Multilevel Fair Allocation under Additive Preferences

Maxime Lucet, Nawal Benabbou, Aurélie Beynier, and Nicolas Maudet
{firstname.lastname}@lip6.fr

LIP6, Sorbonne Université

Abstract. We introduce multilevel fair resource allocation with tree-structured hierarchical relations among agents. At each level, the problem can be viewed locally as allocating an agent's bundle to its children, the overall allocation being a trace of this process iterated down to the leaves, with each internal node potentially having its own local allocation mechanism. Assuming that internal nodes have additive utilities over their children's, and the leaves have classicasl additive utilities over items, we first show that usual envy-based fairness notions (e.g. WEF1) must be adapted to the multilevel setting. We propose three adaptations, and show that choosing between the three is not neutral. We show that under identical preferences the three adapted envy-based notions coincide, and that the multilevel extension of the Weighted Round Robin (11) (MWRR) guarantees them. We then prove that under general preferences, the MWRR may guarantee some, while failing others. Finally, we argue through experimentations that the MWRR may still perform well, even on adaptations which it does not formally guarantee.

Keywords: Fair allocation · Computational Social Choice · Resource allocation.

1 Introduction

A multiagent resource allocation problem consists in determining a fair and/or efficient distribution of a set of items among a set of agents (7). Most existing work assumes that items are allocated directly to the agents or groups who will use them. However, in many real-world settings, agents are embedded within hierarchical organizations – groups nested within larger groups – forming a multilevel hierarchical organization. Such a hierarchy can be naturally modeled as a directed tree, where each node is responsible for allocating the bundle of items it receives to its children.

Consider a food charity operating in a country with multiple administrative levels: regional branches oversee departmental units, which in turn manage distribution centres at the city level. Hence, the food charity operates as a hierarchical structure represented in Fig. 1. To ensure that no geographical area is neglected, resources are typically allocated in stages—first across regions, then among departments within each region, and finally among cities within each department. In such a setting, fairness should be enforced locally at each level of the hierarchy (e.g., among departments within a region or among cities within a department), rather than across levels, as entities at different levels are not directly comparable. This example extends to many allocation problems across territories and hierarchical organizations. Moreover, they can occur even in the absence of an explicit hierarchy. For instance, when allocating jobs to unemployed individuals characterized by gender, age, and education, one may first focus on fairness across gender (e.g., putting more weight on women to address underrepresentation), then across education for women (to promote more qualified women) and across age for men by assigning equal weight to each age group (to balance the representation). Such hierarchical preferences can be naturally represented in the multilevel framework as a directed tree in which the nodes are weighted.

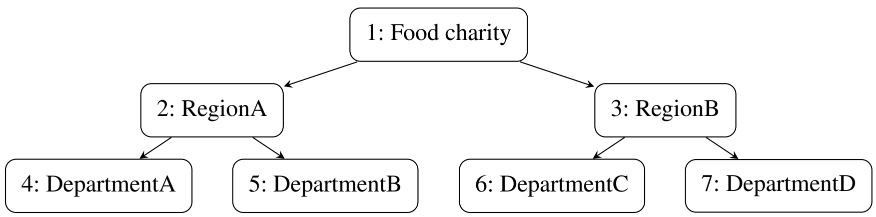

Fig. 1: Hierarchical structure of a food charity.

Different assumptions can be made regarding the utility of the (internal) nodes of the tree: it could be (i) that agents only care about the bundle they receive and are oblivious of the subsequent allocation of those items; (ii) that agents have preferences over the properties of the allocation to their children; or (iii) that agents have preferences that depend on the actual utilities of their children. In this paper, we adopt assumption (iii), which raises notable algorithmic challenges when allocations are computed using a top-down approach: an internal agent must allocate items to satisfy its children, even though their utilities are not yet known, as these depend on subsequent allocations at lower levels.

Contributions. In this paper, we formally introduce multilevel fair allocation problems, assuming that the leaves have additive preferences, and focusing on weighted envy-based fairness notions. We first show that the usual weighted envy-free up to one good (WEF1) notion must be adapted to the multilevel setting, and that the choice of adaptation is not neutral. We propose three adaptations, ranging from the weakest to the strongest: pessimistic, agnostic, and optimistic. Under identical preferences, our three adaptations coincide and can be guaranteed by a multilevel extension of the Weighted Round Robin (11) (MWRR). However, results become more nuanced under general additive preferences. Although the MWRR guarantees our weakest adaptation, it fails to satisfy the stronger ones. Nevertheless, we show in Section 4.3 that the algorithm still performs extremely well experimentally with respect to the medium fairness notion.

Related work. Recently, *group fairness* has become a highly active topic in fair division. Notably, (4; 13) studied group fairness notions requiring fairness to be achieved among *any* group, and introduced different relaxations of group envy-freeness. *Weighted fairness* has also been extensively studied (11; 12; 3). The multilevel setting studied in this paper was recently introduced by (15). While that work focuses on agents with matroid rank valuations, we instead consider additive valuations and study envy-based notions of fairness, which differ from those examined in their framework. (17) study individual and group envy-freeness in settings where groups' valuations depend on the valuations of their constituent agents, and they propose algorithms that provide guarantees at both levels. Our paper extends the setting further, and some of their algorithms. Closely related is the work from (9) which also seeks to reconcile the perspectives of groups and individual agents. They argue that agents at different levels may have distinct preferences, notably that groups' preferences do not necessarily aggregate the valuations of their constituent agents (an assumption we make in this paper). Their model however remains limited to two levels, while ours can be defined for any number of levels. (2) studies a bilevel allocation problem, motivated by the way some food charities operate. In contrast with our paper, their paper focuses on auction mechanism. The multilevel fairness notion we propose in this paper induces a visibility constraint, which may be reminiscent of some lines of work on local fairness (1; 5; 8). Indeed, in our model, only labs affiliated to the same department can compare their situations. Finally, our multilevel model also resembles the recent multilevel apportionment model from (18). Indeed, their hierarchical model is similar to ours, representing the hierarchy through a tree. However, the scope of the paper is specifically on apportionment method, a very specific subfield resource allocation.

2 Model

In this paper, we consider allocation problems in which a set of m items/goods, denoted by $\mathcal{G} = \{g_1, \ldots, g_m\}$, must be distributed among agents organized in a hierarchical structure.

Hierarchical structure. We consider a multilevel allocation problem represented by an *arborescence* (i.e. a directed rooted tree in which all edges point away from the root) denoted by $\mathcal{T} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{1, \ldots, n\}$ is the set of nodes representing agents and $\mathcal{E}$ is the set of arrows representing hierarchical relations among agents. We assume that the nodes are indexed according to a topological ordering of $\mathcal{T}$ (hence the root of $\mathcal{T}$ is node 1). For any node $i \in \mathcal{N}$, let $h(i)$ be the height of i, i.e. the number of arrows in a longest path starting at node i. Let $\mathcal{C}(i)$ denote the set of children of i, which is defined by $\mathcal{C}(i) = \{j \in \mathcal{N} : (i, j) \in \mathcal{E}\}$. Let $\mathcal{P}(i)$ be the parent of i, which is the unique node such that $(\mathcal{P}(i), i) \in \mathcal{E}$. Moreover, let $Anc(i)$ be the set of its ancestors, i.e. nodes belonging to the unique path from 1 to i (including i itself). For any node i, let $\mathcal{T}_i = (\mathcal{N}_i, \mathcal{E}_i)$ denote the subtree of $\mathcal{T}$ rooted at i, consisting of all nodes and edges belonging to paths that start at i. Let $\mathcal{L}(i)$ be the set of leaves of $\mathcal{T}_i$, formally defined as $\mathcal{L}(i) = \{j \in \mathcal{N}_i : \mathcal{C}(j) = \emptyset\}$. Let $\mathcal{I}(i)$ denote the set of internal nodes of $\mathcal{T}_i$, defined by $\mathcal{I}(i) = \mathcal{N}_i \backslash \mathcal{L}(i)$. In particular, $\mathcal{L}$ and $\mathcal{I}$ denote the sets of leaves and internal nodes of the entire tree. For brevity, we write $\mathcal{I} := \mathcal{I}(1)$ and $\mathcal{L} := \mathcal{L}(1)$ throughout the paper, and use $\mathcal{I}(i)$ and $\mathcal{L}(i)$ when referring to a specific subtree.

Allocations. We introduce different types of allocations relevant to our setting:

Definition 1 (Complete multilevel allocation). $\pi : \mathcal{N} \to 2^{\mathcal{G}}$ *is a* multilevel allocation *if it satisfies the following properties: (i)* $\pi(1) = \mathcal{G}$*, (ii)* $\pi(i) \supseteq \cup_{j \in \mathcal{C}(i)} \pi(j), \forall i \in \mathcal{I}$*, and (iii)* $\pi(i) \cap \pi(j) = \emptyset, \forall i, j \in \mathcal{N}$ *such that* $\mathcal{P}(i) = \mathcal{P}(j)$*, where* $\pi(i)$ *denotes the* *bundle of any node $i \in \mathcal{N}$. A multilevel allocation π is* complete *if condition (2) is an equality.*

We require that the root owns all items (i.e. $\pi(1) = \mathcal{G}$) only to ensure that none are discarded a priori. Moreover, we require that each internal node allocate each of its item to at most one child. The set of all multilevel allocations is denoted by Π hereafter.

Definition 2 (Restricted multilevel allocation). *Given a multilevel allocation $\pi \in \Pi$ and a set of nodes $N \subseteq \mathcal{N}$, the* restriction of π to the nodes in N *is denoted by $\pi|_N$ and defined by $\pi|_N = (\pi(i))_{i \in N}$.*

Definition 3 (Local allocation). *Given a set of nodes $N \subseteq \mathcal{N}$ and a set of items $S \subseteq \mathcal{G}$, $A : N \to 2^S$ is a* local allocation *if (1) $\cup_{i \in N} A(i) \subseteq S$ and (2) $A(i) \cap A(j) = \varnothing$ for all $i, j \in N$. Such local allocation A is complete if condition (1) is an equality.*

For any $N \subseteq \mathcal{N}$ and $S \subseteq \mathcal{G}$, let $\mathcal{A}_N^S$ denote the set of corresponding complete local allocations. Note that, for any complete multilevel allocation $\pi \in \Pi$ and any node $i \in \mathcal{I}$, the restricted allocation $\pi|_{\mathcal{C}(i)}$ is a complete local allocation belonging to $\mathcal{A}_{\mathcal{C}(i)}^{\pi(i)}$.

Utility model. Let $v_i : \Pi \to \mathbb{R}_{\geqslant 0}$ be the utility function of node $i \in \mathcal{N}$, and let $v = (v_i)_{i \in \mathcal{N}}$. For any internal node $i \in \mathcal{I}$ and any multilevel allocation $\pi \in \Pi$, $v_i(\pi)$ quantifies some concept of overall welfare derived by the children of i from π. We focus here on the utilitarian social welfare, i.e.

$$v_i(\pi) = \sum_{j \in \mathcal{C}(i)} v_j(\pi)$$

Hence, we assume that internal nodes have additive utilities over their children. This assumption is realistic in many applications where internal nodes represent entities whose utility derives solely from the satisfaction of their members. Note that, by linearity of summation, $v_i(\pi)$ can be rewritten as $v_i(\pi) = \sum_{x \in \mathcal{L}(i)} v_x(\pi)$, where the sum ranges over the leaves of the subtree $\mathcal{T}_i$.

In contrast, leaves do not have children, and hence, we equip them with classical additive utilities, i.e. any leaf $x \in \mathcal{L}$ is equipped with a utility function $u_x : 2^{\mathcal{G}} \to \mathbb{R}_{\geqslant 0}$ such that $v_x(\pi) = u_x(\pi(x))$ for any multilevel allocation $\pi \in \Pi$. Formally, $u_x(\varnothing) = 0$ and $u_x(S) = \sum_{g \in S} u_x(g)$ where $S \subseteq \mathcal{G}$ denotes a bundle of items, and $u_x(g)$ is a slight abuse of notation to denote, for simplicity, the utility of leaf x for item g.

Estimated utility functions. Our goal is to compute a multilevel allocation that is both complete and fair. We focus on envy-based fairness notions, which typically require that each agent values its own bundle at least as much as any other agent's bundle. Such notions rely on agents being able to evaluate bundles directly. However, in our setting, each internal node $i \in \mathcal{I}$ defined over allocations rather than individual bundles. As a result, standard envy-based notions cannot be applied directly.

We therefore begin by introducing a method to evaluate bundles, which will allow us to adapt these fairness notions to our framework. As highlighted by (17), there are two ways to do so: (i) we can assume that i computes an allocation of S in its subtree $\mathcal{T}_i$, and then i's utility would simply be the sum of the utilities of its children (or by linearity of the sum, of its leaves), or (ii) we can also assume that i is agnostic about how her items are allocated, and then we would compute the weighted average of S over all leaves in $\mathcal{L}(i)$.

In this paper, we investigate both propositions. For (i), we propose two natural approaches: the first one, which we call the *optimistic* approach, assumes that i computes a utilitarian-optimal allocation of S to its leaves $\mathcal{L}(i)$, while the *pessimistic* approach assumes that i compute an allocation of S to $\mathcal{L}(i)$ minimizing the utilitarian-welfare.

Definition 4 (Optimistic estimated utility function). *Given a node $i \in \mathcal{N}$ and a subset of items $S \subseteq \mathcal{G}$, the* optimistic estimated utility function *of i for S is*

$$\widehat{v}_i(S) = \max_{A \in \mathcal{A}^S_{\mathcal{L}(i)}} \sum_{x \in \mathcal{L}(i)} \sum_{g \in A(x)} u_x(g)$$

Definition 5 (Pessimistic estimated utility function). *Given a node $i \in \mathcal{N}$ and a subset of items $S \subseteq \mathcal{G}$, the* pessimistic estimated utility function *of i for S is*

$$\check{v}_i(S) = \min_{A \in \mathcal{A}^S_{\mathcal{L}(i)}} \sum_{x \in \mathcal{L}(i)} \sum_{g \in A(x)} u_x(g)$$

For (ii), we refer to the weighted average evaluation of bundle S as the *agnostic* approach. On the contrary of the optimistic and pessimistic approaches, the agnostic approach does not assume a specific allocation to the leaves, and measures the value of a bundle as a weighted average utility over all leaves in $\mathcal{L}(i)$.

Definition 6 (Agnostic estimated utility function). *Given a node $i \in \mathcal{N}$ and a subset of items $S \subseteq \mathcal{G}$, the* agnostic estimated utility function *of i for S is*

$$\bar{v}_i(S) = \sum_{x \in \mathcal{L}(i)} \sum_{g \in S} u_x(g) \cdot W(x, i)$$

where $W(x,i) = \Pi_{k \in Anc(x) \setminus Anc(i)} \frac{w_k}{w_{\mathcal{C}(\mathcal{P}(k))}}$, *where* $w_{\mathcal{C}(\mathcal{P}(k))} = \sum_{k' \in \mathcal{C}(\mathcal{P}(k))} w_{k'}$.

For any internal node $i \in \mathcal{I}$ and any leaf $x \in \mathcal{L}(i)$, $W(x,i)$ can be interpreted as the weight of x in the subtree $\mathcal{T}_i$ (rooted in i). Note that $W(i,j)$ is defined similarly for any node $i \in \mathcal{N}$ and one of its ancestors $j \in Anc(i)$. Moreover, we show in Appendix B that $\sum_{x \in \mathcal{L}(i)} W(x,i) = 1$.

Furthermore, notice that for any leaf $x \in \mathcal{L}$ and any bundle $S \subseteq \mathcal{G}$, we have $\widehat{v}_x(S) = \bar{v}_x(S) = \check{v}_x(S) = u_x(S)$. Indeed, by definition leaf x is a tree with only node, and hence $\mathcal{L}(x) = \{x\}$.

Hereafter, we refer to these utility functions as estimated utility functions.

Example 1. We now illustrate the different estimated utility functions. Assume an instance with $\mathcal{N} = \{1, \ldots, 7\}$ organized as in Tree 1, and $\mathcal{G} = \{g_1, \ldots, g_5\}$. Leaves 4 and 6 have the following preferences $u_x(g) = 2$ for $x \in \{4, 6\}$ and any $g \in \mathcal{G}$; leaves 5 and 7 have the following preferences $u_x(g) = 1$ for $x \in \{5, 7\}$ and any $g \in \mathcal{G}$. Assume the weight of any internal node $i \in \mathcal{I}$ is $w_i = |\mathcal{L}(i)|$, and that of any leaf $x \in \mathcal{L}$ is 1. Suppose we have the multilevel allocation $\pi \in \Pi$ such that $\pi(4) = \{g_1, g_5\}$, $\pi(5) = \{g_3\}$, $\pi(6) = \{g2\}$, and $\pi(7) = \{g_4\}$. By definition, $\pi(2) = \{g_1, g_3, g_5\}$ and $\pi(3) = \{g_2, g_4\}$.

To assess whether node 3 is envious towards node 2, we need to choose an estimated utility function to estimate the value of 2's bundle. Depending on which estimated utility functions we choose, the results might differ. Indeed, in this example, we can see that: $\widehat{v}_3(\pi(2)\backslash\{g_1\}) = 4$ since all items would be allocated to leaf 6 in a utilitarian-optimal allocation ; $\overline{v}_3(\pi(2)\backslash\{g_1\}) = 3$; and finally, $\check{v}_3(\pi(2)\backslash\{g_1\}) = 2$ since all items would be allocation to leaf 7. Hence, the conclusions on whether 3 envies 2 would differ based on which notion you use: according to the optimistic function, 3 is weighted envious towards 2 even up to one good, while according to both agnostic and pessimistic, 3 is weighted envy-free towards 2 up to one good.

Fairness. A central notion of fairness in classical fair division is *envy-freeness* (EF), together with its widely studied relaxations: *envy-freeness up to one good* (EF1) and *envy-freeness up to any good* (EFX). Recent works have also examined *weighted* variants of these criteria (11; 12), namely WEF1 and WEFX. The general idea behind these notions is that no agent should envy the bundle received by another agent, after accounting for the weight differences. We now define a multilevel extension of these fairness notions. This definition is parametrized to accommodate our three estimated utility functions.

Definition 7 **(M[$\mathcal{E}$).** *-WEF1] A multilevel allocation $\pi \in \Pi$ is* estimated multilevel weighted envy-free up to one good *(M[$\mathcal{E}$]-WEF1), for $\mathcal{E} \in \{pess, agno, opt\}$, if for any internal node $i \in \mathcal{I}$, and any pair of children $j, k \in \mathcal{C}(i)$, we have*

$$\frac{v_j(\pi)}{w_j} \geqslant \frac{\tilde{v}_j(\pi(k)\backslash\{g\})}{w_k}$$

for some item $g \in \pi(k)$, and some estimated utility function $\tilde{v}$.

We can then easily derive the definitions for M[opt]-WEF1, M[agno]-WEF1, and M[pess]-WEF1. Moreover, the corresponding *M[$\mathcal{E}$]-WEFX* definition (for $\mathcal{E} \in$ {pess, agno, opt}) is identical to Def. 7 except g can be any item in $\pi(k)$.

Remark 1. If $\mathcal{T}$ has height $h(1)$=1, the problem reduces to a standard *monolevel* allocation setting, and our three fairness notions coincide with the classical WEF1 notion.

3 Identical additive valuations

In (17), the authors propose algorithms achieving strong fairness guarantees under two assumptions: (i) all agents share identical preferences, and (ii) agents within the same group share identical preferences. We show that their algorithms and arguments extend naturally to our multilevel setting. In our framework, these assumptions translate to: (i) all leaves $x \in \mathcal{L}$ have identical preferences, and (ii) all leaves within each subtree rooted at a child of the root—i.e., $x \in \mathcal{L}(i)$ for $i \in \mathcal{C}(1)$—have identical preferences. While both assumptions are strong, (ii) can arise in practice when the tree has limited depth. For instance, in a food charity operating across departments with multiple local branches, branches within the same department may face similar needs and thus share preferences over supplies. We refer to (i) as *all-common* valuations, and to (ii) as *root-child-common* valuations.

Remark 2. Under both all-common and root-child-common valuations, for any internal node $i \in \mathcal{I}$ and any bundle $S \subseteq \mathcal{G}$, we have

$$\sum_{x \in \mathcal{L}(i)} u_x(A(x)) = \sum_{x \in \mathcal{L}(i)} u_x(B(x))$$

for all $A, B \in \mathcal{A}^S_{\mathcal{L}(i)}$. In words, the allocation of S the leaves in $\mathcal{L}(i)$ does not affect the utility derived by node i. Accordingly, in this section, we slightly abuse notation and refer to the value of a bundle even for internal nodes, rather than the value of a multilevel allocation.

