j, M')$.
5: For each $j' \in Group(j, \succ')$, let

$$P_{j',top(j)} \leftarrow \frac{1}{|Group(j, \succ')|}$$

6: $M' \leftarrow M' \setminus \{o \in M : o \text{ is in } top(j)\}$.
7: **return** P

$j \leq n$, agent j comes and invites other agents with the same topological sort $\succ'_j$ to consume her favorite available bundle $Ext(\succ'_j, M')$ with $\frac{1}{|Group(j,\succ')|}$ unit of time and one unit of eating rate. We present MGD as its eating algorithm version here and show that $MGD(R)$ is the expected result of a special random priority algorithm when we prove its decomposability.

Example 3. *Consider the situation in Example 1, the topological sort of agent* 1*'s preference is* $1_F1_B \succ' 1_F2_B \succ' 2_F2_B \succ' 2_F1_B$ *but both* $2_F1_B \succ' 1_F1_B \succ' 2_F2_B \succ' 1_F2_B$ *and* $1_F1_B \succ' 2_F2_B \succ' 2_F1_B \succ' 1_F2_B$ *can be the topological sort of agent* 2*'s preference. If the topological sort of agent* 2 *is the former, in MPS, agent* 1 *consumes* 1_F1_B *and agent* 2 *consumes* 2_F1_B *at the beginning. When they both consume* 0.5 *fraction,* 1_B *is exhausted.* 1_F1_B *and* 2_F1_B *become unavailable. They turn to identify the next bundles in line* 4*. For example, the first two bundles* 2_F1_B *and* 1_F1_B *are unavailable for agent* 2*, so she turns to the*

third bundle 2_F2_B. Agent 1 turns to 1_F2_B. Then MPS goes to next round of consumption until all left items are exhausted at the same time. The result is shown in assignment (1) in Example 2. However, if the topological sort of agent 2 is the latter, agent 1 and agent 2 both get 0.5 of 1_F1_B and 0.5 of 2_F2_B as assignment (2) shows. It is easy to check that MRP has the same conclusion. Suppose agent 1 has the same partial preference with agent 2 in Example 1 in MGD. If their topological sorts are both the former, agent 1 invites agent 2 to consume 2_F1_B in round 1 and agent 2 invites agent 1 to consume 1_F2_B in round 2. But if their topological sorts are both the latter, agent 1 invites agent 2 to consume 1_F1_B in round 1 and agent 2 invites agent 1 to consume 2_F2_B in round 2. □

There is an unique best available bundle w.r.t. any acyclic CP-net preference and remaining supply of items, which can be computed in polynomial time by induction on the types according to the dependency graph as we show in Proposition 1. This is an extension of the well-known result of Boutilier et al. (2004a) that there is a unique best bundle w.r.t. any acyclic CP-net.

Proposition 1. *Let $T' = \{D'_1 \subseteq D_1, \ldots, D'_p \subseteq D_p\}$, $\mathcal{D}' = \Pi_{T'}$, and let $\succ$ be any acyclic CP-net over $\mathcal{D}$. Then, there exists unique $\mathbf{x} \in \mathcal{D}'$ such that for every $\mathbf{y} \neq \mathbf{x} \in \mathcal{D}'$, $\mathbf{x} \succ \mathbf{y}$.*

All missing proofs can be found in a full version on arXiv.

We use $Top(\succ, M')$ to denote the best available bundle in remaining M' given an acyclic CP-net $\succ$. Under the domain restriction of acyclic CP-net preferences, for any topological sorting algorithms, $Ext(\succ'_j, M')$ is exact $Top(\succ_j, M')$ for any M' and $j \leq n$ by Proposition 1. Therefore, we can remove line 1 and instead $Ext(\succ'_j, M')$ with $Top(\succ_j, M')$ in the three mechanisms to save the time and space.

Proposition 2. *MRP, MPS and MGD run in $O(n^{p+1})$ time.*

We note that the size of the preference representation is $O(n^{p+1})$, and forms a part of the input.

Properties under General Partial Preferences

Theorem 1. *Under general partial preferences, MRP satisfies ex-post-efficiency, weak-sd-envy-freeness, equal treatment of equals, weak-sd-strategyproofness, and decomposability.*

The proof of weak-sd-envy-freeness involves showing that for any two agents j and $\hat{j}$, agents j and $\hat{j}$ receive equal shares of every bundle in expectation, due to a bijective mapping from the set of orders where j picks before $\hat{j}$ and vice-versa.

