

Multilevel Fair Allocation under Additive Preferences

Maxime Lucet, Nawal Benabbou, Aurélie Beynier, and Nicolas Maudet

{firstname.lastname}@lip6.fr

LIP6, CNRS, Sorbonne Université, F-75005 Paris, France

Abstract. We study 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. Assuming that internal nodes’ utilities are the utilitarian welfare of their children, and the leaves have classical additive utilities over items, we first propose multilevel adaptations of usual envy-based fairness notions (e.g., WEF1). We present three adaptations and show that the choice among them is not neutral. We prove that, under identical preferences, the three adapted envy-based notions coincide, and that the Multilevel extension of Weighted Round Robin [11] (MWRR) guarantees them. We then show that under general preferences, MWRR may guarantee some notions while failing others. Finally, through experiments, we show that MWRR may still perform well even for adaptations it does not formally guarantee.

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

1 Introduction

Fair allocation has traditionally focused on designing algorithms to allocate fairly resources between individual agents [10, 16], and was then later extended to groups [4, 11, 15, 3, 14]. However, in many real-world settings, agents/groups actually belong to higher level entities, ultimately forming a hierarchical structure. Such hierarchy can be naturally modelled by a directed tree, where the root is responsible for allocating the bundle of items it receives to its children. This new setting, called multilevel fair allocation, was recently formalized in [17].

Multilevel fair allocation problems can arise in multiple real-world situations. For instance, consider a food charity operating in a country with multiple administrative levels: regional branches oversee departmental units, which in turn manage city-level distribution centres with potentially heterogeneous supply needs (see Fig. 1 for an illustration). 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. Note that many allocation problems across territories, or within a hierarchical structure, can be covered by this multilevel setting.

Interestingly, this setting can be useful even in contexts where no explicit hierarchy exists. For instance, when assigning medical resources to patients characterized

by severity, age, and risk level, one may first prioritize fairness across severity (e.g., giving higher priority to the most severe cases), then within severe cases place greater weight on high-risk patients, while for less severe cases one may instead balance allocation across age groups to avoid age bias (see Fig. 2 for an illustration). Such implicit hierarchy may also arise in affirmative action schemes.




Fig. 1: Hierarchical structure of a food charity.




Fig. 2: Example of implicit hierarchy.

In this paper, we assume that intermediary nodes act solely as representative bodies of their constituents (as in [17]), and therefore receive items only to allocate them to their children. These entities derive utility from the satisfaction of their constituents¹, and we adopt the utilitarian social welfare as an aggregation rule (following [17]), which is a standard modelling choice in collective decision-making. However, while [17] assume matroid-rank utilities at the leaves, we assume additive preferences. The additive utility model is the most studied preference model in the literature and is more expressive over singletons than matroid-rank utilities. Coming back to our illustrative example, this assumption is natural in the context of food charities (see e.g., [2]).

Moreover, we further depart from [17] by considering different notions of fairness. More precisely, they consider welfare-based fairness notions (e.g., Lorenz dominance) ensuring that, under matroid-rank utilities, a utilitarian-optimal solution is also fair. In contrast, we consider envy-based fairness notions. A central notion of fairness in classical fair division is *envy-freeness* (EF), together with its relaxations: *envy-freeness up to one good* (EF1) [10] 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. In the multilevel setting, envy-based notions require some adaptations. Indeed, an internal node does not observe how an alternative bundle would be allocated among its children, and thus cannot compute the exact utilitarian welfare it would derive from it, but only an estimate. In this paper, we consider three possible estimation functions (pessimistic, agnostic, and optimistic), and we investigate whether a multilevel version of Weighted Round Robin (WRR) [11] can guarantee some of our fairness notions together with completeness (i.e., all items are allocated). In monolevel settings, the Round Robin [10] and WRR are indeed simple algorithms achieving EF1 and WEF1, respectively.

Related work. *Group fairness* has become a highly active topic in fair division [4, 11, 15, 3, 14], and was recently extended to multilevel settings by [17]. While they focus on matroid-rank valuations and fairness notions unrelated to envy, we consider additive valuations and envy-based fairness notions in multilevel settings. However, some previous works can be seen as bilevel settings. [19] study individual and group envy-freeness when groups' valuations depend on those of their members, and propose al-

¹ Other assumptions could have been made, e.g., they could be oblivious of how the items are allocated to their children, or only care about the properties of the allocation to their children.

gorithms guaranteeing fairness at both levels. Our work extends this setting, as well as some of their algorithmic ideas. Closely related, [8] reconcile group and individual perspectives, allowing group preferences that do not necessarily aggregate those of their members (unlike our assumption). However, their model is limited to two levels, whereas ours allows an arbitrary number of levels. Finally, [2] study a bilevel allocation problem motivated by food charities, but focus on auction mechanisms, in contrast with our setting. Our multilevel model is also related to the multilevel apportionment framework of [20], which uses a similar tree-based hierarchy. However, their work focuses specifically on apportionment, a very specific subfield of resource allocation. Finally, the multilevel fairness notions proposed in this paper induce some visibility constraints, which is reminiscent of work on local fairness [1, 5, 7]; for instance, in our model, only labs affiliated to the same department can compare their situations.

Contributions. In this paper, we formally study multilevel fair allocation problems, assuming that internal nodes follow a utilitarian welfare objective and that leaves have additive preferences, and focusing on weighted envy-based fairness notions. We first show that WEF1 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 the pessimistic 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 agnostic fairness notion.

2 Model

In this paper, we consider an allocation problem where a set of goods $\mathcal{G} = \{g_1, \dots, g_m\}$ 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 $\mathcal{N} = \{1, \dots, n\}$ 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). We also assume that any node $i \in \mathcal{N}$ has a weight $w_i \in \mathbb{R}_{>0}$, which can be arbitrarily fixed, or can depend on the structure of the tree. 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 $(\mathcal{P}(i), i) \in \mathcal{E}$, and $\mathcal{P}(1) = \emptyset$. Moreover, let $\text{Anc}(i)$ 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 $\mathcal{I}(i) = \mathcal{N}_i \setminus \mathcal{L}(i)$. 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). $\pi : \mathcal{N} \rightarrow 2^{\mathcal{G}}$ is a multilevel allocation if it satisfies: (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 (ii) 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 $S \subseteq \mathcal{G}$, $A : N \rightarrow 2^S$ is a local allocation if (i) $\cup_{i \in N} A(i) \subseteq S$ and (ii) $A(i) \cap A(j) = \emptyset$ for all $i, j \in N$. Such a local allocation A is complete if condition (i) 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 $v_i : \Pi \rightarrow \mathbb{R}_{\geq 0}$ 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. 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 have no children and are here equipped with standard additive utilities: any leaf $x \in \mathcal{L}$ is associated with a utility function $u_x : 2^{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$ such that $v_x(\pi) = u_x(\pi(x))$ for any multilevel allocation $\pi \in \Pi$. Formally, $u_x(\emptyset) = 0$ and $u_x(S) = \sum_{g \in S} u_x(g)$ for any bundle $S \subseteq \mathcal{G}$, where $u_x(g)$ is a slight abuse of notation for $u_x(\{g\})$.

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, utility functions v_i , $i \in \mathcal{I}$, are defined over multilevel allocations rather than bundles of items. As a result, standard envy-based notions cannot be applied directly.

We therefore begin by introducing a method to evaluate bundles. As highlighted by [19], there are two ways to do so: (i) we can assume that i computes an allocation of its bundle S within $\mathcal{T}_i$, in which case its estimated utility is the sum of its children's (equivalently, leaves') utilities, or (ii) we can assume that i is agnostic to the allocation, and its estimated utility is the weighted average of S over the 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*

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

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*

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

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*

$$\bar{v}_i(S) = \sum_{x \in \mathcal{L}(i)} \sum_{g \in S} u_x(g) \times W(x, i) \text{ with } W(x, i) = \prod_{k \in \text{Anc}(x) \setminus \text{Anc}(i)} \frac{w_k}{w_{\mathcal{C}(\mathcal{P}(k))}}$$

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 leaf $x \in \mathcal{L}(i)$, $W(x, i)$ can be interpreted as the weight of x in the subtree $\mathcal{T}_i$. The quantity $W(i, j)$ is defined similarly for any node $i \in \mathcal{N}$ and ancestor $j \in \text{Anc}(i)$. Moreover, Appendix B shows that $\sum_{x \in \mathcal{L}(i)} W(x, i) = 1$.

Note that, for any leaf $x \in \mathcal{L}$, we have $\hat{v}_x(S) = \bar{v}_x(S) = \check{v}_x(S) = u_x(S)$ for any bundle $S \subseteq \mathcal{G}$, since leaf x is a tree with only one node (and thus $\mathcal{L}(x) = \{x\}$).

Fairness. Having defined the estimation functions, we can propose some multilevel extensions of some envy-based notions. This definition is parametrized to accommodate our three estimated utility functions.

Definition 7 (M[$\mathcal{E}$]-WEF). *Multilevel allocation $\pi \in \Pi$ is estimated Multilevel Weighted Envy-Free for $\mathcal{E} \in \{\text{pess}, \text{agno}, \text{opt}\}$ (denoted by M[$\mathcal{E}$]-WEF), if for any internal node $i \in \mathcal{I}$, and any pair of children $j, k \in \mathcal{C}(i)$, we have*

$$\frac{v_j(\pi)}{w_j} \geq \frac{\tilde{v}_j(\pi(k))}{w_k}$$

where $\tilde{v}$ denotes $\hat{v}$, $\bar{v}$ or $\check{v}$ for $\mathcal{E} = \text{pess}, \text{agno}, \text{or opt}$, respectively.

The definition of M[$\mathcal{E}$]-WEF1 is obtained by simply requiring $\frac{v_j(\pi)}{w_j} \geq \frac{\tilde{v}_j(\pi(k) \setminus \{g\})}{w_k}$ for some item $g \in \pi(k)$. For M[$\mathcal{E}$]-WEFX, this must hold for all items $g \in \pi(k)$.

Note that, if $\mathcal{T}$ has height $h(1)=1$, the problem reduces to a standard *monolevel* allocation setting, and our three estimated fairness notions coincide with the classical WEF, WEF1, and WEFX notions.

Example 1. We now illustrate the different estimated utility functions. Assume an instance with $\mathcal{N} = \{1, \dots, 7\}$ organized as in Fig 1, and $\mathcal{G} = \{g_1, \dots, g_5\}$. 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 node $i \in \mathcal{N}$ is $w_i = |\mathcal{L}(i)|$. Suppose we have the multilevel allocation $\pi \in \Pi$ such that $\pi(4) = \{g_1, g_5\}$, $\pi(5) = \{g_3\}$, $\pi(6) = \{g_2\}$, 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: $\hat{v}_3(\pi(2) \setminus \{g_1\}) = 4$ since all items are allocated to leaf 6 in a utilitarian-optimal allocation; $\bar{v}_3(\pi(2) \setminus \{g_1\}) = 3$; and finally, $\check{v}_3(\pi(2) \setminus \{g_1\}) = 2$ since all items are allocated to leaf 7 to minimize the utilitarian-welfare. 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.

3 Identical additive valuations

In [19], 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; for instance, in our food charity example, 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 1. Under both all-common and root-child-common valuations, for any internal node $i \in \mathcal{I} \setminus \{1\}$ and any bundle $S \subseteq \mathcal{G}$, we have:

$$\forall A, B \in \mathcal{A}_{\mathcal{L}(i)}^S, \sum_{x \in \mathcal{L}(i)} u_x(A(x)) = \sum_{x \in \mathcal{L}(i)} u_x(B(x)).$$

In other words, the allocation of S to the leaves in $\mathcal{L}(i)$ does not affect the utility derived by i . Consequently, the three estimated utility functions coincide, i.e., $\hat{v}_i(S) = \bar{v}_i(S) = \check{v}_i(S)$. Hence the definitions of M[$\mathcal{E}$]-WEF, for $\mathcal{E} \in \{\text{pess, agno, opt}\}$, also coincide.

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. We also write M-WEF instead of M[$\mathcal{E}$]-WEF for brevity.

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 [19]. 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 $i \in \text{Anc}(x)$ 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 $O(n)$. $\square$

We then show that we can obtain M-WEF1 jointly with some (monolevel) 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 in Section 3.2.

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

3.2 Root-child-common valuations

We now present a multilevel extension of the Weighted Round Robin (WRR) algorithm from [11, 19]. 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 (breaking ties lexicographically). If i is an internal node, 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. If the children of the considered node is a mix of internal nodes and leaves, we use (i). Pseudocode can be found in Algorithm 1.

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

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

Algorithm 1 Multilevel Weighted Round Robin (MWRR)

```

1: Input:  $\mathcal{T}$  - a multilevel tree ;  $\mathcal{G}$  - a set of items ; the leaves' valuations  $(u_x)_{x \in \mathcal{L}}$ 
2: Output:  $\pi$  - a multilevel allocation
3: Initialize  $\pi$  with empty bundle for any  $i \in \mathcal{N} \setminus \{1\}$  and  $\pi(1) = \mathcal{G}$ 
4: Initialize picking scores s.t.  $t_i = 0, \forall i \in \mathcal{N}$ 
5:  $\mathcal{RI} \leftarrow \mathcal{G}$  ▷ Initial set of remaining items
6: while  $\mathcal{RI} \neq \emptyset$  do
7:    $i = \arg \min_{i' \in \mathcal{C}(1)} \frac{t_{i'}}{w_{i'}}$ 
8:   while  $i \notin \mathcal{L}$  do
9:      $i = \arg \min_{i' \in \mathcal{C}(i)} \frac{t_{i'}}{w_{i'}}$ 
10:  end while
11:   $g = \arg \max_{g' \in \mathcal{RI}} u_i(g')$ 
12:  while  $i \neq 1$  do
13:     $\pi(i) \leftarrow \pi(i) \cup \{g\}$ 
14:     $t_i \leftarrow t_i + 1$ 
15:     $i \leftarrow \mathcal{P}(i)$ 
16:  end while
17:   $\mathcal{RI} \leftarrow \mathcal{RI} \setminus \{g\}$ 
18: end while
19: return  $\pi$ 

```

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 $O(m(n^2 + m))$.

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 over items: 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. 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 over items, we can conclude that MWRR is WEF1 w.r.t. $\mathcal{C}(1)$. Then, we can repeat the same argument recursively for any node $i \in \mathcal{N}$. $\square$

Corollary 1. *Under root-child-common valuations and $w_i = |\mathcal{L}(i)|$ for all $i \in \mathcal{N}$, 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 [19] can be extended to our setting and establishes EF1 among the leaves. The full proof can be found in Appendix A. $\square$

In this section, we showed that the algorithms proposed in [19] 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 case where we have general additive valuations.

4.1 Relations among $M[\mathcal{E}]$ -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 $\hat{v}_i(S) \geq \bar{v}_i(S)$.*

Proof. By linearity, $\bar{v}_i(S)$ and $\hat{v}_i(S)$ can be rewritten as follows:

$$\begin{aligned}\bar{v}_i(S) &= \sum_{x \in \mathcal{L}(i)} \sum_{g \in S} u_x(g) W(x, i) = \sum_{g \in S} \sum_{x \in \mathcal{L}(i)} u_x(g) W(x, i) \\ \hat{v}_i(S) &= \max_{A \in \mathcal{A}_{\mathcal{L}(i)}^S} \sum_{x \in \mathcal{L}(i)} \sum_{g \in A(x)} u_x(g) = \sum_{g \in S} \max_{x \in \mathcal{L}(i)} u_x(g)\end{aligned}$$

Note that, for any $g \in S$, both $\sum_{x \in \mathcal{L}(i)} u_x(g) W(x, i)$ and $\max_{x \in \mathcal{L}(i)} u_x(g)$ are convex combinations of $(u_x(g))_{x \in \mathcal{L}(i)}$; the latter places all the weight on some leaf $x^* \in \arg \max_{x \in \mathcal{L}(i)} u_x(g)$, while the former 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 all the weight on the largest element, we obtain:

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

Then, summing over $g \in S$ gives:

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

□

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

Proof. The proof is similar to that of Lemma 1: since $\check{v}_i(S)$ is obtained by placing all weight on the leaf with minimum utility for g , we have $\sum_{x \in \mathcal{L}(i)} u_x(g) W(x, i) \geq \min_{x \in \mathcal{L}(i)} u_x(g)$, and thus $\bar{v}_i(S) \geq \check{v}_i(S)$. □

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. $M[\text{opt}]\text{-WEF1} \Rightarrow M[\text{agno}]\text{-WEF1} \Rightarrow M[\text{pess}]\text{-WEF1}$.

In other words, for any allocation $\pi \in \Pi$, if π is not $M[\text{pess}]\text{-WEF1}$, then it is not $M[\text{agno}]\text{-WEF1}$; similarly, if it is not $M[\text{agno}]\text{-WEF1}$, then it is not $M[\text{opt}]\text{-WEF1}$. However, some allocations that are not $M[\text{opt}]\text{-WEF1}$ may still be $M[\text{agno}]\text{-WEF1}$ (e.g., the allocation in Example 1), and some that are not $M[\text{agno}]\text{-WEF1}$ may still be $M[\text{pess}]\text{-WEF1}$, as shown in the following example.

Example 2. 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 $\mathcal{N} = \{1, \dots, 9\}$ organized in the tree of Fig. 3. Leaves $x = 4, 5$ and 7 have utilities $u_x(g) = 2$ for all $g \in \mathcal{G}$; leaves $x = 6, 8$ and 9 have utilities $u_x(g) = 1$ for all $g \in \mathcal{G}$. Assume the weight of any 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 $\pi(4) = \{g_1\}$, $\pi(5) = \{g_3\}$, $\pi(6) = \{g_5\}$, $\pi(7) = \{g_2\}$, $\pi(8) = \{g_4\}$, and $\pi(9) = \emptyset$. For such an allocation π , we have $v_3(\pi(3)) = 3$, and for any $g \in \pi(2)$, $\bar{v}_3(\pi(2) \setminus \{g\}) = 10/3$, while $\bar{v}_3(\pi(2) \setminus \{g\}) = 2$. Hence, since $w_2 = w_3$, π is M[pess]-WEF1, but not M[agno]-WEF1.




Fig. 3: Tree of Example 2.

4.2 On the existence of M[$\mathcal{E}$]-WEF1 allocation

In this section, we show that MWRR guarantees M[pess]-WEF1 under general additive valuations. In contrast, M[agno]-WEF1 is harder to ensure: some instances admit no such allocation (and thus no M[opt]-WEF1 allocation either), and MWRR may fail to guarantee it even when one exists. Nevertheless, Section 4.3 shows that such failures remain 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 an M[pess]-WEF1 allocation.*

Proof sketch. The proof follows the same structure as the WRR proof of [11], with an additional lemma (Lemma 5, Appendix B) 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. During the first $|\mathcal{C}(1)|$ steps, each child of the root receives at most one item. Hence, the allocation is trivially M[pess]-WEF1 at this stage. Then, as in WRR, the picking rule ensures that for any $i, j \in \mathcal{C}(1)$, we have $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$. 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 M[pess]-WEF1 with respect to i . Applying the same argument recursively to each internal node concludes the proof.

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

Proof. Consider the multilevel instance represented in Fig. 4, and assume weights are $w_i = |\mathcal{L}(i)|$ for $i \in \mathcal{N}$. Let $\mathcal{G} = \{g_1, \dots, g_5\}$. Leaves have additive utilities: children

of 4 and 6 value every item at 2, while children of 5 and 7 value every item at 1. We show that no multilevel allocation can be $M[\text{agno}]\text{-WEF1}$.




Fig. 4: Tree of Example 3

First, notice that no multilevel allocation allocating all items to (leaves of) node 2 can be $M[\text{agno}]\text{-WEF1}$. Indeed, in such case, $v_3(\pi)/w_3 = 0 < \bar{v}_3(\mathcal{G} \setminus \{g\})/w_2 = 20/18$ for any item $g \in \mathcal{G}$. Conversely, no multilevel allocation allocating all items to leaves of node 3 can be $M[\text{agno}]\text{-WEF1}$, by symmetry of the instance.

Similarly, no multilevel allocation allocating exactly one item $g \in \mathcal{G}$ to node 3 can be $M[\text{agno}]\text{-WEF1}$. Indeed, even if this item was allocated to a leaf in $\mathcal{C}(6)$ (i.e. the group of leaves with the highest utilities), we would have $v_3(\pi)/w_3 = 2/6 < \bar{v}_3(\pi(2) \setminus \{g'\})/w_2 = 10/12$, where $\pi(2) = \mathcal{G} \setminus \{g\}$, and $g' \in \pi(2)$. The same applies when allocating exactly one item to node 2, and the rest to node 3, again by symmetry.

Finally, we show that even allocating 3 items to node 2, and 2 items to node 3 (or conversely) cannot yield a $M[\text{agno}]\text{-WEF1}$ allocation. Suppose a multilevel allocation where $|\pi(2)| = 3$ and $|\pi(3)| = 2$. Then, we show that no matter how $\pi(3)$ is allocated to its leaves, it must be agnostic weighted envious towards node 2. Indeed, assume the two items are allocated to leaves in $\mathcal{C}(6)$, then node 7 is agnostic weighted envious towards node 6, since it received no item, and node 6 received 2 items. Then, node 3 is obliged to allocate one item to node 6, and one item to node 7. But then, we have $v_3(\pi)/w_3 = 3/6 < \bar{v}_3(\pi(2) \setminus \{g\})/w_2 = 10/18$. By symmetry, no multilevel allocation allocating 2 items to node 2, and 3 items to node 3 can be $M[\text{agno}]\text{-WEF1}$. $\square$

This example shows that under general valuations, an $M[\text{agno}]\text{-WEF1}$ allocation may fail to exist. However, this is not the reason why MWRR fails: the algorithm may return an allocation that is not $M[\text{agno}]\text{-WEF1}$ even when such an allocation exists.

Theorem 6. *Under general valuations, MWRR does not guarantee an $M[\text{agno}]\text{-WEF1}$ allocation, even when it exists.*

Example 3. We consider the same instance as in the proof of Theorem 5, modifying only the leaves' preferences over singletons. We assume that all leaves in $\mathcal{C}(5)$ and $\mathcal{C}(7)$ have utility $u_x(g) = 1/5$ for any $g \in \mathcal{G}$. For $x \in \mathcal{C}(4)$, we set $u_x(g) = 5/21$ for $g \in \{g_1, g_2, g_4, g_5\}$ and $u_x(g) = 1/21$ for g_3 . Finally, for $x \in \mathcal{C}(6)$, we set $u_x(g) = 5/21$ for $g \in \{g_1, g_2, g_3, g_5\}$ and $u_x(g) = 1/21$ for g_4 .

For this instance, MWRR outputs 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 $\bar{v}_3(\pi(2) \setminus \{g_1\}) = 142/315 \simeq 0.451$. 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 allocation: $\pi(8) = \{g_2\}$, $\pi(12) = \{g_3\}$, $\pi(13) = \{g_4\}$, $\pi(14) = \{g_1\}$, $\pi(18) = \{g_5\}$.

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

Though these results are negative, one may still ask whether an α -approximation of M[agno]-WEF1 can be guaranteed.

Definition 8 (α -M[agno]-WEF1). *A multilevel allocation π is an α -approximation of M[agno]-WEF1, for some $\alpha > 0$, if, for any internal node $i \in \mathcal{I}$ and any children $j, k \in \mathcal{C}(i)$, there exists $g \in \pi(k)$ such that:*

$$v_j(\pi)/w_j \geq \alpha \cdot \bar{v}_j(\pi(k) \setminus \{g\})/w_k$$

Unfortunately, we show that MWRR cannot guarantee any such α -approximation.

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

Proof. Consider the tree in Fig. 5, and assume $w_i = |\mathcal{L}(i)|$ for all $i \in \mathcal{N}$. Let $\mathcal{G} = \{g_1, g_2, g_3\}$ be the set of items. Consider 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 $M > 0$ arbitrarily large.

Let π be the allocation returned by MWRR on this instance: $\pi(8) = \{g_2\}$, $\pi(10) = \{g_3\}$, $\pi(12) = \{g_1\}$. Note that we have $v_3(\pi) = 1$, and $\bar{v}_3(\pi(2) \setminus \{g\}) = W(14, 3) \times M + W(15, 3) \times M = \frac{M}{2}$ for any $g \in \pi(2)$. Hence, for π to be an α -approximation of M[agno]-WEF1, we need in particular $v_3(\pi)/w_3 \geq \alpha \cdot \bar{v}_3(\pi(2) \setminus \{g\})/w_2$. This requires $\alpha \leq 2/M$. Since M can be arbitrarily large, α can be arbitrarily small. $\square$




Fig. 5: Tree of Theorem 7

This result is particularly striking when compared to the $\frac{1}{3}$ -approximation of [19] in the bilevel setting. It highlights a fundamental gap between bilevel instances and multilevel trees of height at least three. While the bilevel case still admits an α -approximation for some $\alpha > 0$, this is no longer possible in the multilevel setting.

4.3 Experimental results

In this section, we run experiments to test how often MWRR fails to return a fair allocation. We show that, while it may frequently violate the optimistic notion, it is almost

always fair under the agnostic one (an experimental finding that significantly refines the negative result of Theorem 6). We describe hereafter the experimental protocol (hardware details are provided in Appendix C).

Trees. We consider three classes of trees in our experiments: (1) balanced binary trees, (ii) comb trees, and (iii) 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: (i) each node $i \in \mathcal{N}$ is assigned weight $w_i = |\mathcal{L}(i)|$, or (ii) the weight of each node is sampled independently and uniformly at random from the integer interval $[1, 6]$. In the tables below, we denote (i) by LW for leaf-count weights and (ii) by RW for random weights.

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, (ii) the resampling Dirichlet model recently introduced by [9], (iii) cost utilities [6], and (iv) correlated utilities as in [13]. 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, some tables are deferred to the appendix.

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 methods, all weighting schemes) is 0.0002 second, while that for $n = 127, m = 254$ is 0.0191 second. Moreover, the running time is consistent across different types of instances.

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

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 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 comb trees with large n and random weights are particularly challenging, with sometimes near 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).

n	m	Comb		Balanced		Part. unbal.	
		RW	LW	RW	LW	RW	LW
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 consider multilevel fair division problems where indivisible goods are allocated to agents organized in a tree-structured hierarchy. We assume that internal nodes have additive utilities over their children, while leaves have additive utilities over items. In this setting, we propose three envy-based fairness notions, depending on how a node estimates the value of a bundle, namely M[pess]-WEF1, M[agno]-WEF1, and M[opt]-WEF1. Under identical preferences, these notions coincide, and can be satisfied together with completeness using MWRR. Under general valuations, MWRR guarantees only M[pess]-WEF1 but still performs very well in practice for M[agno]-WEF1.

A direct extension concerns the existence of M[agno]-WEF1 allocations. In Section 4, we showed that under general additive valuations, such an allocation may not exist. However, this relies on non-normalized utilities: after normalization, the instance admits an M[agno]-WEF1 allocation, and existence under normalized utilities remains open. Still, this does not resolve the issue with MWRR, which may fail to return an M[agno]-WEF1 allocation even when one exists (as Example 3 is already normalized).

A natural extension of this work is to consider broader utility functions for leaves, and alternative utilities for internal nodes. In particular, instead of purely aggregating children's utilities, internal nodes could incorporate their own preferences, as in [8], e.g., preferring balanced allocations across subgroups.

Acknowledgments. This work is supported by the ANR project ANR-24-CE23-1700 AROMATICS.

References

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. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2017)
2. Aggarwal, G., Mertzaniadis, M., Psomas, A., Wang, D.: Mechanism design with delegated bidding (2024), <https://arxiv.org/abs/2409.19087>
3. Aleksandrov, M., Walsh, T.: Group envy freeness and group pareto efficiency in fair division with indivisible items. In: KI 2018: Advances in Artificial Intelligence (2018)
4. Benabbou, N., Chakraborty, M., Elkind, E., Zick, Y.: Fairness towards groups of agents in the allocation of indivisible items. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. p. 95–101. IJCAI’19, AAAI Press (2019)
5. Beynier, A., Chevaleyre, Y., Gourvès, L., Harutyunyan, A., Lesca, J., Maudet, N., Wilczynski, A.: Local envy-freeness in house allocation problems. *Autonomous Agents and Multi-Agent Systems* **33**(5), 591–627 (Sep 2019)
6. Botan, S., Ritossa, A., Suzuki, M., Walsh, T.: Maximin fair allocation of indivisible items under cost utilities. In: *Algorithmic Game Theory*. pp. 221–238 (2023)
7. Brederbeck, R., Kaczmarczyk, A., Niedermeier, R.: Envy-free allocations respecting social networks. *Artificial Intelligence* **305**, 103664 (2022)
8. Bu, X., Li, Z., Liu, S., Song, J., Tao, B.: Fair division with allocator’s preference. In: *Web and Internet Economics: 19th International Conference*. p. 77–94 (2024)
9. Böhm, P., Brederbeck, R., Gözl, P., Kaczmarczyk, A., Szufa, S.: Putting fair division on the map. *Proceedings of AAAI’2026* **40**(20), 16726–16734
10. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.* **7**(3), 12:1–12:32 (2019)
11. