Allocating Indivisible Items in Categorized Domains

Erika Mackin
Computer Science Department
Rensselaer Polytechnic Institute
Troy, NY, USA
mackie2@rpi.edu

Lirong Xia
Computer Science Department
Rensselaer Polytechnic Institute
Troy, NY, USA
xial@cs.rpi.edu

Abstract

We initiate a research agenda of mechanism design for *categorized domain allocation problems (CDAPs)*, where indivisible items from multiple categories are allocated to agents without monetary transfer and each agent gets at least one item per category.

We focus on *basic CDAPs*, where each agent gets exactly one item per category. We first characterize serial dictatorships by a minimal set of three axiomatic properties: strategy-proofness, non-bossiness, and category-wise neutrality. Then, we propose a natural extension of serial dictatorships called *categorial sequential allocation mechanisms (CSAMs)*, which allocate the items in multiple rounds: in each round, the designated agent chooses an item from a designated category. We fully characterize the worst-case *rank efficiency* of CSAMs for optimistic and pessimistic agents.

1 Introduction

Suppose we are organizing a seminar and must allocate 10 discussion topics and 10 dates to 10 students. Students have heterogeneous and combinatorial preferences over (topic, date) bundles. For example, a student may say "I would prefer an early date if I get an easy topic, but I would prefer a late date if I get a hard topic". How should we allocate the topics and dates to students based on their preferences?

This example illustrates a common situation of allocating multiple indivisible items, which we formulate as a *categorized domain*. A categorized domain contains multiple indivisible items, each of which belongs to one of the $p \geq 1$ categories. In *categorized domain allocation problems (CDAPs)*, we want to design a mechanism to allocate the categorized items to agents without monetary transfer, such that each agent gets at least one item per category. In the seminar-organization example there are two categories: topics and dates, and each agent (student) must get one topic and one date.

Many other allocation problems are CDAPs. For example, in cloud computing, agents have heterogeneous preferences over multiple types of items including CPU, memory, and storage [14, 15, 1]; patients must be allocated multiple types of resources including surgeons, nurses, rooms, and equipments [17]; agents may want to reallocate houses and cars [22, 18]; college students compete for course slots from multiple categories, e.g. computer science courses, math courses, social science courses, etc.

The design and analysis of allocation mechanisms for classical *non-categorized* domains have been an active research area at the interface of computer science and economics. In computer science, allocation problems have been studied as *multi-agent resource allocation* [12]. In economics, allocation problems have been studied as *one-sided matching*, also known as *assignment problems* [28]. Previous research faces three main barriers.

• *Preference bottleneck:* When the number of items is not small, it is impractical for the agents to fully express their preferences over all (exponential) bundles of items.

• *Computational bottleneck:* Even if the agents can express their preferences compactly using some preference language, computing an "optimal" allocation is often hard.

• *Threats of agents' strategic behavior:* An agent may have incentive to report untruthfully to obtain a more preferred bundle. This may lead to a socially inefficient allocation.

Our Contributions. We initiate a research agenda of mechanism design for CDAPs towards breaking the three aforementioned barriers. CDAPs naturally generalize classical non-categorized allocation problem (which are CDAPs with one category). CDAPs are our main conceptual contribution.

As a first step, we focus on *basic categorized domain allocation problems (basic CDAPs)*, where the number of items in each category is exactly the same as the number of agents, so that each agent gets exactly one item from each category. The seminar-organization example is a basic CDAP.

Our technical contributions are three-fold. First, we characterize *serial dictatorships* for any basic CDAPs with at least two categories by a minimal set of three axiomatic properties:

• *strategy-proofness*, which states that no agent has incentive to misreport her preferences;

• *non-bossiness*, which states that no agent can change the allocation of any other agent without changing her own by reporting differently;

• *category-wise neutrality*, which states that after permuting the names of items in any category in all agents' preferences, the outcome allocation is permuted in the same way.

This characterization helps us understand the possibility of designing strategy-proof mechanisms to overcome the third barrier, i.e. threats of agents' strategic behavior.

Second, to overcome the preference bottleneck and the computational bottleneck, and to go beyond serial dictatorships, we propose *categorial sequential allocation mechanisms (CSAMs)*, which are a large class of indirect mechanisms that naturally extend serial dictatorships [29], sequential allocation protocols [5, 10, 11], and the draft mechanism [9]. For n agents and p categories, a CSAM is characterized by an ordering over all (agent, category) pairs: in each round, the designated agent picks an available item from the designated category.

Third, we completely characterize the worst-case *rank efficiency* of CSAMs for any combination of two types of *myopic* agents: *optimistic* agents, who always choose the item in their top-ranked bundle that is still available, and *pessimistic* agents, who always choose the item that have the best worst-case guarantee. This characterization naturally leads to useful corollaries that help us choose the optimal CSAMs w.r.t. their worst-case rank efficiency. For example, we show that while serial dictatorships with all-optimistic agents have the best worst-case utilitarian rank, they have the worst worst-case egalitarian rank (Proposition 1).

A previous version of this paper was presented at ISAIM-14 special session on Computational Social Choice [30]. **The full version with all missing proofs can be found on arXiv.**

Related Work and Discussions. We are not aware of previous work that explicitly formulates CDAPs. Previous work on multi-type resource allocation often assume that items of the same type are interchangeable, and agents have a specific type of preferences, e.g. *Leontief preferences* [14] or *threshold preferences* [17]. CDAPs are different because agents' preferences are only required to be rankings but not otherwise restricted.

From the modeling perspective, if we ignore the categorial information then CDAPs become standard multi-agent resource allocation problems. However, the categorial information can be very useful in designing natural allocation mechanisms. For example, CSAMs are not even well-defined without the categorial information. As another example, agents can naturally use graphical languages such *CP-nets* [3] to represent their preferences. Then natural solutions for *combinatorial voting* [8] such as sequential voting can be modified to allocate items for CDAPs.

Technically, one-sided matching problems are basic CDAPs with one category. Our characterization of serial dictatorships for basic CDAPs are different from characterizations of serial dictatorships and similar mechanisms for one-sided matching [29, 23, 24, 25, 13, 16]. This is because the category-wise neutrality used in our characterization is different from (and arguably weaker than) the neutrality used in previous work.

Our analysis of the worst-case rank efficiency of categorial sequential allocation mechanisms resembles the *price of anarchy* [19], which is defined for strategic agents together with a social welfare function that numerically evaluates the quality of outcomes. Our theorem is also related to *distortion* in the voting setting [27, 4], which concerns the social welfare loss caused by agents reporting a ranking instead of a utility function. Nevertheless, our approach is significantly different because we focus on allocation problems for myopic and strategic agents, and we do not assume the existence of agents' cardinal preferences nor a social welfare function, even though our theorem can be easily extended to study worst-case social welfare loss given a social welfare function as in Proposition 1.

Finally, we are certainly not the first to study (optimistically or pessimistically) myopic agents in social choice [6, 20, 7, 21]. However, previous work focused on voting or cake cutting, while we focus on allocation of indivisible items. Analyzing outcomes for other types of agents, especially strategic agents, is a promising direction for future research.

2 Categorized Domain Allocation Problems

Definition 1 *A* categorized domain *is composed of $p \geq 1$ categories of indivisible items, denoted by $\{D_1, \ldots, D_p\}$. In a* categorized domain allocation problem (CDAP)*, the items are allocated to n agents without monetary transfer, such that each agent gets at least one item from each category.*

In a basic categorized domain *for n agents, for each $i \leq p$, $|D_i| = n$, $\mathfrak{D} = D_1 \times \cdots \times D_p$, and each agent's preferences are represented by a linear order over $\mathfrak{D}$. In a* basic categorized domain allocation problem (basic CDAP)*, each agent gets exactly one item from each category.*

In this paper, we focus on basic categorized domains and basic CDAPs for *non-sharable* items [12], that is, each item can only be allocated to one agent. For any $i \leq p$, we denote $D_i = \{1, \ldots, n\}$. Each element in $\mathfrak{D}$ is called a *bundle*. For any $j \leq n$, let R_j denote a linear order over $\mathfrak{D}$ and let $P = (R_1, \ldots, R_n)$ denote the agents' *(preference) profile*. For any pair of bundles $\vec{d}, \vec{e} \in \mathfrak{D}$, we use $\vec{d} \succ_{R_j} \vec{e}$ to denote that $\vec{d}$ is ranked higher than $\vec{e}$ by agent j. An *allocation* A is a mapping from $\{1, \ldots, n\}$ to $\mathfrak{D}$, such that $\bigcup_{j=1}^{n} [A(j)]_i = D_i$, where for any $j \leq n$ and any $i \leq p$, $A(j)$ is the bundle allocated to agent j and $[A(j)]_i$ is the item in category i allocated to agent j. An *allocation mechanism* f is a mapping that takes a profile as input, and outputs an allocation. We use $f^j(P)$ to denote the bundle allocated to agent j by f for profile P.

We now define three axioms for allocation mechanisms. The first two are common in the literature [29].

• A direct mechanism f satisfies *strategy-proofness* if no agent benefits from misreporting her preferences. That is, for any profile P, any agent j, and any linear order R'_j over $\mathfrak{D}$, $f^j(P) \succ_{R_j} f^j(R'_j, R_{-j})$, where R_{-j} is composed of preferences of all agents except agent j.

• f satisfies *non-bossiness* if no agent is *bossy*. An agent is bossy if she can misreport her preferences to change the allocation of some other agent without changing her own. That is, f is non-bossy if for any profile P, any agent j, and any linear order R'_j over $\mathfrak{D}$, $[f^j(P) = f^j(R'_j, R_{-j})] \Rightarrow [f(P) = f(R'_j, R_{-j})]$.

• f satisfies *category-wise neutrality* if after applying a permutation over the items in any given category, the allocation is also permuted in the same way. That is, for any profile P, any category i, and any permutation M_i over D_i, we have $f(M_i(P)) = M_i(f(P))$, where for any bundle $\vec{d} \in \mathfrak{D}$, $M_i(\vec{d}) = (M_i([\vec{d}]_i), [\vec{d}]_{-i})$.

When there is only one category, category-wise neutrality degenerates to the traditional neutrality for one-sided matching [29]. When $p \geq 2$, category-wise neutrality is weaker than the traditional neutrality. This is because neutrality requires that the allocation is insensitive to *all* permutations over items (bundles), while category-wise neutrality only requires such insensitivity for a specific class of permutations (the permutations that can be decomposed into multiple category-wise permutations).

A *serial dictatorship* is characterized by a linear order $\mathcal{K}$ over $\{1, \ldots, n\}$ such that agents choose items in turns according to $\mathcal{K}$. In each step, the designated agent chooses her top-ranked bundle that is still available.

Example 1 Let $n = 3$ and $p = 2$. $\mathfrak{D} = \{1, 2, 3\} \times \{1, 2, 3\}$. Agents' preferences are as follows.

$$R_1 = [12 \succ 21 \succ 32 \succ 33 \succ 31 \succ 22 \succ 23 \succ 13 \succ 11]$$
$$R_2 = [32 \succ 12 \succ 21 \succ 13 \succ 33 \succ 11 \succ 31 \succ 23 \succ 22]$$
$$R_3 = [13 \succ 12 \succ 11 \succ 22 \succ 32 \succ 21 \succ 33 \succ 31 \succ 23]$$

Let $\mathcal{K} = [1 \rhd 2 \rhd 3]$. In the first round of the serial dictatorship, agent 1 chooses 12; in the second round, agent 2 cannot choose 32 or 12 because item 2 in D_2 is unavailable, so she chooses 21; in the final round, agent 3 chooses 33. □

3 An Axiomatic Characterization of Serial Dictatorships

Theorem 1 *For any basic CDAP with $p \geq 2$ 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, the three axioms are minimal for characterizing serial dictatorships.*

Proof sketch: We first present four lemmas that will be frequently used in the proof. The first three lemmas are standard whose proofs are omitted. The last one (Lemma 4) is new, whose proof uses new techniques involving the categorial structure.

The first lemma (roughly) says that for all strategy-proof and non-bossy mechanisms f and all profiles P, if every agent j reports a different ranking without enlarging the set of bundles ranked above $f^j(P)$, then the allocation to all agents does not change in the new profile. This resembles *(strong) monotonicity* in voting.

Lemma 1 *Let f be a strategy-proof and non-bossy allocation mechanism over a basic categorized domain with $p \geq 2$. For any pair of profiles P and P' such that for all $j \leq n$, $\{\vec{d} \in \mathfrak{D} : \vec{d} \succ_{R'_j} f^j(P)\} \subseteq \{\vec{d} \in \mathfrak{D} : \vec{d} \succ_{R_j} f^j(P)\}$, we have $f(P') = f(P)$.*

For any linear order R over $\mathfrak{D}$ and any bundle $\vec{d} \in \mathfrak{D}$, we say a linear order 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. The next lemma states that for any strategy-proof and non-bossy mechanism f, if an agent reports her preferences differently by only pushing up a bundle $\vec{d}$, then either the allocation to all agents does not change, or she gets $\vec{d}$.

Lemma 2 *Let f be a strategy-proof and non-bossy allocation mechanism over a basic categorized domain with $p \geq 2$. 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) $f^j(R'_j, R_{-j}) = \vec{d}$.*

The next lemma states that strategy-proofness, non-bossiness, and category-wise neutrality altogether imply *Pareto-optimality*, which states that for any profile P, there does not exist an allocation A where (1) all agents prefer their bundles in A to their bundles in $f(P)$, and (2) some agents strictly prefer their bundles in A.

Lemma 3 *For any basic categorized domains with $p \geq 2$, any strategy-proof, non-bossy, and category-wise neutral allocation mechanism is Pareto optimal.*

The fourth lemma states that for any strategy-proof and non-bossy allocation mechanism f, any profile P, and any pair of agents j_1, j_2, there is no bundle $\vec{c}$ that is contained in the items allocated to j_1, j_2 by f such that both j_1 and j_2 prefer $\vec{c}$ to their bundles allocated by f, respectively.

Lemma 4 *Let f be a strategy-proof and non-bossy allocation mechanism over a basic categorized domain with $p \geq 2$ and $n \geq 2$. For any 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 does not exist $\vec{c} \in \{a_1, b_1\} \times \{a_2, b_2\} \times \cdots \times \{a_p, b_p\}$ such that $\vec{c} \succ_{R_{j_1}} \vec{a}$ and $\vec{c} \succ_{R_{j_2}} \vec{b}$, where a_i is the i-th component of $\vec{a}$.*

Proof: Suppose for the sake of contradiction that such a bundle $\vec{c}$ exists for some P, j_1, and j_2. Let $\vec{d}$ denote the bundle such that $\vec{c} \cup \vec{d} = \vec{a} \cup \vec{b}$. More precisely, for all $i \leq p$, $\{c_i, d_i\} = \{a_i, b_i\}$. For example, if $\vec{a} = 1213$, $\vec{b} = 2431$, and $\vec{c} = 1211$, then $\vec{d} = 2433$.