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

Given an assignment P and a partial preference profile R, for any $\mathbf{x}, \hat{\mathbf{x}} \in \mathcal{D}$, $(\mathbf{x}, \hat{\mathbf{x}})$ is an *improvable tuple*, denoted by $Imp(P, R)$, if there exists an agent $j < n$ such that $\mathbf{x} \succ_j \hat{\mathbf{x}}$ and $p_{j,\hat{\mathbf{x}}} > 0$. We use $Imp(P)$ for short when the preferences are clear from the context. Bogomolnaia and Moulin (2001) show that an assignment P is sd-efficient if and only if the binary relation $Imp(P)$ has no cycle in single-type resources allocations, but the sufficient condition fails to hold for MTRAs. As Example 2 shows, assignment (3) is sd-inefficient, but the set of the improvable tuples, $\{(1_F1_B, 1_F2_B), (1_F1_B, 2_F1_B), (1_F2_B, 2_F1_B), (2_F2_B, 2_F1_B)\}$, has no cycle.

Theorem 2. *Under general partial preferences, MPS satisfies sd-efficiency, weak-sd-envy-freeness, and equal treatment of equals.*

To prove the sd-efficiency of MPS, we relax the cycle from bundles to items and find a sufficient condition for sd-efficiency in MTRAs under general partial preferences. For example in assignment (3), agent 1 can extract 1_B from 2_F1_B to match 1_F and extract 2_B from 1_F2_B to match 2_F to improve her result. We show MPS satisfies sd-efficiency by proving MPS satisfies the sufficient condition. We find that one agent does not envy others at the first bundle of her topological sort since she consumes the bundle from the beginning to the end when the bundle is exhausted. Based on that, We can prove she does not envy others at any upper contour set under MPS by induction on bundles under the topological sorting.

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

PS is both sd-efficient and decomposable in single-type resources allocations (Bogomolnaia and Moulin 2001), but not decomposable in MTRAs. Serial dictatorship (Mackin and Xia 2016; Hosseini and Larson 2019) maintains both sd-efficiency and decomposability even in MTRAs but perform badly in fairness. Thankfully, MGD not only satisfies sd-efficiency and decomposability but the common requirement for fairness, equal treatment of equals.

Theorem 3. *Under general partial preferences, MGD satisfies sd-efficiency, ex-post-efficiency, equal treatment of equals, and decomposability.*

When agents have different topological sorts with each other, MGD comes to serial dictatorship. Since the dictator only invite agents who have the same topological sort with her to share her bundle, the sd-efficiency of serial dictatorship maintains in MGD. Since agents with the same preference have the same topological sorts, MGD satisfies equal treatment of equals. MGD is decomposable since its output can be seen as an expected result of a special random priority algorithm. Decomposability and sd-efficiency induce ex-post-efficiency (Bogomolnaia and Moulin 2001).

Remark 3. *MGD is not ordinally fair, weak-sd-envy-free, upper invariant and weak-sd-strategyproof even in single-type resources allocations under linear preferences.*

PS is both sd-efficient and sd-envy-free under linear preferences (Bogomolnaia and Moulin 2001) but this is no longer true under general partial preferences, as the impossibility result in Proposition 3 shows.

Proposition 3. *No mechanism can satisfy both sd-efficiency and sd-envy-freeness under general partial preferences.*

Properties under Acyclic CP-net Preferences

Theorem 4. *Given any CP-profile R, $MRP(R)$ is sd-strategyproof.*

Proof. Let P and P' be the expected assignments of MRP given the CP-profile R and $R' \in \mathcal{R}$ such that $R' = (\succ'_j, \succ_{-j})$ for some $j \leq n$. For an arbitrary fixed priority order, in any agent j's turn, the set of available bundles is the same under R and R'. By Proposition 1, the best available bundle is also the same and unique. If j lies, she may get a smaller share of her best available bundle. Therefore, the result of a lie is stochastic dominated by the result of truthfulness. Since this is true for any fixed priority order, P' is stochastic dominated by P. □

Proposition 4. *Given any CP-profile R, $MRP(R)$ is upper invariant for any other CP-profile R'.*

Thanks to Proposition 1, given any upper invariant transformation at some $\mathbf{y} \in \mathcal{D}$, if $\mathbf{y}$ is available in the misreport agent's turn, she gets the same bundle despite her lie. Otherwise if $\mathbf{y}$ is unavailable, the lie does not affect the assignment of $\mathbf{y}$.