Consequently, under both preferences, the three estimated utility functions coincide. In particular, for every node $i \in \mathcal{N}$, every bundle $S \subseteq \mathcal{G}$, and a (all- or root-child-) common utility function u, we have:

$$\widehat{v}_i(S) = \overline{v}_i(S) = \check{v}_i(S) = u(S).$$

Hence, it follows that the pessimistic, agnostic, and optimistic adaptations of m-WEF1 all coincide, and can be rewritten in the following classical form. A multilevel allocation π is M-WEF1 if, for every internal node $i \in \mathcal{I}$, every pair of children $j, k \in \mathcal{C}(i)$, and some item $g \in \pi(k)$,

$$\frac{u(\pi(j))}{w_j} \geqslant \frac{u(\pi(k) \backslash \{g\})}{w_k}.$$

3.1 All-common valuations

Theorem 1. *Under all-common valuation, a M-WEFX allocation can be computed in polynomial time.*

Proof. The following algorithm is a multilevel extension of the algorithm presented in the proof of Theorem 3.1 in (17). Order the goods in decreasing order of preferences. Starting at the root, select the child with the least weighted bundle value (i.e. the utility of the bundle divided by the weight of the node). If this child is not a leaf, repeat the selection process until selecting a leaf, denoted $x \in \mathcal{L}$. Allocate the first remaining item (i.e. the highest-value remaining item) to the selected leaf. At any point in time, the leaf we selected (or any of its ancestors) cannot be weighted envied before the allocation of the new item by one of its siblings (as it has the least valued bundle). Any envy that forms towards any of the nodes $i \in Anc(x)$ can only result of the latest good allocated, and any envy will disappear upon dropping this good. Moreover, this good is also the least valued one in i's bundle, by construction. The resulting multilevel allocation is thus M-WEFX. Furthermore, the algorithm is polynomial: sorting the goods is polynomial-time doable, the leaf-selection process is in $O(n)$, and allocating the chosen item at an iteration also takes $O(n)$. □

We then show that we can obtain M-WEF1 jointly with some fairness notion at the leaves. The full proof and the algorithm can be found in the appendix. One component of the algorithm presented in the appendix is the Multilevel Weighted Round Robin, presented in Algorithm 1.

Theorem 2. *Under all-common valuations and* $w_i = |\mathcal{L}(i)|, \forall i \in \mathcal{N}$*, there exists a polynomial-time algorithm that returns an allocation that is M-WEF1 and EFX between all leaves in* $\mathcal{L}$*.*

3.2 Child-root-common valuations

We now introduce a multilevel extension of the Weighted Round Robin (WRR) algorithm from (11). In the standard monolevel setting, WRR is a polynomial-time algorithm that produces complete WEF1 allocations. The algorithm operates as follows: each agent is assigned a picking score, which counts the number of times the agent has been selected to choose an item. At each step, the agent with the lowest weighted picking score (i.e., the picking score divided by the agent's weight) selects her most preferred item among the remaining ones. (17) propose a variant of this procedure, called Iterative Weighted Round Robin (IWRR), in a setting where agents are partitioned into groups. In addition to agents, groups are also assigned picking scores and weights: a group's weight corresponds to its number of agents, while each agent has weight 1. The algorithm first selects the group with the lowest weighted picking score, and then selects, within that group, the agent with the lowest picking score. The IWRR can be viewed as a special case of the multilevel algorithm we introduce in this paper. The simplicity and strong fairness guarantees of these procedures make them natural candidates for extension to the multilevel setting.

The principle of our algorithm, which we call Multilevel Weighted Round Robin (MWRR) is the following: each node $i \in \mathcal{N}$ is equipped with a picking score t_i which counts the number of times i was picked by the MWRR, and this picking score is weighted by i's weight. Then, starting at the root, MWRR picks the child $i \in \mathcal{C}(1)$ minimizing this picking score. If i is an internal node, i.e. if $i \in \mathcal{I}$, we repeat the selection process, i.e. we select the child $j \in \mathcal{C}(i)$ which minimizes the picking score, until the selected node is a leaf. Once it is the case, the selected leaf gets to choose her preferred item among the remaining ones. We repeat this procedure until all items are allocated. For the leaf selection procedure, ties are broken differently depending on whether they occur between internal nodes or leaves: (i) among internal nodes, ties are broken lexicographically; (ii) among leaves, they are broken in favor of the leaf with the highest utility for any remaining item. Pseudocode can be found in Algorithm 1.

Our goal is to study what fairness properties can MWRR guarantee. We first look at its performance under child-root-common additive valuations, and in Section 4 we study thoroughly how fair it is in the general additive case.

Theorem 3. *Under root-children-common valuations and arbitrary weights, the MWRR returns an M-WEF1 allocation in polynomial time.*

Proof. We first show that the MWRR runs in polynomial time: the while loop (Line 6) runs in $O(m)$ steps. Finding the child of the root with minimum weighted picking score (Line 7) is at most in $O(n)$ (for tree of height 1), and a loose bound for the while loop (Line 8) is $O(n^2)$. Finding the item with maximum utility (Line 11) can be done in $O(m)$, and updating the bundle (Line 12) requires at most $O(n)$ (in a comb tree). Hence, the complexity of the algorithm is $O(m(n^2 + m))$.

Algorithm 1 Multilevel Weighted Round Robin (MWRR)

Input: $\mathcal{T}$ - a multilevel tree ; $\mathcal{G}$ - a set of items ; the leaves' valuations $(u_x)_{x\in\mathcal{L}}$
Output: π - a multilevel allocation
Initialize π with empty bundle for any $i \in \mathcal{N}\backslash\{1\}$ and $\pi(1) = \mathcal{G}$
Initialize picking scores s.t. $t_i = 0, \forall i \in \mathcal{N}$
$\mathcal{RI} \leftarrow \mathcal{G}$ ▷ Initial set of remaining items
while $\mathcal{RI} \neq \emptyset$ **do**
  $i = \arg\min_{i'\in\mathcal{C}(1)} \frac{t_{i'}}{w_{i'}}$
  while $i \notin \mathcal{L}$ **do**
    $i = \arg\min_{i'\in\mathcal{C}(i)} \frac{t_{i'}}{w_{i'}}$
  end while
  $g = \arg\max_{g'\in\mathcal{RI}} u_i(g')$
  while $i \neq 1$ **do**
    $\pi(i) \leftarrow \pi(i) \cup \{g\}$
    $i \leftarrow \mathcal{P}(i)$
  end while
  $\mathcal{RI} \leftarrow \mathcal{RI}\backslash\{g\}$
end while
return π

Then, recall that for any child of the root, $i \in \mathcal{C}(1)$, all leaves in $\mathcal{L}(i)$ have the same preferences over singletons, i.e. $\forall x, y \in \mathcal{L}(i), \forall g \in \mathcal{G}, u_x(g) = u_y(g)$. Hence, any child of the root i can be seen as an agent with additive utility with the same preferences over singletons, as no matter which of its leaves is chosen at any iteration, node i will receive the same item (as all leaves in $\mathcal{L}(i)$ have the same preferences) which will yield the same utility. Hence, at the root, choosing a node $i \in \mathcal{C}(1)$ reduces to choosing an agent with additive utility. Moreover, MWRR at the root chooses a child according to the "least weight-adjusted frequent picker" criterion, exactly like in (11). Since their algorithm is known to be WEF1 w.r.t. agents with additive utilities, we can conclude that MWRR is WEF1 w.r.t. $\mathcal{C}(1)$. Then, we can repeat the same argument recursively for any node $i \in \mathcal{N}$. □

Corollary 1. *Under root-children-common valuations and* $w_i = |\mathcal{L}(i)|, \forall i \in \mathcal{N}$*, the MWRR is M-WEF1 and EF1 between all leaves in* $\mathcal{L}$*.*

Proof. From Theorem 3, we know that MWRR satisfies M-WEF1. Furthermore, the proof of Theorem 3.9 in (17) readily extends to our setting and establishes EF1 among the leaves. The full proof can be found in Appendix A. □

In this section, we showed that the algorithms proposed in (17) could be extended easily to satisfy our multilevel fairness properties, namely M-WEF1 and M-WEFX, sometimes even jointly with fairness at the leaves. In Section 4 we discuss how to extend our multilevel fairness notions under general additive valuations.

4 General additive valuations

We now focus on the more general, and arguably more interesting case where we have general additive valuations. Hence, we drop the identical valuations assumption.

4.1 Relations among M-WEF1 notions

We formally establish the hierarchy between the adaptations we presented.

Lemma 1. *For any internal node $i \in \mathcal{I}$ and bundle $S \subseteq \mathcal{G}$, we have $\widehat{v}_i(S) \geqslant \overline{v}_i(S)$*

Proof. Recall that:

$$\overline{v}_i(S) = \sum_{x \in \mathcal{L}(i)} \sum_{g \in S} u_x(g)\, W(x,i),$$
$$\widehat{v}_i(S) = \max_{A \in \mathcal{A}^S_{\mathcal{L}(i)}} \sum_{x \in \mathcal{L}(i)} \sum_{g \in A(x)} u_x(g) = \sum_{g \in S} \max_{x \in \mathcal{L}(i)} u_x(g).$$

For any $g \in S$, both $\max_{x \in \mathcal{L}(i)} u_x(g)$ and $\sum_{x \in \mathcal{L}(i)} u_x(g)\, W(x,i)$ are convex combinations of $(u_x(g))_{x \in \mathcal{L}(i)}$. The first places all weight on some leaves $x^* \in \arg\max_{x \in \mathcal{L}} u_x(g)$, while the second uses weights $(W(x,i))_{x \in \mathcal{L}(i)}$, which satisfy $\sum_{x \in \mathcal{L}(i)} W(x,i) = 1$ (as shown in Appendix B).

Since a convex combination is maximized by putting full weight on a largest coordinate, we obtain:

$$\max_{x \in \mathcal{L}(i)} u_x(g) \geqslant \sum_{x \in \mathcal{L}(i)} u_x(g)\, W(x,i).$$

Summing over $g \in S$ gives:

$$\widehat{v}_i(S) = \sum_{g \in S} \max_{x \in \mathcal{L}(i)} u_x(g) \geqslant \sum_{g \in S} \sum_{x \in \mathcal{L}(i)} u_x(g)\, W(x,i) = \overline{v}_i(S).$$

Lemma 2. *For any internal node $i \in \mathcal{I}$ and bundle $S \subseteq \mathcal{G}$, we have $\overline{v}_i(S) \geqslant \check{v}_i(S)$*

Proof. The proof is similar to Lemma 1, except we show that $\check{v}_i(S)$ can be written as a convex combination of the utilities of the leaves in $\mathcal{L}(i)$ for item g, where all the weight is put on the leaf with minimum utility for item g. Hence, we have that $\sum_{x \in \mathcal{L}(i)} u_x(g) \cdot W(x,i) \geqslant \min_{x \in \mathcal{L}(i)} u_x(g)$. Summing over all leaves, we obtain that $\overline{v}_i(S) \geqslant \check{v}_i(S)$. □

These relationships between the estimated utility functions induce an implication structure among the fairness notions. The proof trivially results from Lemmas 1 and 2.

Proposition 1. *For any allocation $\pi \in \Pi$, if π is not M[pess]-WEF1, then it is not M[agno]-WEF1. Similarly, if π is not M[agno]-WEF1, then it is not M[opt]-WEF1.*

Remark 3. It might be the case that some allocation that are not M[opt]-WEF1 are M[agno]-WEF1, and some that are not M[agno]-WEF1 are M[pess]-WEF1.

Example 2. In Example 1, the allocation presented is M[agno]-WEF1, but not M[opt]-WEF1.

Example 3. We can slightly modify the instance of Example 1 to exhibit a multilevel allocation that is M[pess]-WEF1 but not M[agno]-WEF1. We now have $\mathcal{N} = \{1, \ldots, 9\}$ organized in the tree of Fig. 2. Leaves 4, 5 and 7 have the following preferences $u_x(g) = 2$ for $x \in \{4, 5, 7\}$ and any $g \in \mathcal{G}$; leaves 6 and 9 have the following preferences $u_x(g) = 1$ for $x \in \{6, 9\}$ and any $g \in \mathcal{G}$. Assume the weight of any internal node $i \in \mathcal{I}$ is $w_i = |\mathcal{L}(i)|$, and that of any leaf $x \in$ is 1. Suppose we have the multilevel allocation $\pi \in \Pi$ such that $\pi(4) = \{g_1\}$, $\pi(5) = \{g_3\}$, $\pi(6) = \{g5\}$, $\pi(7) = \{g_2\}$, $\pi(8) = \{g_4\}$, and $\pi(7) = \varnothing$.

For such allocation π, we know have $\overline{v}_3(\pi(2)\backslash\{g_1\} = 10/3$, while $\breve{v}_3(\pi(2)\backslash\{g_1\}) = 2$. Hence, π is M[pess]-WEF1, but not M[agno]-WEF1.

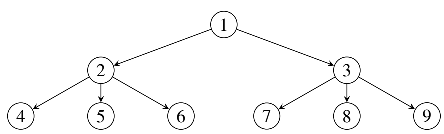

Fig. 2: Tree of Example 3.

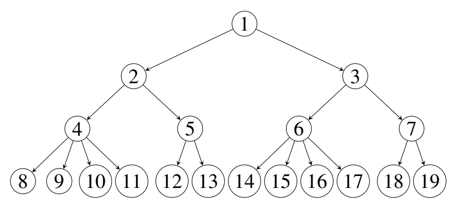

Fig. 3: Tree of Example 4

4.2 Formal guarantees

In this section, we show that the MWRR is guaranteed to return a M[pess]-WEF1 allocation under general additive valuations. Yet, we show that, going up in the hierarchy of notions, the MWRR fails to ensure M[agno]- and M[opt]-WEF1 allocations. We show however in Section 4.3 that MWRR's failures to ensure M[agno]-WEF1 remains rare in practice.

We begin by showing that MWRR guarantees M[pess]-WEF1. Our approach relies on a key lemma, which allows us to adapt the proof of WEF1 for the Weighted Round Robin algorithm from (11) to our multilevel setting.

Theorem 4. *MWRR always returns a M[pess]-WEF1 allocation.*

Proof sketch. The proof follows the same structure as the WRR analysis of (11), with an additional lemma to handle the multilevel pessimistic setting. We first prove the result for the children of the root. The same argument then applies recursively to any internal node, since allocations within each subtree are independent. During the first $|\mathcal{C}(1)|$ steps, each child of the root receives at most one item. Hence, the allocation is trivially pessimistic WEF1 at this stage. As in WRR, the picking rule ensures that for any $i, j \in \mathcal{C}(1)$, we have:

$$\frac{t_j}{t_i} \geqslant \frac{w_j}{w_i}.$$

We show a lemma stating that whenever node j is selected, the utility it derives from the chosen item (through its selected leaf) is at least the worst utility that any leaf in $\mathcal{L}(i)$

can obtain from any item allocated in subsequent steps. In particular, it dominates the worst item eventually received by any sibling i. Combining the picking ratio property with the mentioned lemma, we obtain that the total utility accumulated by j (excluding its first item) is at least the total worst-case utility of i. This implies that j is pessimistic WEF1 with respect to i.

Applying the same argument recursively to each internal node concludes the proof. □

Though, MWRR is guaranteed to return a M[pess]-WEF1 allocation, we show that ensuring M[agno]-WEF1 proves more challenging. In fact, we show there exists some instance where no M[agno]-WEF1 allocation exists (and hence, no M[opt]-WEF1 either). We also show that MWRR may fail to guarantee an M[agno]-WEF1 (hence optimistic) allocation, even when it exists.

Theorem 5. *Under general valuations, an M[agno]-WEF1 allocation is not guaranteed to exist.*

Proof. Consider the multilevel instance represented in Fig. 3. Let $\mathcal{G} = \{g_1, \ldots, g_5\}$ be the set of items. Each leaf is endowed with *general additive utilities* over $\mathcal{G}$: children in $\mathcal{C}(4)$ and $\mathcal{C}(6)$ assign value 2 to every item, while children in $\mathcal{C}(5)$ and $\mathcal{C}(7)$ assign value 1 to every item.

Assume we run MWRR on this instance, and denote π the returned allocation. We give the bundles of the leaves (which suffices to know the multilevel allocation). We have: $\pi(8) = \{g_1\}, \pi(9) = \{g_5\}, \pi(12) = \{g_3\}, \pi(14) = \{g_2\}, \pi(18) = \{g_4\}$, and all other leaves received nothing. One can check that $v_3(\pi) = 3$ while $\bar{v}_3(\pi(2)\backslash\{g\}) = 10/3$ for any $g \in \pi(2)$. Since nodes 2 and 3 have the same weights, the allocation is not M[agno]-WEF1.

Moreover, one can also check that there actually exists no allocation that is M[agno]-WEF1 in this instance. □

This example shows that under general valuations, an M[agno]-WEF1 allocation need not exist. However, this impossibility is not what causes MWRR to fail: the algorithm may return an allocation that is not M[agno]-WEF1 even in instances where such an allocation does exist.

Theorem 6. *Under general valuations, the MWRR is not guaranteed to compute an M[agno]-WEF1 allocation, even when it exists.*

Example 4. We take the same instance problem than in the proof of Theorem 5, and only modify the preferences of the leaves over the singletons. We assume all children x in $\mathcal{C}(5)$ and $\mathcal{C}(7)$ have utility $u_x(g) = 1/5$ for any item in $\mathcal{G}$. Children $x \in \mathcal{C}(4)$ have $u_x(g) = 5/21$ for any $g \in \{g_1, g_2, g_4, g_5\}$ and $u_x(g) = 1/21$ for any $g \in \{g_3\}$. Finally, children $x \in \mathcal{C}(6)$ have $u_x(g) = 5/21$ for any $g \in \{g_1, g_2, g_3, g_5\}$ and $u_x(g) = 1/21$ for any $g \in \{g_4\}$.

For this instance, the MWRR ouput multilevel allocation π, where we have : $\pi(8) = \{g_1\}, \pi(9) = \{g_5\}, \pi(12) = \{g_3\}, \pi(14) = \{g_2\}, \pi(18) = \{g_4\}$, and all other leaves received nothing. One can check that $v_3(\pi) = 46/105 \simeq 0.438$ and $\bar{v}_i(\pi(2)\backslash\{g\}) = 142/315 \simeq 0.451$. Since nodes 2 and 3 have the same weights, we can deduce that π is not M[agno]-WEF1. However, there does exist an M[agno]-WEF1: $\pi(8) = \{g_2\}, \pi(12) = \{g_3\}, \pi(13) = \{g_4\}, \pi(14) = \{g_1\}, \pi(18) = \{g_5\}$. □

Corollary 2. *Under general valuations, the MWRR is not guaranteed to compute an M[opt]-WEF1 allocation, even when it exists.*

Though these results are disappointing, we may still wonder about whether an α-approximation of the M[agno]-WEF1 fairness notion can be guaranteed. Formally, we want to know whether MWRR returns an allocation π guaranteeing that for any internal node $i \in \mathcal{I}$, and any of its children $j, k \in \mathcal{C}(i)$, we have $v_j(\pi) \geqslant \alpha \cdot \overline{v}_j(\pi(k) \backslash \{g\})$ for some item $g \in \pi(k)$, and for some value $\alpha > 0$. Unfortunately, we show that no α-approximation can be guaranteed.

Theorem 7. *For any constant factor $\alpha > 0$, there exists an instance where the allocation returned by MWRR is not M[agno]-WEF1 up to α.*

Proof. Suppose the tree in Fig. 4, where $w_i = |\mathcal{L}(i)|, \forall i \in \mathcal{N}$, and a set of items $\mathcal{G} = \{g_1, g_2, g_3\}$. Assume the following utilities: for any $x \in \mathcal{L}(2)$, $u_x(\{g_1\}) = 0$ and $u_x(\{g_2\}) = u_x(\{g_3\}) = 1$; for any $x \in \mathcal{C}(6)$, $u_x(\{g_1\}) = 1$ and $u_x(\{g_2\}) = u_x(\{g_3\}) = 0$; and for any $x \in \mathcal{C}(7)$, $u_x(\{g_1\}) = 0$ and $u_x(\{g_2\}) = u_x(\{g_3\}) = M$ for M arbitrarily large.

Let π be the allocation returned by MWRR on this instance, we get π such that $\pi(8) = \{g_2\}, \pi(10) = \{g_3\}, \pi(12) = \{g_1\}$. Notice that under this allocation, we have $v_3(\pi) = 1$, and $\overline{v}_3(\pi(2) \backslash \{g_1\}) = W(14, 3) \times M + W(15, 3) \times M = \frac{M}{2}$.

Hence, if we want π to be α M[agno]-WEF1, we need for example $v_3(\pi)/w_3 \geqslant \alpha \cdot \overline{v}_3(\pi(2) \backslash \{g_1\})$. For this inequality to be true, we need $\alpha \leqslant 2/M$. Since M can be arbitrary large, α can be arbitrarily low. □

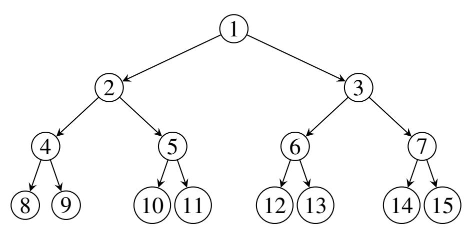

Fig. 4: Tree of Theorem 7

These results are particularly striking when compared to the $\frac{1}{3}$-approximation established by (17) in the bilevel setting. We therefore highlight a fundamental gap between the bilevel case and the multilevel case with trees of height at least three. While in the bilevel setting it remains possible to guarantee an α-approximation for some $\alpha > 0$, this is no longer achievable as soon as we move to three levels, where such guarantees become impossible.

4.3 Numerical tests

In this section, we run experiments to test how often MWRR fails to return an allocation fulfilling these notions. We show that, while MWRR performs may fail often the optimistic notion, it actually is almost always fair under the agnostic one – an experimental finding that substantially nuance the earlier theoretical result. We describe hereafter the experimental protocol.

Protocol. Experiments were conducted on a server equipped with two Intel Xeon E5-2690v3 CPUs running at 2.60GHz, and 192 GB of RAM (each experiment was run on 4 cores). The program is written in Python. All results were obtained over 200 instances.

Trees. We consider three classes of trees in our experiments: (1) Balanced binary trees of branching factor 2, (2) Comb trees, (3) Partially unbalanced trees, which are binary except at the last internal level: for any pair of siblings at this level, one internal node has two children while the other has five. For balanced and comb trees, we tested for $n = \{15, 31, 63, 127\}$, and for partially unbalanced trees, we tested for $n = \{21, 43, 87, 175\}$. Since the number of nodes varies between comb/balanced trees and partially unbalanced trees, we will refer to those number of nodes as {small, medium-, medium+, large}. In all cases, we tested for $m = \{n, 2n\}$ items.

Moreover, we evaluate two weighting schemes: (1) In the first configuration, each node $i \in \mathcal{N}$ is assigned weight $w_i = |\mathcal{L}(i)|$, (2) In the second configuration, the weight of every node is sampled independently and uniformly at random from the integer interval $[1, 6]$.

Preference generation. To ensure a robust experimental evaluation, we generate preferences using four different methods. Specifically, we propose multilevel adaptations of (i) the Mallows model (16), (ii) the resampling Dirichlet model recently introduced by (10), (iii) cost utilities (6), and (iv) correlated utilities as in (14). Importantly, our results are mostly consistent across all generation methods, which is why we do not devote extensive space to their detailed presentation in the main paper. The generation methods are nevertheless detailed in Appendix C.

Results. For each class of instances, we report the 95% confidence interval for the proportion of cases in which the MWRR algorithm produces an M[agno]-WEF1 or an M[opt]-WEF1 allocation. We also provide the average running time of MWRR and its standard deviation. All results are based on 200 instances per class. Due to space constraints, the detailed tables are deferred to the appendix, while we summarize and discuss the main findings below.

Running time. Experimental results show that MWRR is extremely fast and scales well (see Table 1). For instance, for balanced trees, the average running time for $n = 15, m = 15$ (over all generation method, all weighting scheme) is 0.0002 sec., while that for $n = 127, n = 254$ is 0.0191 sec. Moreover, the running time is consistent across different types of instances.

M[agno]-WEF1. Interestingly, our experiments strongly mitigate the impossibility results of Section 4. Although we cannot formally guarantee that MWRR always returns an M[agno]-WEF1 allocation, we observe that it does so in almost all cases. Out of the 192,000 generated instances, fewer than 50 result in allocations that are not M[agno]-WEF1. The few instances that led to unfair allocations were in general for comb trees

Table 1: Average running time (s) $\pm$ std.

n	m	**Comb**	**Balanced**	**Partially unbalanced**
small	n	0.0002 ± 0.0001	0.0002 ± 0.0000	0.0003 ± 0.0001
small	$2n$	0.0003 ± 0.0001	0.0003 ± 0.0001	0.0007 ± 0.0001
medium-	n	0.0006 ± 0.0003	0.0005 ± 0.0001	0.0010 ± 0.0001
medium-	$2n$	0.0012 ± 0.0006	0.0011 ± 0.0001	0.0025 ± 0.0003
medium+	n	0.0027 ± 0.0018	0.0017 ± 0.0001	0.0038 ± 0.0003
medium+	$2n$	0.0057 ± 0.0034	0.0041 ± 0.0003	0.0099 ± 0.0012
large	n	0.0158 ± 0.0126	0.0067 ± 0.0004	0.0155 ± 0.0015
large	$2n$	0.0387 ± 0.0296	0.0191 ± 0.0034	0.0493 ± 0.0085

with random weights (tables can be found in the appendix). This suggests that MWRR is extremely likely, in practice, to return a fair allocation in the agnostic sense.

M[opt]-WEF1. In contrast, the optimistic notion of fairness appears significantly more challenging to satisfy (see Table 2). The performance of MWRR varies substantially depending on the tree structure, the size of the instance, and whether weights are randomly generated. For example, instances based on large comb trees (i.e., for large n) with random weights are particularly difficult, with up to 100% of computed allocations failing to satisfy M[opt]-WEF1. On the other hand, for balanced trees with weights defined as $w_i = |\mathcal{L}(i)|$ for all $i \in \mathcal{N}$, achieving M[opt]-WEF1 appears somewhat less challenging. Nevertheless, the proportion of unfair allocations remains highly variable and can still be substantial. Overall, the choice of weights — random or related to the number of leaves — emerges as the most influential factor affecting fairness performance.

Table 2: Proportion of non-M[opt]-WEF1 allocations (95% CI).
Sub-columns: RW=T for random weights and RW=F for $w_i = |\mathcal{L}(i)|, \forall i \in \mathcal{N}$.

		Comb		**Balanced**		**Part. unbal.**	
n	m	RW= T	RW= F	RW= T	RW= F	RW= T	RW= F
small	n	[0.25, 0.27]	[0.02, 0.03]	[0.02, 0.03]	[0.03, 0.04]	[0.05, 0.06]	[0.04, 0.05]
small	$2n$	[0.24, 0.27]	[0.01, 0.02]	[0.02, 0.03]	[0.02, 0.03]	[0.06, 0.07]	[0.04, 0.05]
medium-	n	[0.53, 0.56]	[0.02, 0.03]	[0.09, 0.11]	[0.09, 0.11]	[0.14, 0.17]	[0.12, 0.14]
medium-	$2n$	[0.57, 0.60]	[0.01, 0.01]	[0.08, 0.09]	[0.09, 0.10]	[0.14, 0.16]	[0.11, 0.12]
medium+	n	[0.70, 0.73]	[0.01, 0.02]	[0.21, 0.23]	[0.21, 0.23]	[0.29, 0.32]	[0.22, 0.25]
medium+	$2n$	[0.71, 0.74]	[0.01, 0.01]	[0.18, 0.20]	[0.12, 0.15]	[0.27, 0.30]	[0.16, 0.18]
large	n	[0.74, 0.77]	[0.01, 0.01]	[0.40, 0.43]	[0.37, 0.40]	[0.48, 0.51]	[0.38, 0.41]
large	$2n$	[0.91, 0.94]	[0.02, 0.04]	[0.55, 0.61]	[0.28, 0.34]	[0.62, 0.69]	[0.51, 0.57]

5 Conclusion

In this paper, we studied a multilevel fair allocation problem consisting in allocating indivisible items fairly in a hierarchical organization. We showed that the usual envy-based fairness notions required to be adapted to the multilevel setting – an adaptation which is far from neutral. Indeed, we showed that the MWRR, a multilevel extension of the Weighted Round Robin (11), could guarantee our weakest adaptation (M[pess]-WEF1), but failed even for the agnostic one. Nevertheless, we showed experimentally that the MWRR still performed extremely well on the agnostic fairness notion, thereby mitigating the lack of formal guarantee.

Open question. In Section 4, we showed that under general additive valuations, there may exist instances where no M[agno]-WEF1 exists. However, our counterexample takes advantage of unbalanced utilities, in the sense that the sum of the utilities over singletons of all leaves did not necessarily sum to the same value. One can check that if we normalized the utilities in this counterexample, the problem would vanish as the instance would no longer be doomed to be unfair. Indeed, an M[agno]-WEF1 solution could now be found. Hence, the existence of an M[agno]-WEF1 allocation remains open for general, *normalized* utilities.

Nevertheless, this does not solve the problem we raised with MWRR as even under normalized utilities, it may fail to return an M[agno]-WEF1 allocation, even when it exist. Indeed, one can see that Example 2 is already normalized.

Bibliography

[1] Abebe, R., Kleinberg, J., Parkes, D.C.: Fair division via social comparison. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems. p. 281–289. AAMAS '17, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2017)

[2] Aggarwal, G., Mertzanidis, M., Psomas, A., Wang, D.: Mechanism design with delegated bidding (2024), https://arxiv.org/abs/2409.19087

[3] Aziz, H., Moulin, H., Sandomirskiy, F.: A polynomial-time algorithm for computing a pareto optimal and almost proportional allocation. Operations Research Letters **48**(5), 573–578 (2020). https://doi.org/https://doi.org/10.1016/j.orl.2020.07.005, https://www.sciencedirect.com/science/article/pii/S0167637720301024

[4] Aziz, H., Rey, S.: Almost group envy-free allocation of indivisible goods and chores. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence. IJCAI'20 (2021)

[5] Beynier, A., Chevaleyre, Y., Gourvès, L., Harutyunyan, A., Lesca, J., Maudet, N., Wilczynski, A.: Local envy-freeness in house allocation problems. Auton. Agents Multi Agent Syst. **33**(5), 591–627 (2019). https://doi.org/10.1007/S10458-019-09417-X, https://doi.org/10.1007/s10458-019-09417-x

[6] Botan, S., Ritossa, A., Suzuki, M., Walsh, T.: Maximin fair allocation of indivisible items under cost utilities. In: Deligkas, A., Filos-Ratsikas, A. (eds.) Algorithmic Game Theory. pp. 221–238. Springer Nature Switzerland, Cham (2023)

[7] Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D.: Handbook of Computational Social Choice. Cambridge University Press, USA, 1st edn. (2016)

[8] Bredereck, R., Kaczmarczyk, A., Niedermeier, R.: Envy-free allocations respecting social networks. Artificial Intelligence **305**, 103664 (2022). https://doi.org/https://doi.org/10.1016/j.artint.2022.103664, https://www.sciencedirect.com/science/article/pii/S0004370222000042

[9] Bu, X., Li, Z., Liu, S., Song, J., Tao, B.: Fair division with allocator's preference. In: Web and Internet Economics: 19th International Conference, WINE 2023, Shanghai, China, December 4–8, 2023, Proceedings. p. 77–94. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-48974-7_5, https://doi.org/10.1007/978-3-031-48974-7_5

[10] Böhm, P., Bredereck, R., Gölz, P., Kaczmarczyk, A., Szufa, S.: Putting fair division on the map. In: Proc. of AAAI'2026, to appear (2026)

[11] Chakraborty, M., Igarashi, A., Suksompong, W., Zick, Y.: Weighted envy-freeness in indivisible item allocation. ACM Trans. Econ. Comput. **9**(3) (Aug 2021). https://doi.org/10.1145/3457166, https://doi.org/10.1145/3457166

[12] Chakraborty, M., Segal-Halevi, E., Suksompong, W.: Weighted fairness notions for indivisible items revisited. CoRR **abs/2112.04166** (2021), https://arxiv.org/abs/2112.04166

[13] Conitzer, V., Freeman, R., Shah, N., Vaughan, J.W.: Group fairness for the allocation of indivisible goods. In: Proceedings of the Thirty-Third AAAI

Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence. AAAI'19/IAAI'19/EAAI'19, AAAI Press (2019). https://doi.org/10.1609/aaai.v33i01.33011853, https://doi.org/10.1609/aaai.v33i01.33011853

[14] Dickerson, J.P., Goldman, J.R., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: Brodley, C.E., Stone, P. (eds.) Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada. pp. 1405–1411. AAAI Press (2014). https://doi.org/10.1609/AAAI.V28I1.8884, https://doi.org/10.1609/aaai.v28i1.8884

[15] Lucet, M., Benabbou, N., Beynier, A., Maudet, N.: Multilevel fair allocation (2025), https://arxiv.org/abs/2512.24105

[16] Mallows, C.L.: Non-null ranking models. Biometrika **44**, 114–130 (1957)

[17] Scarlett, J., Teh, N., Zick, Y.: For one and all: Individual and group fairness in the allocation of indivisible goods. In: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems. p. 2466–2468. AAMAS '23, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2023)

[18] Schmidt-Kraepelin, U., Suksompong, W., Wijaya, S.: On multi-level apportionment. Theory and Decision **abs/2511.10000** (2025), https://doi.org/10.1007/s11238-025-10106-3