Assigning indivisible and categorized items

Lirong Xia
Computer Science Department
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180, USA
xial@cs.rpi.edu

Abstract

In this paper, we study the problem of assigning indivisible items under the following constraints: 1) each item belongs to one of the p categories, 2) each agent is required to get exactly one item from each category, 3) no free disposal, and 4) no monetary transfer.

We first characterize serial dictatorships by a minimal set of 3 properties: *strategy-proofness*, *non-bossiness*, and *category-wise neutrality*. Then, we analyze a natural extension of serial dictatorships called *sequential allocation mechanisms*, which allocates the items in muitple stages according to a given order over all (agent,category) pairs, so that in each stage, the designated agent chooses an item from the designated category. We then focus on the cases where each agent is either *optimistic* or *pessimistic*, and for any sequential allocation, we characterize a tight lower bound on the rank of the allocated bundle for each agent.

Introduction

Suppose you are organizing a conference and you want to allocate catered lunch to n agents (conference attendees). The lunch consists of three categories of foods: wraps, deserts, and bottles of drinks. Each category contains exactly n indivisible and heterogeneous items,[1] and each agent must get exactly one item in each category. Agents may have different preferences over the (wrap,desert,drink) combinations. How would you allocate these items?

Many other allocation problems share this categorized feature: allocating rooms, tables, chairs, and computers among a group of students; allocating tent, car, GPS among a group of families for camping out; allocating time slots and topics for seminar paper presentation.

These are a special case of a more general (and much harder) problem that has been extensively studied by researchers in economics and artificial intelligence. Suppose there are m indivisible items and n agents. The agents may have different private preferences over bundles of items, and we want to design a *mechanism* to allocate these items to agents without monetary transfer. This problem in its most general form is too hard due to high complexity of preference representation and communication. Therefore, most previous work focused on designing sensible mechanisms with various constraints on agents' preference structures.

In economics, these problems are known as *assignment problems*. Most previous research focused on designing mechanisms that satisfy some desired properties, including Pareto-optimality and strategy-proofness, for assignment problems with different constraints on agents' preferences [15, 14, 16, 9, 12, 8, 18]. Many of these papers characterize *serial dictatorships* or their variants. A mechanism is a serial dictatorship if the agents pick their top-ranked available bundles in turn, according to a given order.

In artificial intelligence, this type of problems are known as *multiagent resource allocation (MARA)* [6]. Most previous research focused on designing *centralized* and *decentralized* (usually negotiation schemes) mechanisms with an emphasis on computational aspects of preference representation, communication, and computing the optimal assignment.

In this paper, we focus on designing mechanisms to assign indivisible items that are categorized, and each agent is required to get exactly one item from each category. Following the convention in economics, we call such problems *assignment problems with indivisible and categorized items*. Formally, suppose there are $n \geq 2$ agent and $p \geq 1$ categories of items, denoted by $\mathcal{I} = \{D_1, \ldots, D_p\}$, where each category D_i is a set of n indivisible items, that is, $|D_1| = |D_2| = \cdots = |D_p| = n$. Agents' preferences are represented by linear orders over $\mathcal{D} = D_1 \times \cdots \times D_p$. We want to design a mechanism f to assign these np items to the agent such that for every $i \leq p$, every agent gets exactly one item in D_i.

Our contributions

We first give a characterization of serial dictatorships by showing that a mechanism that allocates indivisible and categorized items and satisfies *strategy-proofness*, *non-bossiness*, and *category-wise neutrality* if and only if it is a serial dictatorship. Moreover, we also show that the three properties are a minimal set of properties that characterizes serial dictatorships.

Then, we move on to analyze a natural extension of serial dictatorships called *sequential allocation mechanisms*, which allocates the items in np stages according to a given order over all (agent,category) pairs in the following way: in

[1] For example, wraps with different meats are served. Even for wraps with the same kind of meat, we still assume that an agent will have (slightly) different preferences.

stage $i = 1, \ldots, np$, suppose (a, c) is ranked at the ith place in the order, then agent a picks an item in category c. We consider two types of agents: in each stage, an *optimistic* agent always chooses the item in her top-ranked bundle that is still available, while a pessimistic agents always chooses the item that is best in the worst-case scenario. Our main results are: When each agent is either optimistic or pessimistic, for any sequential allocation mechanism, we give a lower bound on the rank of the allocated bundle for each agent (where the worst-case is taken over all agents' preference profile). Moreover, we show that all these lower bounds are tight and can be reached in the same preference profile, and more surprisingly, in this preference profile, there is a way to assign the items so that $n-1$ agents get their top-ranked bundle, and the remaining agent gets her top-ranked or second-ranked bundle.

Most previous work in economics focused on assignment problems with different constraints on the agents' preferences, e.g. quota constraints [15], monotonic preferences [8, 9], and additive preferences [9]. Most previous work in multiagent systems focused on various compact representations, including k-additive preferences [7], CI-nets [3], and SCI-nets [4]. Recently welfare properties and computational complexity of equilibrium under serial dictatorships have been studied for additive preference [5, 10, 11]. Our work is also related to preference representation and aggregation in multi-issue domains, especially combinatorial voting [13].

We are not aware of any previous work that focuses on the assignment problem with categorized items (or equivalently, agents' preferences are represented by linear orders over bundles). We do not see any previous results characterization results for serial dictatorships imply ours. We are also not aware of similar results as ours for sequential allocations. The closest might be [5, 10] but these papers focused on the social welfare under specific utility functions, while our results reveals the worst-case rank of the allocated bundle for each single agent.

Preliminaries

Suppose there are $n \geq 2$ agent and $p \geq 1$ categories of items, denoted by $\mathcal{I} = \{D_1, \ldots, D_p\}$, where each category D_i is a set of n indivisible items $\{1_i, \ldots, n_i\}$. Sometimes the subscripts are omitted. Agents' preferences are represented by linear orders over $\mathcal{D} = D_1 \times \cdots \times D_p$. Each element in $\mathcal{D}$ is called a *bundle*. For any $d, j \leq n$ and any $i \leq p$, we use $([j]_i, [d]_{-i})$ denote the bundle where the ith component is j and the other components are all d, where "$[\cdots]$" is used to distinguish $[j]_i$ from the variable j_i. For any $j \leq n$, let R_j denote a linear order over $\mathcal{D}$, and let $P = (R_1, \ldots, R_n)$ denote a *preference profile*. An *assignment* (or *allocation*) A is a mapping from $\{1, \ldots, n\}$ to $\mathcal{D}$ such that for every $j \leq n$, $A(j)$ is the bundle assigned to agent j, which means that for any $i \leq p$, we have $\cup_{j=1}^{n}(A(j))_i = D_i$, where $(A(j))_i$ is the item in category i that is assigned to agent j. A *mechanism* f is a mapping that takes any preference profile as input, and outputs an assignment. In this paper we sometimes use $f^j(P)$ to denote $(f(P))(j)$, that is, the bundle allocated to agent j when the preference profile is R.

We say a mechanism f satisfies *strategy-proofness*, if no agent can benefit from misreporting here preferences. f satisfies *non-bossiness*, if no agent is *bossy* in f. An agent is bossy if she can report differently to change the bundles allocated to some other agents, while keeping her own allocation unchanged. That is, j is bossy if there exists a profile P and R'_j such that $f^j(P) = f^j(P_{-j}, R'_j)$ and there exists l with $f^l(P) \neq f^l(P_{-j}, R'_j)$. f satisfies *category-wise neutrality*, if we apply any permutations to the preference profile such that each of these permutations only permutes the names of items within the same category, then the final allocation is permuted in the same way.

A mechanism is a *serial dictatorship*, if there exists a linear order $\mathcal{K}$ over $\{1, \ldots, n\}$ such that for any preference profile P, agents pick their top-ranked available bundle sequentially according to $\mathcal{K}$.

A characterization of serial dictatorships

In this section we characterize serial dictatorship by a minimal set of 3 properties: strategy-proofness, non-bossiness, and category-wise neutrality. Due to the space constraints, many proofs are omitted. All omitted proofs can be found on the author's homepage.

We will frequently use the following two classical lemmas for strategy-proof mechanisms. Similar ones were proved for other situations.

Lemma 1. *For any n, m, let f be a strategy-proof and non-bossy allocation mechanism. For any preference profile P and any preference profile P' such that for all $j \leq n$ and all bundles $\vec{d}$, $[f^j(P) \succ_{R_j} \vec{d}] \implies [f^j(P) \succ_{R'_j} \vec{d}]$, we have $f(P') = f(P)$.*

Proof. We first prove the lemma for the special case where P and P' only differ on one agent's preferences. Let j be an agent such that $R'_j \neq R_j$ and for all $\vec{d}$, $\vec{d}$, $[f^j(P) \succ_{R_j} \vec{d}] \implies [f^j(P) \succ_{R'_j} \vec{d}]$. We will prove that $f^j(R'_j, R_{-j}) = f^j(R_j, R_{-j})$.

Suppose for the sake of contradiction $f^j(R'_j, R_{-j}) \neq f^j(R_j, R_{-j})$. If $f^j(R'_j, R_{-j}) \succ_{R_j} f^j(R_j, R_{-j})$ then it means that f is not strategy-proof since j has incentive to report R'_j when her true preferences are R_j. If $f^j(R_j, R_{-j}) \succ_{R_j} f^j(R'_j, R_{-j})$ then $f^j(R_j, R_{-j}) \succ_{R'_j} f^j(R'_j, R_{-j})$, which means that when agent j's preferences are R'_j she has incentive to report R_j, which again contradicts the assumption that f is strategy-proof. Therefore $f^j(R_j, R_{-j}) = f^j(R'_j, R_{-j})$.

By non-bossiness $f(R_j, R_{-j}) = f(R'_j, R_{-j})$. Then we recursively apply this for $j = 1, \ldots, n$, which leads to the lemma. $\square$

For any ranking R over $\mathcal{D}$ and any bundle $\vec{d} \in \mathcal{D}$, we say ranking R' is a *pushup* of $\vec{d}$ from R if R' can be obtained from R by raising the position of $\vec{d}$ while keeping the relative positions of other bundles unchanged.

Lemma 2. *Let f be a stragyproof and non-bossy mechanism. For any profile P, any $j \leq n$, any bundle $\vec{d}$, and any R'_j that is a pushup of $\vec{d}$ from R_j, either 1) $f(R'_j, R_{-j}) = f(R)$ or 2) $f(R'_j, R_{-j}) = \vec{d}$.*

Proof. Suppose on the contrary that $f(R'_j, R_{-j})$ is neither $f(R)$ nor $\vec{d}$. If $f(R'_j, R_{-j}) \succ_R f(R)$, then f is not strategy-proof since when agent j's true preferences are R_j and other agents' preferences are R_{-j}, she has incentive to report R'_j to make her allocation better. If $f(R) \succ_R f(R'_j, R_{-j})$, then since $\vec{d} \neq f(R'_j, R_{-j})$, we have $f(R) \succ_{R'} f(R'_j, R_{-j})$. In this case when agent j's true preferences are R'_j and other agents' preferences are R_{-j}, she has incentive to report R_j to make her allocation better, which means that f is not strategy-proof. □

We prove that strategy-proofness, non-bossiness, and category-wise neutrality imply Pareto-optimality.

Proposition 1. *For any m and n, any strategy-proof, non-bossy, and category-wise neutral allocation mechanism is Pareto optimal.*

Proof. We prove the proposition by contradiction. Suppose f is a strategy-proofness, non-bossiness, category-wise neutral allocation mechanism and f is not Pareto-optimal. Let P denote a profile such that $f(P)$ is Pareto dominated by an allocation A. For any $i \leq m$, let M_i denote the permutation over D_i so that for every $j \leq n$, $(f^j(P))_i$ is permuted to $(A^j)_i$. Since bother $f(P)$ and A allocates all items, M_i is well defined. Let $M = (M_1, \ldots, M_m)$. It follows that for all $j \leq n$, $M(f^j(P)) = A^j$.

Let R'_j denote any ranking obtained from R_j by raising A^j to the top place, and then raising $f^j(P)$ to the second place if it is different from A^j, and the other bundles are ranked arbitrarily. Let R^*_j denote any ranking obtained from R_j by raising $f^j(P)$ to the top place, and then raising A^j to the second place if it is different from $f^j(P)$, and the other bundles are ranked arbitrarily. Let $P' = (R'_1, \ldots, R'_n)$ and $P^* = (R^*_1, \ldots, R^*_n)$. Since A Pareto dominates $f(P)$, by Lemma 1 we have $f(P') = f(P)$. Also by Lemma 1 we have $f(P^*) = f(P)$. By category-wise neutrality $f(M(P')) = M(f(P')) = A$. However, the only differences between $M(P')$ and P^* are the orderings among $\mathcal{D} \setminus \{A^j, f^j(P)\}$. By Lemma 1 $f(P) = f(P^*) = f(M(P')) = A$, which is a contradiction. □

Lemma 3. *Given any m, n and any allocation mechanism f that satisfies strategy-proofness and non-bossiness. For any preference profile P, and any $j_1 \neq j_2 \leq n$, let $\vec{a} = f^{j_1}(P)$ and $\vec{b} = f^{j_2}(P)$, there is no $\vec{c} \in \{a_1, b_1\} \times \{a_2, b_2\} \times \cdots \times \{a_m, b_m\}$ such that $\vec{c} \succ_{R_{j_1}} \vec{a}$ and $\vec{c} \succ_{R_{j_2}} \vec{b}$.*

Theorem 1. *For any $m \geq 2$ and $n \geq 2$, an allocation mechanism is strategy-proof, non-bossy, and category-wise neutral if and only if it is a serial dictatorship. Moreover, strategy-proofness, non-bossiness and category-wise neutrality are a minimal set of properties that characterize serial dictatorships.*

Proof. The proof is inspired by proofs in [15, 16, 9].[2] We first prove that strategy-proofness, non-bossiness and category-wise neutrality imply the mechanism is a serial dictatorship. Let R be a linear order over $\mathcal{D}$ that satisfies the following conditions.

- $(1, \ldots, 1) \succ (2, \ldots, 2) \succ \cdots \succ (n, \ldots, n)$.
- For any $i < n$, the bundles ranked between $(i, \ldots, i)$ and $(i+1, \ldots, i+1)$ are those satisfying the following two conditions: 1) at least one component is i, and 2) all components are in $\{i, i+1, \ldots, n\}$. Let K_i denote these bundles. That is, $K_i \subseteq \mathcal{D}$ and $K_i = \{\vec{d} : \forall l, d_l \geq i \text{ and } \exists l', d_{l'} = i\}$.
- For any i and any $\vec{d}, \vec{e} \in K_i$, if the number of i's in d is strictly larger than the number of i's in $\vec{e}$, then $\vec{d} \succ \vec{e}$.

Let $P = (R, \ldots, R)$. We will first prove the following claim.

Claim 1. *For any $l \leq n$, there exists $j_l \leq n$ such that $f^{j_l}(P) = (l, \ldots, l)$.*

Proof. We prove the claim by induction on l. When $l = 1$. For the sake of contradiction suppose there is no j_l with $f^{j_l}(P) = (1, \ldots, 1)$. The there exist a pair of agents j_1 and j_2 such that both $\vec{a} = f^{j_1}(P)$ and $\vec{b} = f^{j_2}(P)$ contain 1 in some components.

Let $\vec{c}$ be the bundle obtained from $\vec{a}$ by changing all components where $\vec{b}$ takes 1 to 1. More precisely,

$$\vec{c} = (c_1, \ldots, c_n) \text{ s.t. } c_i = \begin{cases} 1 & \text{if } a_i = 1 \text{ or } b_i = 1 \\ a_i & \text{otherwise} \end{cases}$$

It follows that in R, $\vec{c} \succ_R \vec{a}$ and $\vec{c} \succ_R \vec{b}$ since the number of 1's in $\vec{c}$ is strictly larger than the number of 1's in $\vec{a}$ or $\vec{b}$. However, this contradicts the assumption that f is strategy-proof and non-bossy by Lemma 3. Hence there exists $j_1 \leq n$ such that $f^{j_1}(P) = (1, \ldots, 1)$.

Suppose the claim is true for $l \leq l'$. We next prove that there exists $j_{l'+1}$ such that $f^{j_{l'+1}}(P) = (l'+1, \ldots, l'+1)$. This follows from a similar reasoning as above for the $l = 1$ case. Formally, suppose for the sake of contradiction there does not exist such a $j_{l'+1}$. Then, there exist two agents who get $\vec{a}$ and $\vec{b}$ in $f(P)$, where both $\vec{a}$ and $\vec{b}$ contain $l'+1$ for some categories (but not the same). By the induction hypothesis, all components of $\vec{a}$ and $\vec{b}$ are at least $l'+1$. Let $\vec{c}$ be the bundle obtained from $\vec{a}$ by changing all components where $\vec{b}$ takes $l'+1$ to $l'+1$. We have $\vec{c} \succ_R \vec{a}$ and $\vec{c} \succ_R \vec{b}$, which is a contradiction by Lemma 3. Therefore, the claim holds for $l = l'+1$, which completes the proof. □

We are ready to prove the theorem. Let $j_1, \ldots, j_n$ denote the agents in Claim 1. For any profile $P' = (R'_1, \ldots, R'_n)$, we define n bundles as follows. Let $\vec{d}^1$ denote the top-ranked bundle at R'_{j_1}, and for any $l \geq 2$, let $\vec{d}^l$ denote j_l's

[2] However, we do not see an easy way to extend proofs in these papers to our setting.

top-ranked available bundle given that $\vec{d^1}, \ldots, \vec{d^{l-1}}$ are already chosen. That is, $\vec{d_l}$ is the bundle that is ranked in top of $\{\vec{d} : \forall l' < l, \vec{d} \cap \vec{d_{l'}} = \emptyset\}$ by R'_{j_l}. It follows that $\vec{d^1}, \ldots, \vec{d^n}$ is an allocation. Then, for any $l \leq m$, we define a category-wise permutation M_l such that for all $j \leq n$, $M_l(j) = d^j_l$. Let $M = (M_1, \ldots, M_m)$. We have that for all $j \leq n$, $M(j, \ldots, j) = \vec{d^j}$. By category-wise neutrality and Claim 1, in $f(M(P))$ agent j gets $\vec{d^j}$. We note that for all $j \leq n$, if $\vec{e} \succ_{M(R'_j)} \vec{d^j}$ then $\vec{e} \succ_{R'_j} \vec{d^j}$, otherwise it is a contradiction to the selection of $\vec{d^j}$. Therefore by Lemma 1, $f(P') = f(M(P)) = M(f(P))$, which proves that f is a serial dictatorship.

Next, we show that strategy-proofness, non-bossiness, and category-wise neutrality are a minimal set of properties that characterize serial dictatorships by giving the following examples.

strategy-proofness is necessary: Consider the allocation mechanism that maximizes the social welfare w.r.t. the following utility functions. For any $i \leq m$ and $j \leq n$, the bundle that is ranked at the ith position of agent j's preferences gets $(m-i)(1+(\frac{1}{2m})^j)$ points. It is not hard to check that for any pair of different allocations, the social welfare are different. Therefore, the allocation mechanism satisfies non-bossiness since if agent j's allocation is the same when only agent j report differently, the set of items allocated to other agents is also the same, which implies that the optimum allocation is the same. Since the definition of utilities only depends on the positions (but not on the names of the items), the allocation mechanism satisfies category-wise neutrality. It is not hard to verify that this mechanism is not a serial dictatorship. For example, consider the case of $m = n = 2$ with $R'_1 = [11 \succ 12 \succ 22 \succ 21]$ and $R'_2 = [12 \succ 21 \succ 11 \succ 22]$. A serial dictatorship will either give 11 to agent 1 and give 22 to agent 2, or give 21 to agent 1 and give 12 to agent 2, but the allocation that maximizes social welfare is to give 12 to agent 1 and give 21 to agent 2.[3]

non-bossiness is necessary: Consider the allocation mechanism that is a serial dictatorship where agent 1 picks first, and the order over agents $\{2, \ldots, n\}$ depends on the first agents' preferences in the following way: if first component of agent 1's second choice is the same as the first component of her top choice, then the order over the rest of agents is $2 \rhd 3 \rhd \cdots \rhd n$; otherwise it is $n \rhd n-2 \rhd \cdots \rhd 2$. It is not hard to verify that this mechanism satisfies strategy-proofness and category-wise neutrality, and is not a serial dictatorship (where the order must be fixed before seeing the preference profile).

category-wise neutrality is necessary: Consider the allocation mechanism that is a serial dictatorship where agent 1 picks first, and the order over agents $\{2, \ldots, n\}$ depends on the allocation of agent 1 in the following way: if agent 1 gets $(1, \ldots, 1)$, then the order over the rest of agents is $2 \rhd 3 \rhd \cdots \rhd n$; otherwise it is $n \rhd n-2 \rhd \cdots \rhd 2$. It is not hard to verify that this mechanism satisfies strategy-proofness and non-bossiness, and is not a serial dictatorship (where the order must be fixed before seeing the preference profile). □

[3] The utility functions are only used to avoid ties in allocations. Any utility functions where there are no ties satisfy non-bossiness and category-wise neutrality, but some of them are serial dictatorships.

Sequential allocation mechanisms

A sequential allocation is defined to be a total order $\mathcal{O}$ over $\{1, \ldots, n\} \times \{1, \ldots, p\}$, where (j, i) means that it is agent j's turn to choose an item from D_i. Let $\mathcal{O}^{-1}(j, i) \in [1, np]$ denote the position of (j, i) in $\mathcal{O}$.

Given $\mathcal{O}$, the sequential allocation is a distributed protocol that works in np steps as follows. For each step t, suppose the t-th element in $\mathcal{O}$ is (j, i), that is, $t = \mathcal{O}^{-1}(j, i)$. Then, agent j is required to choose an item from D_i that is still available (meaning that no agent has chosen that value before). Then, agent j's choice is sent to the center and broadcast among all agents. The protocol is illustrated in Algorithm 1.

Algorithm 1: Sequential allocation protocol.

Input: An order $\mathcal{O}$ over $\{1, \ldots, n\} \times \{1, \ldots, p\}$.
1 Broadcast $\mathcal{O}$ to all agents.
2 **for** $t = 1$ *to* np **do**
3 | Let (j, i) be the t-th in $\mathcal{O}$.
4 | Ask agent j to choose an available item $d_{j,i}$ from D_i.
5 | Broadcast $d_{j,i}$ to all agents.
6 **end**

As for combinatorial auctions and combinatorial voting, the advantage of Algorithm 1 are two-fold.

First, it has a low communication cost compared to the straightforward centralized system, where reporting a full ranking over $\mathcal{D}$ requires $O(n^p p \log n)$ bits for each agent, so totally $O(n^{p+1} p \log n)$. For sequential allocation, broadcasting $\mathcal{O}$ (Step 1) uses $O(n(np \log np)) = O(n^2 p \log np)$ bits, then in each round communicate the allocation from the active agent, and then broadcast it to the other agents takes $O(n \log n)$. Since there are totally np rounds, so the total communication complexity for Algorithm 1 is $O(n^2 p \log n + n(np \log np + p \log p) = O(n^2 p \log np)$.

Second, agents are more comfortable choosing a value for an category per step compared to reporting a full ranking over the n^p bundles in $\mathcal{D}$, when p is not too small.

Therefore, sequential allocation is an *indirect mechanism*, since the message space for the agents are different from their preference space. This is the main reason why we can have significant communicational savings. For indirect mechanisms, it is hard (if not impossible) to define the behavior of a truthful agent.

In this paper, we will investigate two types of "truthful" agents under sequential allocation mechanisms.

- Setting 1: *optimistic* agents. When an optimistic agent j is active and is asked to report a value from D_i, she will choose the ith component of the top-ranked bundle that is still available, given the results of allocation in previous steps.

- Setting 2: *pessimistic* agents. When a pessimistic agent j is active and is asked to report a value from D_i, she will choose a value $d_{j,i}$ so that the worst-case available bundle is optimized given the results of previous rounds and $d_{j,i}$.

Example 1. *Let $n = 3$, $p = 2$. Three agents' preferences are as follows, where 12 represents $(1, 2)$ and for each agent, "others" represents any order over the bundles not specified in the context (and it will not affect the outcome of sequential allocation in this example).*

Agent 1 (optimistic):	$12 \succ 21 \succ$ others $\succ 11$
Agent 2 (optimistic):	$32 \succ$ others $\succ 22$
Agent 3 (pessimistic):	$13 \succ$ others $\succ 33 \succ 31 \succ 23$

Let $\mathcal{O} = [(1,1) \rhd (2,2) \rhd (3,1) \rhd (3,2) \rhd (2,1) \rhd (1,2)]$. Suppose agent 1 and agent 2 are optimistic and agent 3 is pessimistic.

When $t = 1$, $(1, 2)$ is the top-ranked available item for agent 1. Since agent 1 is optimistic, she chooses 1. When $t = 2$, $(3, 2)$ is the top-ranked available item for agent 2. Since agent 2 is optimistic, she chooses 2. When $t = 3$, the available bundles are $\{2, 3\} \times \{1, 3\}$. If agent 3 chooses 2, then the worst-case available bundle is $(2, 3)$, and if agent 3 chooses 3 then the worst-case available bundle is $(3, 1)$. Since agent 3 prefers $(3, 1)$ to $(2, 3)$, she will choose 3 from D_1. When $t = 4$, the available bundles are $\{3\} \times \{1, 3\}$, and agent 3 will choose 3 from D_2. Then when $t = 5$, agent 2 choses 2_1 and when $t = 6$, agent 1 choses 1_2. The final allocation is: agent 1 gets $(1, 1)$, agent 2 gets $(2, 2)$, and agent 3 gets $(3, 3)$.

Efficiency of sequential allocation

In this section, we assume that each agent is either optimistic or pessimistic. We note that this is much more general than assuming that the agents are either all optimistic or all pessimistic. Given any sequential allocation $\mathcal{O}$, we will characterize the worst-case rank of the allocation by $\mathcal{O}$ for all agents. To do so, given $\mathcal{O}$, we define the following notation for any j.

- Let $\mathcal{O}_j$ denote the order over $\{1, \ldots, p\}$ according to which agent j choose values for categories in $\mathcal{I}$.
- Let K_j denote the smallest natural number such that no agent can "interrupt" agent j from choosing her top-ranked bundle that is still available for the categories $\mathcal{O}_j(K_j), \mathcal{O}_j(K_j + 1), \ldots, \mathcal{O}_j(p)$. Formally, K_j is the smallest number so that for any i' with $K_j < i' \le p$, between the round when agent j chooses $\mathcal{O}_j(K_j)$ and the round when agent j chooses $\mathcal{O}_j(i')$ according to $\mathcal{O}$ in Algorithm 1, no agent chooses a value for $\mathcal{O}_j(i')$. Note that K_j is defined solely by $\mathcal{O}$, which means that it does not depend on the agents' preferences (which will be considered in the worst-case analysis). Also note that the K_j represents a position in $\mathcal{O}_j$, which is not necessarily the K_jth category.
- For any $i \le p$, let $k_{j,i}$ denote the number of values of D_i that have *not* been chosen by other agents before agent j chooses a value from D_i. Formally, $k_{j,i} = 1 + |\{(j', i) : (j, i) \rhd_{\mathcal{O}} (j', i)\}|$. Equivalently $k_{j,i} = n - |\{(j', i) : (j', i) \rhd_{\mathcal{O}} (j, i)\}|$.

Example 2. *Let $\mathcal{O} = [(1,1) \rhd (1,2) \rhd (1,3) \rhd (2,1) \rhd (2,2) \rhd (2,3) \rhd (3,1) \rhd (3,2) \rhd (3,3)]$. That is, Algorithm 1 with $\mathcal{O}$ is a serial dictatorship. Then $\mathcal{O}_1 = \mathcal{O}_2 = \mathcal{O}_3 = 1 \rhd 2$. $K_1 = K_2 = K_3 = 1$. $k_{1,1} = k_{1,2} = 3$, $k_{2,1} = k_{2,2} = 2$, $k_{3,1} = k_{3,2} = 1$.*

Let $\mathcal{O}$ be the order in Example 1, that is, $\mathcal{O} = [(1,1) \rhd (2,2) \rhd (3,1) \rhd (3,2) \rhd (2,1) \rhd (1,2)]$.

$\mathcal{O}_1 = 1 \rhd 2$. $K_1 = 2$ since $(2, 2)$ is between $(1, 1)$ and $(1, 2)$ in $\mathcal{O}$. $k_{1,1} = 3$, $k_{1,2} = 1$.

$\mathcal{O}_2 = 2 \rhd 1$. $K_2 = 2$ since $(3, 1)$ is between $(2, 2)$ and $(2, 1)$. $k_{2,1} = 1$, $k_{2,2} = 3$.

$\mathcal{O}_3 = 1 \rhd 2$. $K_3 = 1$ since between $(3, 1)$ and $(3, 2)$ in $\mathcal{O}$ no agent chooses an item from D_2. $k_{3,1} = 2$, $k_{3,2} = 2$.

Proposition 2. *For any $j \le n$, we have the following bounds.*

- **Lower bounds for optimistic agents:** *the bundle assigned to an optimistic agent j is ranked no lower than the $\prod_{l=K_j}^{p} k_{j,O_j(l)}$-th position from the bottom.*
- **Lower bounds for pessimistic agents:** *the bundle assigned to a pessimistic agent j is ranked no lower than the $(1 + \sum_{l=1}^{p}(k_{j,O_j(l)} - 1))$-th position from the bottom.*

We note that in Proposition 2, K_j's are only used to define lower bounds for optimistic agents, and $k_{j,\mathcal{O}_j(i)}$ for all $i < K_j$ are only used to define lower bounds if agent j is pessimistic.

Our main theorem in this section states that all these lower bounds can be achieved in one preference profile. Moreover, for the same profile there exists an allocation where almost all agents get their top-ranked bundle (and the only person who may not get her top-ranked bundle gets her second-ranked bundle).

Theorem 2. *For any sequential allocation $\mathcal{O}$, and each agent is fixed to be optimistic or pessimistic, there exists a preference profile P such that:*

1. *an optimistic agent j's allocation is ranked at the $\prod_{l=K_j}^{p} k_{j,O_j(l)}$-th position from the bottom;*
2. *a pessimistic agent j's allocation is ranked at the $(1 + \sum_{l=1}^{p}(k_{j,O_j(l)} - 1))$-th position from the bottom;*
3. *there exists an allocation where at least $n - 1$ agents get their top-ranked bundles, and the remaining agent is allocated her top-ranked or second-ranked bundle. Moreover, if the first agent in $\mathcal{O}$ is pessimistic, then there exists an allocation where all agents get their top-ranked bundle.*

We emphasize that in the theorem, whether an agent is optimistic or pessimistic is fixed before we construct the preference profile.

Example 3. *Let $\mathcal{O}$ be the order in Example 1 and suppose agent 1 and 2 are optimistic while agent 3 is pessimistic.*

By Theorem 2 and Example 2, there exists a profile P such that after applying the sequential allocation $\mathcal{O}$, the bundle agent 1 receives is ranked at $k_{1,2} = 1$the last position; the bundle agent 2 receives is ranked at the last position; and the *bundle agent 3 receives is ranked at the 3rd position from the bundle. Moreover, there exists an allocation agent 2 and 3 get their top-ranked bundles and agent 1 gets her second-ranked bundle. In fact, the preference profile described in Example 1 satisfies all these conditions.*

Future work

We feel that the assignment problems with categorized items provide a framework to apply many techniques developed in other fields of combinatorial preference representation and aggregation. In the future we plan to extend our results to other assignment problems with category constraints, for example the assignment problems where agents have indifference preferences, or some category contains more/less than n items, and the agents can get more/less than 1 item from each category. We have already mentioned that the structure of category constraints allow us to explore other preference representation languages that were not suitable for non-categorized domains, especially CP-nets [2], LP-trees [1], and soft constraints [17]. We also plan to analyze the behavior and effect of strategic agents or regret-minimizing agents. We can also study probabilistic assignment mechanisms.

References

[1] Richard Booth, Yann Chevaleyre, Jérôme Lang, Jérôme Mengin, and Chattrakul Sombattheera. Learning conditionally lexicographic preference relations. In *Proceeding of the 2010 conference on ECAI 2010: 19th European Conference on Artificial Intelligence*, pages 269–274, Amsterdam, The Netherlands, The Netherlands, 2010.

[2] 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.

[3] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Conditional importance networks: A graphical language for representing ordinal, monotonic preferences over sets of goods. In *Proceedings of the 21st international jont conference on Artifical intelligence*, IJCAI'09, pages 67–72, Pasadena, California, USA, 2009.

[4] 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.

[5] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78, Barcelona, Catalonia, Spain, 2011.

[6] 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.

[7] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Multiagent resource allocation in k-additive domains: Preference representation and complexity. *Annals of Operations Research*, 163(1):49–62, 2008.

[8] Lars Ehlers and Bettina Klaus. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice Welfare*, 21:265—280, 2003.

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

[10] Thomas Kalinowski, Nina Narodytska, and Toby Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of IJCAI-2013*, 2013.

[11] Thomas Kalinowski, Nina Narodytska, Toby Walsh, and Lirong Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of AAAI-2013*, 2013.

[12] Bettina Klaus and Eiichi Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30(3):421–435, 2002.

[13] Jérôme Lang and Lirong Xia. Sequential composition of voting rules in multi-issue domains. *Mathematical Social Sciences*, 57(3):304–324, 2009.

[14] Szilvia Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433, 2000.

[15] Szilvia Pápai. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5:91–105, 2000.

[16] Szilvia Pápai. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory*, 3(3):257–71, 2001.

[17] Giorgio Dalla Pozza, Maria Silvia Pini, Francesca Rossi, and K. Brent Venable. Multi-agent soft constraint aggregation via sequential voting. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, pages 172–177, Barcelona, Catalonia, Spain, 2011.

[18] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.