The rest of the proof derives a contradiction by proving a series of observations illustrated in Table 1. In each step, we prove that the boxed bundles are allocated to agent j_1 and agent j_2 respectively, and all other agents get their top-ranked bundles.

Step 1. Let $\hat{R}_{j_1} = [\vec{c} \succ \vec{a} \succ \vec{d} \succ \vec{b} \succ \text{others}]$, $\hat{R}_{j_2} = [\vec{c} \succ \vec{b} \succ \vec{a} \succ \vec{d} \succ \text{others}]$, where "others" represents an arbitrary linear order over the remaining bundles, and for any $j \neq j_1, j_2$, let $\hat{R}_j = [f^j(P) \succ \text{others}]$. By Lemma 1, $f(\hat{P}) = f(P)$.

Step 2. Let $\bar{R}_{j_2} = [\vec{c} \succ \vec{a} \succ \vec{b} \succ \vec{d} \succ \text{others}]$ be a pushup of $\vec{a}$ from $\hat{R}_{j_2}$. We will prove that $f(\bar{R}_{j_2}, \hat{R}_{-j_2}) = f(\hat{P}) = f(P)$. Since $\bar{R}_{j_2}$ is a pushup of $\vec{a}$ from $\hat{R}_{j_2}$, by Lemma 2, $f^{j_2}(\bar{R}_{j_2}, \hat{R}_{-j_2})$ is either $\vec{a}$ or $\vec{b}$. We now show that the former case is impossible. Suppose for the sake of contradiction $f^{j_2}(\bar{R}_{j_2}, \hat{R}_{-j_2}) = \vec{a}$, then $f^{j_1}(\bar{R}_{j_2}, \hat{R}_{-j_2})$ cannot be $\vec{c}$, $\vec{a}$, or $\vec{d}$ since otherwise some item will be allocated twice. This

Table 1: The 6 steps in the proof for Lemma 4.

Step 1	Step 2	Step 3
$\hat{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others $\hat{R}_{j_2}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others Other j : $f^j(P) \succ$ others	$\hat{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others Other j : $f^j(P) \succ$ others	$\bar{R}_{j_1}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{b} \succ \vec{d} \succ$ others Other j : $f^j(P) \succ$ others

Step 4	Step 5	Step 6
$\bar{R}_{j_1}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others $\mathring{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others Other j : $f^j(P) \succ$ others	$\mathring{R}_{j_1}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others $\mathring{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others Other j : $f^j(P) \succ$ others	$\mathring{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{b} \succ \vec{d} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others Other j : $f^j(P) \succ$ others

means that $f(\bar{R}_{j_2}, \hat{R}_{-j_2})$ is Pareto dominated by the allocation where j_1 gets $\vec{d}$, j_2 gets $\vec{c}$, and all other agents get their top-ranked bundles. This contradicts the Pareto-optimality of f (Lemma 3). Hence $f^{j_2}(\bar{R}_{j_2}, \hat{R}_{-j_2}) = \vec{b} = f^{j_2}(\hat{P})$. By non-bossiness we have $f(\bar{R}_{j_2}, \hat{R}_{-j_2}) = f(\hat{P}) = f(P)$.

Step 3. Let $\bar{R}_{j_1} = [\vec{c} \succ \vec{b} \succ \vec{a} \succ \vec{d} \succ \text{others}]$ be a pushup of $\vec{b}$ from $\hat{R}_{j_1}$. We will prove that in $f(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$, j_1 gets $\vec{b}$, j_2 gets $\vec{a}$, and all other agents get the same items as in $f(P)$. Since $\bar{R}_{j_1}$ is a pushup of $\vec{b}$ from $\hat{R}_{j_1}$, by Lemma 2, $f^{j_1}(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$ is either $\vec{a}$ or $\vec{b}$. We now show that the former case is impossible. Suppose for the sake of contradiction that $f^{j_1}(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = \vec{a}$. By non-bossiness, $f^{j_2}(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = \vec{b}$. This means that $f(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$ is Pareto-dominated by the allocation where j_1 gets $\vec{b}$, j_2 gets $\vec{a}$, and all other agents get their top-ranked bundles. This contradicts the Pareto-optimality of f (Lemma 3).

Step 4. Let $\mathring{R}_{j_2} = [\vec{c} \succ \vec{a} \succ \vec{d} \succ \vec{b} \succ \text{others}]$ be a pushup of $\vec{d}$ from $\bar{R}_{j_2}$. By Lemma 1, $f(\bar{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = f(\bar{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$.

Step 5. Let $\mathring{R}_{j_1} = [\vec{c} \succ \vec{a} \succ \vec{b} \succ \vec{d} \succ \text{others}]$ be a pushup of $\vec{a}$ from $\bar{R}_{j_1}$. We will prove that $f(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = f(\bar{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$. Since $\mathring{R}_{j_1}$ is a pushup of $\vec{a}$ from $\bar{R}_{j_1}$, by Lemma 2, $f^{j_1}(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$ is either $\vec{a}$ or $\vec{b}$. We now show that the former case is impossible. Suppose for the sake of contradiction that $f^{j_1}(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = \vec{a}$. Then in $f(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$, agent j_2 cannot get $\vec{c}$, $\vec{a}$, or $\vec{d}$, which means that $f(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$ is Pareto-dominated by the allocation where j_1 gets $\vec{c}$, j_2 gets $\vec{d}$, and all other agents get their top-ranked bundles. This contradicts the Pareto-optimality of f. Hence, $f^{j_1}(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = \vec{b}$. By non-bossiness $f(\mathring{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = f(\bar{R}_{j_1}, \mathring{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$.

Step 6. We note that $\mathring{R}_{j_1}$ is a pushup of $\vec{b}$ from $\hat{R}_{j_1}$ (and $\vec{b}$ is still below $\vec{a}$). By Lemma 1, $f(\mathring{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}}) = f(\hat{R}_{j_1}, \bar{R}_{j_2}, \hat{R}_{-\{j_1,j_2\}})$. We note that the right hand side is the profile in Step 2.

Contradiction. Finally, the observations in Step 5 and Step 6 imply that when agents' preferences are $\mathring{R}_{j_1}$ and $\bar{R}_{j_2}$ as in Step 6, agent j_2 has incentive to report $\mathring{R}_{j_2}$ as in Step 5 to improve the bundle allocated to her (from $\vec{b}$ to $\vec{a}$). This contradicts the strategy-proofness of f and completes the proof of Lemma 4. ■

It is easy to check that any serial dictatorship satisfies strategy-proofness, non-bossiness and category-wise neutrality. We now prove that any mechanism f satisfying the three axioms must be a serial dictatorship. Let R^* be a linear order over $\mathfrak{D}$ that satisfies the following conditions:

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

The next claim states that f agrees with a serial dictatorship on the profile $(R^*, \ldots, R^*)$ where all agents have the same preferences R^* that we have just defined. We will later show that f agrees with the same serial dictatorship on all profiles.

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

Proof: The claim is proved by induction on l. Let $l = 1$, for the sake of contradiction suppose there is no j_1 with $f^{j_1}(P^*) = (1, \ldots, 1)$. Then there exist a pair of agents j and j' such that both $\vec{a} = f^j(P^*)$ and $\vec{b} = f^{j'}(P^*)$ contain 1 in at least one category.

Let $\vec{c}$ be the bundle obtained from $\vec{a}$ by replacing items in categories where $\vec{b}$ takes 1 to 1. More precisely, we let $\vec{c} = (c_1, \ldots, c_p)$, where $c_i = \begin{cases} 1 & \text{if } a_i = 1 \text{ or } b_i = 1 \\ a_i & \text{otherwise} \end{cases}$.

It follows that $\vec{c} \succ_{R^*} \vec{a}$ and $\vec{c} \succ_{R^*} \vec{b}$ because the number of 1's in $\vec{c}$ is strictly larger than the number of 1's in $\vec{a}$ and $\vec{b}$. By Lemma 4, this contradicts the assumption that f is strategy-proof and non-bossy. Hence there exists $j_1 \leq n$ with $f^{j_1}(P^*) = (1, \ldots, 1)$.

Suppose that 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 after a similar reasoning to 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 $f(P^*)$ such that both $\vec{a}$ and $\vec{b}$ contain $l'+1$ in at least one category. By the induction hypothesis, items $\{1,\ldots,l'\}$ in all categories have been allocated, which means that all components of $\vec{a}$ and $\vec{b}$ are at least as large as $l'+1$. Let $\vec{c}$ be the bundle obtained from $\vec{a}$ by replacing items in all categories where $\vec{b}$ takes $l'+1$ to $l'+1$. We have $\vec{c} \succ_{R^*} \vec{a}$ and $\vec{c} \succ_{R^*} \vec{b}$, leading to a contradiction by Lemma 4. Therefore, the claim holds for $l = l'+1$. This completes the proof of Claim 1. ■

W.l.o.g. let $j_1 = 1$, $j_2 = 2$, ..., $j_n = n$ denote the agents in Claim 1. We next show that for all profiles, f agrees with the serial dictatorship $1 \rhd 2 \rhd \cdots \rhd n$. For any profile $P' = (R'_1,\ldots,R'_n)$, we define n bundles as follows. Let $\vec{d^1}$ denote the top-ranked bundle in R'_1. For any $l \geq 2$, let $\vec{d^l}$ denote agent l's top-ranked available bundle supposing that items in $\vec{d^1},\ldots,\vec{d^{l-1}}$ have already been allocated. That is, $\vec{d^l}$ is the most preferred bundle in $\{\vec{d} : \forall l' < l, \vec{d} \cap \vec{d^{l'}} = \emptyset\}$ according to R'_l. In other words, $\vec{d^1},\ldots,\vec{d^n}$ are the bundles allocated to agents 1 through n by the serial dictatorship $1 \rhd 2 \cdots \rhd n$. We next prove that this is exactly the allocation by f.

For any $i \leq m$, we define a category-wise permutation M_i such that for all $l \leq n$, $M_i(l) = [\vec{d^l}]_i$, where we recall that $[\vec{d^l}]_i$ is the item in the i-th category in $\vec{d^l}$. Let $M = (M_1,\ldots,M_m)$. It follows that for all $l \leq n$, $M(l,\ldots,l) = \vec{d^l}$. By category-wise neutrality and Claim 1, in $f(M(P^*))$ agent l gets $M(f^l(P^*)) = \vec{d^l}$.

Comparing $M(P^*)$ to P', we notice that for all $l \leq n$ and all bundles $\vec{e}$, if $\vec{d^l} \succ_{M(R^*)} \vec{e}$ then $\vec{d^l} \succ_{R'_l} \vec{e}$. This is because if there exists $\vec{e}$ such that $\vec{d^l} \succ_{M(R^*)} \vec{e}$ but $\vec{e} \succ_{R'_l} \vec{d^l}$, then $\vec{e}$ is still available after $\{\vec{d^1},\ldots,\vec{d^{l-1}}\}$ have been allocated, and $\vec{e}$ is ranked higher than $\vec{d^l}$ in R'_l. This contradicts the selection of $\vec{d^l}$. By Lemma 1, $f(P') = f(M(P^*)) = M(f(P^*)) = (\vec{d^1},\ldots,\vec{d^n})$. We recall that by definition $(\vec{d^1},\ldots,\vec{d^n})$ is the output of the serial dictatorship $1 \rhd 2 \cdots \rhd n$. This proves that f is the serial dictatorship w.r.t. the order $1 \rhd 2 \rhd \cdots \rhd n$. ■

Remarks. The theorem is somewhat negative because it shows that we have to sacrifice one of strategy-proofness, category-wise neutrality, or non-bossiness. Among the three axiomatic properties, we feel that non-bossiness is the least natural one.

4 Categorial Sequential Allocation Mechanisms

Given a linear order $\mathcal{O}$ over $\{1,\ldots,n\} \times \{1,\ldots,p\}$, the *categorial sequential allocation mechanism (CSAM)* $f_{\mathcal{O}}$ allocates the items in np steps as illustrated in Protocol 1. In each step t, suppose the t-th element in $\mathcal{O}$ is (j,i), (equivalently, $t = \mathcal{O}^{-1}(j,i)$). Agent j is called the *active agent* in step t and she chooses an item $d_{j,i}$ that is still available from D_i. Then, $d_{j,i}$ is broadcast to all agents and we move on to the next step.

Protocol 1: Categorial sequential allocation mechanism (CSAM) $f_{\mathcal{O}}$.

Input: An order $\mathcal{O}$ over $\{1,\ldots,n\} \times \{1,\ldots,p\}$.
Broadcast $\mathcal{O}$ to all agents.[1]
for $t = 1$ *to* np **do**
  Let (j,i) be the t-th element in $\mathcal{O}$.
  Agent j chooses an available item $d_{j,i} \in D_i$.
  Broadcast $d_{j,i}$ to all agents.
end

CSAMs are different from sequential allocation protocols [5] and the draft mechanism [9], where in each step the active agent can choose any available item from any category.

Example 2 The serial dictatorship w.r.t. $\mathcal{K} = [j_1 \rhd \cdots \rhd j_n]$ is a CSAM w.r.t. $(j_1,1) \rhd (j_1,2) \rhd \cdots \rhd (j_1,p) \rhd \cdots \rhd (j_n,1) \rhd (j_n,2) \rhd \cdots \rhd (j_n,p)$. □

Similar to sequential allocation protocols [5], CSAMs can be implemented in a distributed way. Communication cost for CSAMs is much lower than that for direct mechanisms, where agents report their preferences in full to the center, which requires $\Theta(n^p p \log n)$ bits per agent, and thus the total communication cost is $\Theta(n^{p+1} p \log n)$. For CSAMs, the total communication cost of Protocol 1 is $\Theta(n^2 p \log n + np(n \log n)) = \Theta(n^2 p \log np)$, which has $\Theta(n^{p-2} \cdot \frac{\log n}{\log n + \log p})$ multiplicative saving. In light of this, CSAMs preserve more privacy as well.

To analyze the outcome and efficiency of CSAMs, we consider two types of myopic agents. For any $1 \leq i \leq p$, we let $D_{i,t}$ denote the set of available items in D_i at the beginning of round t.

- *Optimistic* agents. An optimistic agent chooses the item in her top-ranked bundle that is still available.
- *Pessimistic* agents. A pessimistic agent j in round t chooses an item $d_{j,i}$ from $D_{i,t}$, such that for all $d'_i \in D_{i,t}$ with $d'_i \neq d_{j,i}$, agent j prefers the worst available bundle whose i-th component is $d_{j,i}$ to the worst available bundle whose i-th component is d'_i.

In this paper, we assume that whether an agent is optimistic or pessimistic is fixed before applying a CSAM.

Example 3 Let $n = 3$, $p = 2$. Consider the same profile as in Example 1, which can be simplified as follows.

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 $t = 1$, agent 1 (optimistic) chooses item 1 from D_1. When $t = 2$, item 32 is the top-ranked available bundle for agent 2 (optimistic), so she chooses 2 from D_2. When $t = 3$, the available bundles are $\{2,3\} \times \{1,3\}$. If agent 3 chooses 2 from D_1, then the worst-case available bundle is 23, and if agent 3 chooses 3 from D_1, then the worst-case

[1] As one of the anonymous reviewers suggested, the order $\mathcal{O}$ does not need to be broadcast in the beginning. This gives a better justification for myopic agents and may also reduce manipulation.

available bundle is 31. Since agent 3 prefers 31 to 23, she chooses 3 from D_1. When $t = 4$, agent 3 chooses 3 from D_2. When $t = 5$, agent 2 choses 2 from D_1 and when $t = 6$, agent 1 choses 1 from D_2. Finally, agent 1 gets 11, agent 2 gets 22, and agent 3 gets 33. □

5 Rank Efficiency of CSAMs for Myopic Agents

For any linear order R over $\mathfrak{D}$ and any bundle $\vec{d}$, we let $\text{Rank}(R, \vec{d})$ denote the rank of $\vec{d}$ in R, such that the highest position has rank 1 and the lowest position has rank n^p. Given a profile $P = (R_1, \ldots, R_n)$ and a mechanism f, the *rank efficiency* of f is a vector $(\text{Rank}(R_1, f^1(P)), \ldots, \text{Rank}(R_n, f^n(P)))$ that is composed of the ranks of bundles agents receive.

We recall that a CSAM $f_{\mathcal{O}}$ is characterized by a linear order over all (agent,category) pairs. Given $f_{\mathcal{O}}$, we introduce the following notation for all $j \leq n$ to characterize the worst-case rank efficiency of $f_{\mathcal{O}}$.

• Let $\mathcal{O}_j$ denote the linear order over the categories $\{1, \ldots, p\}$ according to which agent j chooses items.

• For any $i \leq p$, let $k_{j,i}$ denote the number of items in D_i that are still available right before agent j chooses from D_i. Formally, $k_{j,i} = 1 + |\{(j', i) : (j, i) \rhd_{\mathcal{O}} (j', i)\}|$.

• Let K_j denote the smallest index in $\mathcal{O}_j$ such that no agent can "interrupt" agent j from choosing all items in her top-ranked bundle that is available in round $(j, \mathcal{O}_j(K_j))$. Formally, K_j is the smallest number such that for any l with $K_j < l \leq p$, between the round when agent j chooses an item from category $\mathcal{O}_j(K_j)$ and the round when agent j chooses an item from category $\mathcal{O}_j(l)$, no agent chooses an item from category $\mathcal{O}_j(l)$. We note that K_j is defined only by $\mathcal{O}$ and is thus independent of agents' preferences.

Example 4 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, $f_{\mathcal{O}^*}$ is a serial dictatorship. Then $\mathcal{O}^*_1 = \mathcal{O}^*_2 = \mathcal{O}^*_3 = 1 \rhd 2 \rhd 3$. $K_1 = K_2 = K_3 = 1$. $k_{1,1} = k_{1,2} = k_{1,3} = 3$, $k_{2,1} = k_{2,2} = k_{2,3} = 2$, $k_{3,1} = k_{3,2} = k_{3,3} = 1$.

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

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

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

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

Theorem 2 *For any CSAM $f_{\mathcal{O}}$, any combination of optimistic and pessimistic agents, any $j \leq n$, and any profile:*

- **Upper bound for optimistic agents:** *if j is optimistic, then the rank of the bundle allocated to her is at most $n^p + 1 - \prod_{l=K_j}^{p} k_{j,\mathcal{O}_j(l)}$.*
- **Upper bound for pessimistic agents:** *if j is pessimistic, then the rank of the bundle allocated to her is at most $n^p - \sum_{l=1}^{p}(k_{j,\mathcal{O}_j(l)} - 1)$.*

Moreover, there exists a profile P where the bounds for all agents are simultaneously tight. For the same profile P, there exists an allocation where at least $n - 1$ agents get their top-ranked bundles, and the remaining agent gets her top-ranked or second-ranked bundle.

Proof sketch: For any optimistic agents, the upper bound is calculated by counting how many bundles must be ranked below the final allocation once no agent can interrupt her from choosing the top bundle (among available ones). For pessimistic agents, in each step j we know that at least $k_{j,\mathcal{O}_j(l)} - 1$ other bundles must be ranked below the final allocation. The proof for the matching lower bounds is by construction and is quite involved. Intuitively, the construction ensures that for all agents, the bundles mentioned in the proof for the upper bounds are the only bundles ranked below the final allocation. ■

We note that Theorem 2 works for any combination of optimistic and pessimistic agents, which is much more general than the setting with all-optimistic agents and the setting with all-pessimistic agents.

Theorem 2 can be used to compare various CSAMs with optimistic and pessimistic agents. Given a CSAM $f_{\mathcal{O}}$, the worst-case *utilitarian rank* is the worst (largest) total rank of the bundles (w.r.t. respective agent's preferences) allocated by $f_{\mathcal{O}}$. The worst-case *egalitarian rank* is the worst (largest) rank of the least-satisfied agent. The worst case is taken over all profiles of n agents. Due to the space limit, we only present one proposition without proof.

Proposition 1 *Among all CSAMs, serial dictatorships with all-optimistic agents have the best worst-case utilitarian rank and the worst worst-case egalitarian rank.*

The proposition is proved by applying Theorem 2. For any serial dictatorship with all-optimistic agents, the worst-case utilitarian rank is $n(n^p + 1) - \sum_{j=1}^{n} j^p$ and the worst-case egalitarian rank is n^p (when all agents have the same preferences).

6 Summary and Future Work

We have initiated a research agenda to study item allocation in categorized domains. There are many open questions and directions for future research, including strategic agents and minimax-regret agents. We also plan to work on theoretical analysis of the expected utilitarian rank and egalitarian rank for randomized allocation mechanisms. For general CDAPs, we are excited to explore generalizations of CP-nets [3], LP-trees [2], and soft constraints [26] for preference representation. Based on these new languages we can analyze fairness and computational aspects of CSAMs and other mechanisms. Mechanism design for CDAPs with sharable, non-sharable, and divisible items is also an important and promising topic for future research.

Acknowledgments

This work is supported by the National Science Foundation under grant IIS-1453542. We thank Vincent Conitzer, Jerome Lang, David Parkes, anonymous reviewers of this paper, and participants of ISAIM-14 and the Economics and Computation Program at Simons Institute for the theory of computing for helpful comments and suggestions.

References

[1] Arka A. Bhattacharya, David Culler, Eric Friedman, Ali Ghodsi, Scott Shenker, and Ion Stoica. Hierarchical scheduling for diverse datacenter workloads. In *Proc. SoCC*, pages 4:1–4:15, 2013.

[2] Richard Booth, Yann Chevaleyre, Jérôme Lang, Jérôme Mengin, and Chattrakul Sombattheera. Learning conditionally lexicographic preference relations. In *Proc. ECAI*, pages 269–274, 2010.

[3] 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.

[4] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. In *Proc. ACM EC*, pages 197–214, 2012.

[5] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proc. IJCAI*, pages 73–78, 2011.

[6] Steven J. Brams, D. Marc Kilgour, and William S. Zwicker. The paradox of multiple elections. *Social Choice and Welfare*, 15(2):211–236, 1998.

[7] Steven J. Brams, Michael A. Jones, and Christian Klamler. Better Ways to Cut a Cake. *Notices of the AMS*, 53(11):1314–1321, 2006.

[8] Felix Brandt, Vincent Conitzer, and Ulle Endriss. Computational social choice. In G. Weiss, editor, *Multiagent Systems*. MIT Press, 2013.

[9] Eric Budish and Estelle Cantillon. The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. *American Economic Review*, 102(5):2237–71, 2012.

[10] Katarína Cechlárová, Pavlos Eirinakis, Tamás Fleiner, Dimitrios Magos, Ioannis Mourtos, and Eva Potpinková. Pareto optimality in many-to-many matching problems. In *Discrete Optimization*, volume 14, pages 160–169, 2014.

[11] Katarína Cechlárová, Pavlos Eirinakis, Tamás Fleiner, Dimitrios Magos, David F. Manlove, Ioannis Mourtos, Eva Oceľáková, and Baharak Rastegari. Pareto Optimal Matchings in Many-to-Many Markets with Ties. In *SAGT*, pages 27–39. 2015.

[12] 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.

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

[14] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proc. NSDI*, pages 323–336, 2011.

[15] Ali Ghodsi, Vyas Sekar, Matei Zaharia, and Ion Stoica. Multi-resource Fair Queueing for Packet Processing. In *Proc. ACM SIGCOMM*, volume 42, pages 1–12, 2012.

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

[17] Woonghee Tim Huh, Nan Liu, and Van-Anh Truong. Multiresource Allocation Scheduling in Dynamic Environments. *Manufacturing and Service Operations Management*, 15(2):280–291, 2013.

[18] Hideo Konishi, Thomas Quint, and Jun Wako. On the Shapley–Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics*, 35(1):1—15, 2001.

[19] Elias Koutsoupias and Christos Papadimitriou. Worst-case Equilibria. In *Proc. STACS*, pages 404–413, 1999.

[20] Dean Lacy and Emerson M.S. Niou. A problem with referendums. *Journal of Theoretical Politics*, 12(1):5–31, 2000.

[21] Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A Local-Dominance Theory of Voting Equilibria. In *Proc. ACM EC*, pages 313–330, 2014.

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

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

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

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

[26] Giorgio Dalla Pozza, Maria Silvia Pini, Francesca Rossi, and K. Brent Venable. Multi-agent soft constraint aggregation via sequential voting. In *Proc. IJCAI*, pages 172–177, 2011.

[27] Ariel D. Procaccia and Jeffrey S. Rosenschein. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents*, volume 4149 of *LNAI*, pages 317–331. 2006.

[28] Tayfun Sönmez and M. Utku Ünver. Matching, Allocation, and Exchange of Discrete Resources. In Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, editors, *Handbook of Social Economics*, chapter 17, pages 781–852. North-Holland, 2011.

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

[30] Lirong Xia. Assigning indivisible and categorized items. ISAIM-14 special session on Computational Social Choice, 2014.