Theorem 5. *Given any CP-profile R, $MPS(R)$ is sd-envy-free and ordinally fair.*

Proof. W.l.o.g we only prove the case between agent 1 and agent 2. Let $P = MPS(R)$. Let $\mathcal{D}_1 = \{\mathbf{x} \in \mathcal{D} : P_{1,\mathbf{x}} > 0\}$ and $n_1 = |\mathcal{D}_1|$. By Proposition 1, we have an order over $\mathcal{D}_1$ such that $\mathbf{x}_1 \succ_1 \mathbf{x}_2 \succ_1 \cdots \succ_1 \mathbf{x}_{n_1}, \mathbf{x}_i \in \mathcal{D}_1$ for any $i \leq n_1$. For agent 2, we can define $\mathcal{D}_2$ and n_2 and have an order over $\mathcal{D}_2$, $\hat{\mathbf{x}}_1 \succ_2 \hat{\mathbf{x}}_2 \succ_2 \cdots \succ_2 \hat{\mathbf{x}}_{n_2}$, similarly.

(1) sd-envy-freeness. We need to prove $P(1) \succ_1^{sd} P(2)$. For any $\mathbf{y} \in \mathcal{D}$, let $\mathbf{x}_i$ be the least favorable bundle of agent 1 in $U(\succ_1, \mathbf{y}) \cap \mathcal{D}_1$. If $i = n_1$, $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{1,\mathbf{x}} = \sum_{k=1}^{n_1} P_{1,\mathbf{x}_k} = 1 \geq \sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{2,\mathbf{x}}$. If $i < n_1$, $\mathbf{x}_{i+1} \notin U(\succ_1, \mathbf{y})$. When agent 1 starts to consume $\mathbf{x}_{i+1}$, $\hat{\mathbf{y}}$ is unavailable for any $\hat{\mathbf{y}} \in U(\succ_1, \mathbf{y})$. Otherwise by Proposition 1, we have $\mathbf{x}_{i+1} \succ_1 \hat{\mathbf{y}}$ or $\mathbf{x}_{i+1} = \hat{\mathbf{y}}$ both indicating $\mathbf{x}_{i+1} \in U(\succ_1, \mathbf{y})$, a contradiction. Suppose $t(\mathbf{x}_i)$ be the time when $\mathbf{x}_i$ is exhausted and agent 1 starts to consume $\mathbf{x}_{i+1}$. Then $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{1,\mathbf{x}} = \sum_{k=1}^{i} P_{1,\mathbf{x}_k} = t(\mathbf{x}_i) \geq \sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{2,\mathbf{x}}$ for any $\mathbf{y} \in \mathcal{D}$. Hence $P(1) \succ_1 P(2)$.

(2) ordinal fairness. For any $\mathbf{y} \in \mathcal{D}_1$, we need to prove $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{1,\mathbf{x}} \leq \sum_{\mathbf{x} \in U(\succ_2, \mathbf{y})} P_{2,\mathbf{x}}$. For any $\mathbf{x}_i \in \mathcal{D}_1$, suppose that when $\mathbf{x}_i$ is exhausted, agent 2 is consuming $\hat{\mathbf{x}}_{\hat{i}} \in \mathcal{D}_2$ or $\hat{\mathbf{x}}_{\hat{i}}$ is exhausted at the same time. If $\hat{\mathbf{x}}_{\hat{i}} = \mathbf{x}_i$, we have $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{x}_i)} P_{1,\mathbf{x}} = \sum_{\mathbf{x} \in U(\succ_2, \mathbf{x}_i)} P_{2,\mathbf{x}}$. If $\hat{\mathbf{x}}_{\hat{i}} \neq \mathbf{x}_i$, since $\mathbf{x}_i$ is available when agent 2 starts to consume $\hat{\mathbf{x}}_{\hat{i}}$, we have $\hat{\mathbf{x}}_k \succ_2 \mathbf{x}_i$ for any $k \leq \hat{i}$. Hence, $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{x}_i)} P_{1,\mathbf{x}} = \sum_{k=1}^{i} P_{1,\mathbf{x}_k} \leq \sum_{k=1}^{\hat{i}} P_{2,\hat{\mathbf{x}}_k} \leq \sum_{\mathbf{x} \in U(\succ_2, \mathbf{x}_i)} P_{2,\mathbf{x}}$. Therefore, for any $\mathbf{y} \in \mathcal{D}$ with $P_{1,\mathbf{y}} > 0$, we have $\sum_{\mathbf{x} \in U(\succ_1, \mathbf{y})} P_{1,\mathbf{x}} \leq \sum_{\mathbf{x} \in U(\succ_2, \mathbf{y})} P_{2,\mathbf{x}}$. □

Hashimoto et al. (2014) shows that ordinal fairness characterizes PS when we come to single-type resources. However, this is not the case in MTRAs: for example, if two agents' preferences are both as (c) in Figure 1 shows, the result that two agents both get 0.5 $1_F 2_B$ and 0.5 $2_F 1_B$ is ordinally fair but not the output of MPS.

Proposition 5. *Given any CP-profile R, $MPS(R)$ is upper invariant for any other CP-profile R'.*

Given any upper invariant transformation at some $\mathbf{y} \in \mathcal{D}$, we show the consumption processes are identical until $\mathbf{y}$ is exhausted whether the misreport agent is truthful or lies, by induction on the bundles consumed by the misreport agent under R. Details are in the Appendix.

Remark 4. *MPS is not weak-sd-strategyproof under CP-net preferences.*

Proposition 6. *Given any CP-profile R with an shared dependency graph, $MPS(R)$ is weak-sd-strategyproof for any other CP-profile R' with the shared dependency graph.*

Proof sketch. Let agent 1 be the manipulator w.l.o.g. The proof involves showing that for any R' obtained from R by changing 1's preference, without modifying the dependency graph, such that $P' = MPS(R') \succ_1^{sd} MPS(R) = P$, it holds that $P'_1 = P_1$, in two steps: (Step 1) We show that agent 1 has no incentive to lie about preferences involving any type, D_i, for which preferences are conditionally independent of other types, i.e. with no incoming edges in the dependency graph. This involves showing that the assignment computed by MPS restricted to any such type D_i is the same as the output of PS applied to the preference profile restricted to preferences over such a type D_i. The claim follows from the weak-sd-strategyproofness property of PS (Bogomolnaia and Moulin 2001). (Step 2) We show by induction over the dependency graph, that agent 1 has no incentive to lie, due to two important facts: (1) agent 1 does not misreport the dependency graph, and (2) due to CP-net preferences, for any type D_i, and given an assignment for types $Pa(D_i)$, preferences over D_i are conditionally preferentially independent of all types. The formal proof and details are in the Appendix. □

Conclusion and Future Work

We proposed and studied MRP and MPS as extensions of RP and PS to MTRAs. We also proposed MGD that is both sd-efficient and decomposable. For future work, we are interested in axiomatic characterization of MRP, MPS and MGD. More generally, designing desirable mechanism for MTRAs is still an interesting and challenging open question.

Acknowledgments

We are grateful to the anonymous reviewers for their helpful comments. LX acknowledges NSF #1453542 and #1716333 for support. YC acknowledges NSFC under Grants 61772035, 61751210, and 61932001, and the National Science and Technology Major Project for IND (investigational new drug) under Grant 2018ZX09201-014 for support. HW acknowledges NSFC under Grants 61572003 and 61972005, and the National Key R&D Program under Grants 2018YFB1003904 and 2018YFC1314200 for support.

References

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

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

Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.

Bogomolnaia, A., and Heo, E. J. 2012. Probabilistic assignment of objects: Characterizing the serial rule. *Journal of Economic Theory* 147(5):2072–2082.

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

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

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

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004b. Preference-based constrained optimization with CP-nets. *Computational Intelligence* 20(2):137–157.

Bouveret, S.; Endriss, U.; and Lang, J. 2010. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI-2010)*.

Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

Brams, S. J.; Jones, M. A.; and Klamler, C. 2006. Better ways to cut a cake. *Notices of the AMS* 53(11):1314–1321.

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.

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

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.

Ghodsi, A.; Zaharia, M.; Hindman, B.; Konwinski, A.; Shenker, S.; and Stoica, I. 2011. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, 323–336.

Hashimoto, T.; Hirata, D.; Kesten, O.; Kurino, M.; and Ünver, M. U. 2014. Two axiomatic approaches to the probabilistic serial mechanism. *Theoretical Economics* 9(1):253–277.

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

Heo, E. J., and Yılmaz, Ö. 2015. A characterization of the extended serial correspondence. *Journal of Mathematical Economics* 59:102–110.

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

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

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

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 *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, 1366–1371.

Mackin, E., and Xia, L. 2016. Allocating indivisible items in categorized domains. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 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.

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

Rossi, F.; Venable, K. B.; and Walsh, T. 2004. mCP nets: Representing and reasoning with preferences of multiple agents. In *Proceedings of AAAI-04*, 729–734.

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

Sandomirskiy, F., and Segal-Halevi, E. 2019. Fair division with minimal sharing. *arXiv preprint arXiv:1908.01669*.

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*.

Sikdar, S.; Adali, S.; and Xia, L. 2018. Top-trading-cycles mechanisms with acceptable bundles. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI*, volume 18.

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