Allocating Indivisible Items in Categorized Domains

Paper 972

Abstract

We formulate a general class of allocation problems called *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 the number of items in each category equals to the number of agents. We characterize serial dictatorships for basic CDAPs by a minimal set of three desired properties: strategy-proofness, non-bossiness, and category-wise neutrality. Then, we propose a natural extension of serial dictatorships called *categorical sequential allocation mechanisms (CSAMs)*, which allocate the items in multiple rounds: in each round, the active agent chooses an item from a designated category. We fully characterize the worst-case ordinal efficiency of CSAMs for optimistic and pessimistic agents. We believe that these constitute a promising first step towards theoretical foundations and applications of general CDAPs.

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: their preferences over the topics may depend on the date and vice versa, because she may prefer an early date if she gets an easy topic and may prefer a late date if she gets a hard topic.

This example illustrates a common setting for 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 items to agents without monetary transfer, such that each agent gets at least one item from each category. In the above example, there are two categories: topics and dates, and each agent (student) must get a topic and a 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[1] [15, 14, 1]; patients must be allocated multiple types of resources including surgeons, nurses, rooms, and equipments [17]; college students need to choose courses from multiple categories per semester, e.g. computer science courses, math courses, social science courses, etc.

The design and analysis of allocation mechanisms for 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* [25]. Previous research faces three main barriers.

- *Preference bottleneck:* When the number of items is not too small, it is impractical for the agents to 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 a hard combinatorial optimization problem.
- *Threats of agents' strategic behavior:* An agent may have incentive to report untruthfully to get a more preferred bundle. This may lead to a socially inefficient allocation.

Our Contributions. We initiate the study of mechanism design under the novel framework of CDAPs towards breaking the three aforementioned barriers. CDAPs naturally generalize classical non-categorized allocation problems, 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. See e.g. the seminar-organization example. As we will show, mechanism design for basic CDAPs is already highly non-trivial.

Our technical contributions are two-fold. First, we characterizes *serial dictatorships* for any basic CDAPs

[1]Suppose each type contains discrete units of resources that are essentially indivisible for operational convenience.

with at least two categories by a minimal set of three axiomatic properties: *strategy-proofness*, *non-bossiness*, and *category-wise neutrality*. This 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 *categorical sequential allocation mechanisms (CSAMs)*, which are a large class of indirect mechanisms that naturally extend serial dictatorships [26], sequential allocation protocols [6], and the draft mechanism [11]. For n agents and p categories, a CSAM is defined by an ordering over all (agent, category) pairs: in each round, the active agent picks an item that has not been chosen yet from the designated category. CSAMs have low communication complexity and can be implemented in a distributed manner.

We completely characterize the worst-case ordinal efficiency of CSAMs, measured by agents' *ranks* of the bundles they receive, 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 gives them best worst-case guarantee.[2] This characterization naturally leads to useful corollaries on worst-case efficiency of various CSAMs. 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. On the other hand, balanced CSAMs with all-pessimistic agents have good worst-case utilitarian rank.

Related Work and Discussions. We are not aware of previous work that explicitly formulates CDAPs. Previous work on multi-type resource allocation assumes that items of the same type are interchangeable, and agents have specific preferences, e.g. *Leontief preferences* [15] and threshold preferences [17]. CDAPs are more general as agents' preferences are only required to be rankings but not otherwise restricted.

From the modeling perspective, ignoring the categorical information, CDAPs become standard centralized multi-agent resource allocation problems. However, the categorical information opens more possibilities for designing natural allocation mechanisms such as CSAMs. More importantly, we believe that CDAPs provide a natural framework for cross-fertilization of ideas and techniques from other fields of preference representation and aggregation. For example, the combinatorial structure of categorized domains naturally allows agents to use graphical languages (e.g. *CP-nets* [4]) to represent their preferences, which is otherwise hard [7]. Approaches in *combinatorial voting* [10] can also be naturally considered in CDAPs.

Technically, one-sided matching problems are basic CDAPs with one category. Our characterization of serial dictatorships for basic CDAPs may look similar to characterizations of serial dictatorships and similar mechanisms for one-sided matching [26, 20, 21, 22, 13, 16]. However, our theorem is stronger as the category-wise neutrality used in our characterization is weaker than the neutrality used in previous work.

Our analysis of the worst-case ordinal efficiency of categorical sequential allocation mechanisms resembles the *price of anarchy* [18], which is defined for strategic and self-interested agents, with the presence of a social welfare function that numerically evaluates the quality of outcomes. Our theorem is also related to *distortion* in the voting setting [24, 5], 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 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 2 through 5.

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)*, we want to allocate the items 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)*, we want to allocate the items to n agents without monetary tranfer, such that every agent gets exactly one item in 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. Therefore, for all $i \leq p$, we write $D_i = \{1, \ldots, n\}$. 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*. 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 $i \leq p$, $A(j)$ 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 desired axiomatic properties for allocation mechanisms. The first two properties are common in the literature [26], and the third is new.

• 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

[2] Similar types of agents have been studied in other social choice settings [8, 19, 9].

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 report differently to change the bundles allocated to some other agents without changing her own allocation. That is, 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 a 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 [26]. When $p \geq 2$, category-wise neutrality is much weaker than the traditional neutrality.

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

Example 1 Let $n = 3$ and $p = 2$. $\mathfrak{D} = \{1, 2, 3\} \times \{1, 2, 3\}$. Agents' preferences are as follows.
$R_1 = [12 \succ 21 \succ 32 \succ 33 \succ 31 \succ 22 \succ 23 \succ 13 \succ 11]$
$R_2 = [32 \succ 12 \succ 21 \succ 13 \succ 33 \succ 11 \succ 31 \succ 23 \succ 22]$
$R_3 = [13 \succ 12 \succ 11 \succ 22 \succ 32 \succ 21 \succ 33 \succ 31 \succ 23]$

Suppose the agents are truthful. 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. □

An Axiomatic Characterization

Theorem 1 *For any $p \geq 2$ and $n \geq 2$, an allocation mechanism for basic categorized domain 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: It is easy to check that any serial dictatorship satisfies strategy-proofness, non-bossiness, and category-wise neutrality. We prove the converse by four lemmas. The first three lemmas are standard and the last one (Lemma 4) is novel, whose proof is more involved and heavily depends on the categorical structure. **Due to the space constraint, most proofs are omitted. All missing proofs can be found in the supplementary material.**

The first lemma resembles *strong monotonicity* in voting theory: for all strategy-proof and non-bossy mechanism f and all profile P, if each agent j reports differently without enlarging the set of bundles ranked above $f^j(P)$, then the allocation does not change.

Lemma 1 *Let f be a strategy-proof and non-bossy allocation mechanism. 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 $\vec{d}$ without changing the orders of other bundles. The second lemma states that for any strategy-proof and non-bossy mechanism f, if an agent reports differently by only pushing up a bundle $\vec{d}$, then either the allocation does not change, or she gets $\vec{d}$.

Lemma 2 *Let f be a strategy-proof and non-bossy allocation 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) $f(R'_j, R_{-j}) = f(P)$ or (2) $f^j(R'_j, R_{-j}) = \vec{d}$.*

The third lemma states that strategy-proofness, non-bossiness, and category-wise neutrality altogether imply *Pareto-optimality*, which means that for any profile P, there is no allocation A such that all agents prefer their bundles in A than their bundles in $f(P)$, and some of these preferences are strict.

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 says 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 only contains items allocated to agent j_1 and j_2, and both agents prefer $\vec{c}$ to their allocated bundles respectively.

Lemma 4 *Let f be a strategy-proof and non-bossy allocation mechanism. For any profile P and any $j_1 \neq j_2$, let $\vec{a} = f^{j_1}(P)$ and $\vec{b} = f^{j_2}(P)$, there is no $\vec{c} \in \{a_1, b_1\} \times \{a_2, b_2\} \times \cdots \times \{a_p, b_p\}$ 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 sketch: Suppose for the sake of contradiction that such a bundle $\vec{c}$ exists. Let $\vec{d}$ denote the bundle such that $\vec{c} \cup \vec{d} = \vec{a} \cup \vec{b}$. For example, if $\vec{a} = 1213$, $\vec{b} = 2431$, and $\vec{c} = 1211$, then $\vec{d} = 2433$.

We derive a contradiction in 6 steps 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. The first two steps are shown as an example.

Step 1. Let $\hat{R}_{j_1} = [\vec{c} \succ \vec{a} \succ \vec{d} \succ \vec{b} \succ \text{others}]$, $\hat{R}_{j_2} = [\vec{c} \succ \vec{b} \succ \vec{a} \succ \vec{d} \succ \text{others}]$. 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 $\bar{R}_{j_2} = [\vec{c} \succ \vec{a} \succ \vec{b} \succ \vec{d} \succ \text{others}]$. We will prove that $f(\bar{R}_{j_2}, \hat{R}_{-j_2}) = f(\hat{P}) = f(P)$. Because $\bar{R}_{j_2}$ 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 $\vec{b}$. The former case is impossible, otherwise $f^{j_1}(\bar{R}_{j_2}, \hat{R}_{-j_2})$ cannot be $\vec{c}$, $\vec{a}$, or $\vec{d}$ because otherwise some item will be allocated twice. This 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

$\hat{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others $\hat{R}_{j_2}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others Step 1	$\mathring{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others Step 2	$\bar{R}_{j_1}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{b} \succ \vec{d} \succ$ others Step 3
$\bar{R}_{j_1}: \vec{c} \succ \boxed{\vec{b}} \succ \vec{a} \succ \vec{d} \succ$ others $\mathring{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others Step 4	$\mathring{R}_{j_1}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others $\mathring{R}_{j_2}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{d} \succ \vec{b} \succ$ others Step 5	$\mathring{R}_{j_1}: \vec{c} \succ \boxed{\vec{a}} \succ \vec{b} \succ \vec{d} \succ$ others $\bar{R}_{j_2}: \vec{c} \succ \vec{a} \succ \boxed{\vec{b}} \succ \vec{d} \succ$ others Step 6

Table 1: Proof sketch for Lemma 4. In all steps any other agents j's preferences are $f^j(P) \succ$ others.

$f^{j_2}(\bar{R}_{j_2}, \hat{R}_{-j_2}) = \vec{b} = f^{j_2}(\hat{P})$. By non-bossiness we have $f(\bar{R}_{j_2}, \hat{R}_{-j_2}) = f(\hat{P}) = f(P)$.

Contradiction. The observations in Step 5 and Step 6 (see Table 1) imply that when agents' preferences are as in Step 6, agent j_2 has incentive to report $\mathring{R}_{j_2}$ in Step 5 to improve her allocation from $\vec{b}$ to $\vec{a}$. This contradicts the strategy-proofness of f. ■

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 that satisfy 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.
- For any j and any $\vec{d}, \vec{e} \in B_j$, if there are more j's in $\vec{d}$ than in $\vec{e}$, then $\vec{d} \succ \vec{e}$.

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

The proof of Claim 1 uses Lemma 4. W.l.o.g. let $j_1 = 1$, $j_2 = 2$, ..., $j_n = n$ 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 in R'_1, and for any $l \geq 2$, let $\vec{d^l}$ denote agent l's top-ranked available bundle after $\{\vec{d^1}, \ldots, \vec{d^{l-1}}\}$ have been allocated. Then, 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, $f^l(M(P^*)) = M(f^l(P^*)) = \vec{d^l}$.

Comparing $M(P^*)$ to P', we have that for all $l \leq n$ and all bundle $\vec{e}$, 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, $f(P') = f(M(P^*)) = M(f(P^*))$, which proves that f is the serial dictatorship w.r.t. the order $1 \rhd 2 \rhd \cdots \rhd n$.

Finally, we show the minimality of {strategy-proofness, non-bossiness, category-wise neutrality}.

Strategy-proofness is necessary by considering the allocation mechanism that maximizes the social welfare w.r.t. the following utility functions. For any $i \leq n^p$ and $j \leq n$, the bundle ranked at the i-th position in agent j's preferences gets $(n^p - i)(1 + (\frac{1}{2n^p})^j)$ points.

Non-bossiness is necessary by considering the following "conditional serial dictatorship": agent 1 chooses her favorite bundle in the first p rounds, and if the first component of agent 1's second-ranked bundle is the same as the first component of her top-ranked bundle, then the order over the rest of agents is $2 \rhd 3 \rhd \cdots \rhd n$; otherwise the order is $n \rhd n-1 \rhd \cdots \rhd 2$.

Category-wise neutrality is necessary by considering the following "conditional serial dictatorship": agent 1 chooses her favorite bundle in the first p rounds, and if agent 1 gets $(1, \ldots, 1)$, then the order over the rest of agents is $2 \rhd 3 \rhd \cdots \rhd n$; otherwise the order is $n \rhd n-1 \rhd \cdots \rhd 2$. ■

Categorical Sequential Allocation Mechanisms

Given a linear order $\mathcal{O}$ over $\{1, \ldots, n\} \times \{1, \ldots, p\}$, the *categorical 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: Categorical 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.
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

In CSAMs, in each step the active agent must choose an item from the designated category. Hence, CSAMs are different from sequential allocation protocols [6] and the draft mechanism [11], 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. $(j_1, 1) \rhd (j_1, 2) \rhd \cdots \rhd (j_1, p) \rhd \cdots \rhd (j_n, 1) \rhd (j_n, 2) \rhd \cdots \rhd (j_n, p)$.

For any even number p, given any linear order $\mathcal{K} = [j_1 \rhd \cdots \rhd j_n]$ over the agents, we define the *balanced CSAM* to be the mechanism where agents choose items in p phases, such that for each $i \leq p$, in phase i all agents choose from D_i w.r.t. $\mathcal{K}$ if i is odd, and w.r.t. inverse K if i is even.

For example, when $n = 3$, $p = 2$, and $\mathcal{K} = [1 \rhd 2 \rhd 3]$, the balanced CSAM uses the order $(1, 1) \rhd (2, 1) \rhd (3, 1) \rhd (3, 2) \rhd (2, 2) \rhd (1, 2)$. □

Similar to sequential allocations [6], CSAMs can be implemented in a distributed manner. Communication cost for CSAMs is much lower than 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 a $\Theta(n^{p-2} \cdot \frac{\log n}{\log n + \log p})$ multiplicative saving. In light of this, CSAMs preserve more privacy as well.

To analyze the outcomes of CSAMs, we focus on 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, given the items she chose in previous steps.
- *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 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. □

Ordinal Efficiency of CSAMs

In this section, we focus on characterizing the *ordinal efficiency* of CSAMs measured by agents' ranks of the bundles they receive.[3] 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 CSAM $f_{\mathcal{O}}$, we introduce the following notation for any $j \leq n$.

- Let $\mathcal{O}_j$ denote the linear order over the categories $\{1, \ldots, p\}$ according to which agent j chooses items from in $\mathcal{O}$.
- 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, $\mathcal{O} = [(1, 1) \rhd (2, 2) \rhd (3, 1) \rhd (3, 2) \rhd (2, 1) \rhd (1, 2)]$.

$\mathcal{O}_1 = 1 \rhd 2$. $K_1 = 2$ 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. $k_{3,1} = k_{3,2} = 2$. □

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

Proof sketch: W.l.o.g. let $\mathcal{O}_j = 1 \rhd 2 \rhd \cdots \rhd p$. If j is optimistic, then we let $t_j = \mathcal{O}^{-1}(j, K_j)$ and let $(d_{j,1}, \ldots, d_{j,K_j-1}) \in D_1 \times \cdots \times D_{K_j-1}$ denote the items agent j chose in the previous rounds. It follows that at the beginning of round t_j, the following $\prod_{l=K_j}^{p} k_{j,l}$ bundles are available for agent j: $\mathfrak{D}_j = (d_{j,1}, \ldots, d_{j,K_j-1}) \times \prod_{l=K_j}^{p} D_{l,t_j}$. By the definition of K_j, no agent can

[3] This is different from the ordinal efficiency for randomized allocation mechanisms [2].

interrupt agent j from choosing the items in her top-ranked bundle in $\mathfrak{D}_j$, and $|\mathfrak{D}_j| = \prod_{l=K_j}^{p} k_{j,l}$.

If j is pessimistic, then we let $\vec{d_j} = (d_{j,1}, \ldots, d_{j,p}) = f^j_{\mathcal{O}}(P)$ denote her allocation by $f_{\mathcal{O}}$. By the definition of pessimism and the assumption that for any $1 \leq l \leq p$, in round $t^* = \mathcal{O}^{-1}(j, l)$ agent j chose $d_{j,l}$ from D_{l,t^*}, we must have that for all $d'_l \in D_{l,t^*}$ with $d'_l \neq d_{j,l}$, there exists an bundle $(d_{j,1}, \ldots, d_{j,l-1}, d'_l, \ldots, d'_p)$ that is ranked below $\vec{d_j}$. Such bundles are all different and the number of them is $\sum_{l=1}^{p}(k_{j,l} - 1)$, which proves the bound for pessimistic agents. ■

We note that Proposition 1 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. In addition, once the CSAM and the properties of the agents (that is, whether each agent is optimistic or pessimistic) is given, the bounds hold for all preference profile.

Our main theorem in this section states that, surprisingly, for all combinations of optimistic and pessimistic agents, all upper bounds described in Proposition 1 can be matched in a same profile. Even more surprisingly, for the same profile there exists an allocation where almost all agents get their top-ranked bundle, and the only agent who may not get her top-ranked bundle gets her second-ranked bundle. Therefore, the theorem not only provides a worst-case analysis in the absolute sense in that all upper bounds in Proposition 1 are matched in the same profile, but also in the comparative sense w.r.t. the optimal allocation of the profile.

Theorem 2 *For any CSAM $f_{\mathcal{O}}$ and any combination of optimistic and pessimistic agents, there exists a profile P such that for all $j \leq n$:*

1. *if agent j is optimistic, then the rank of the bundle allocated to her is $n^p + 1 - \prod_{l=K_j}^{p} k_{j,O_j(l)}$;*
2. *if agent j is pessimistic, then the rank of the bundle allocated to her is $n^p - \sum_{l=1}^{p}(k_{j,O_j(l)} - 1)$;*
3. *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.*

The proof is quite involved and can be found in the supplementary material.

Example 5 The profile in Example 3 is an example of the profile guaranteed by Theorem 2: agent 1 (optimistic) gets her bottom bundle ($K_1 = 2$ and $k_{1,2} = 1$), agent 2 (optimistic) gets her bottom bundle ($K_2 = 2$ and $k_{2,1} = 1$), and agent 3 (pessimistic) gets her third bundle ($k_{3,1} = k_{3,2} = 2$). Moreover, there exists an allocation where agent 2 and agent 3 get their top bundles and agent 1 gets her second bundle. □

Theorem 2 can be used to compare various CSAMs with optimistic and pessimistic agents w.r.t. worst-case utilitarian rank and worst-case egalitarian rank.

Definition 2 *Given any CSAM $f_{\mathcal{O}}$ and any n, the* worst-case utilitarian rank *is* $\max_{P_n} \sum_{R_j \in P_n} Rank(R_j, f^j_{\mathcal{O}}(P_n))$, *and the* worst-case egalitarian rank *is* $\max_{P_n} \max_{R_j \in P_n} Rank(R_j, f^j_{\mathcal{O}}(P_n))$, *where P_n is a profile of n agents.*

In words, 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, which is also a well-accepted measure of fairness. The worst case is taken over all profiles of n agents.

Proposition 2 *Among all CSAMs, serial dictatorships with all-optimistic agents have the best (smallest) worst-case utilitarian rank and the worst (largest) worst-case egalitarian rank.*

Proposition 3 *Any CSAM with all-optimistic agents has the worst (largest) worst-case egalitarian rank, which is n^p.*

Proposition 4 *For any even number p, the worst-case egalitarian rank of any balanced CSAM (defined in Example 2) with all-pessimistic agents is $n^p - (n-1)p/2$. These are the CSAMs with the best worst-case egalitarian rank among CSAMs with all-pessimistic agents.*

A natural question after Proposition 4 is: do the balanced CSAMs with all-pessimistic agents have optimal worst-case egalitarian rank, among all CSAMs *for any combination of optimistic and pessimistic agents*? The answer is negative.

Proposition 5 *For any even number p with $2^p > 1 + (n-1)p/2$, there exists a CSAM with both optimistic and pessimistic agents, whose worst-case egalitarian rank is strictly better (smaller) than $n^p - (n-1)p/2$.*

Summary and Future Work

In this paper we propose CDAPs to model allocation problems for indivisible and categorized items without monetary transfer, when agents have heterogenous and combinatorial preferences. We characterize serial dictatorships for basic CDAPs, propose CSAMs and characterize worst-case ordinal efficiency for CSAMs with any combination of optimistic and pessimistic agents, which leads to characterizations of utilitarian rank and egalitarian rank of various CSAMs.

There are many open questions and directions for future research, including analyzing the outcomes and ordinal efficiency for CSAMs for other types of agents, e.g. strategic agents and minimax-regret agents. We also plan to work on expected utilitarian rank and egalitarian rank (some simulation results are included in the supplementary material), and randomized allocation mechanisms. For general CDAPs, we are excited to explore generalizations of CP-nets [4], LP-trees [3], and soft constraints [23] 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.

References

[1] Bhattacharya, Arka A. and Culler, David and Friedman, Eric and Ghodsi, Ali and Shenker, Scott and Stoica, Ion. Hierarchical scheduling for diverse datacenter workloads. In *Proceedings of the 4th Annual Symposium on Cloud Computing*, pages 4:1–4:15, Santa Clara, CA, USA, 2013.

[2] Anna Bogomolnaia and Hervé Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

[3] 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, 2010.

[4] 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.

[5] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. In *ACM Conference on Electronic Commerce*, pages 197–214, Valencia, Spain, 2012.

[6] 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.

[7] 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.

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

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

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

[11] 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.

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

[15] Ghodsi, Ali and Zaharia, Matei and Hindman, Benjamin and Konwinski, Andy and Shenker, Scott and Stoica, Ion. Dominant Resource Fairness: Fair Allocation of Multiple Resource Types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, pages 323–336, Boston, MA, USA, 2011.

[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] Elias Koutsoupias and Christos Papadimitriou. Worst-case Equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, pages 404–413, Trier, Germany, 1999.

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

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

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

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

[23] 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.

[24] 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.

[25] 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.

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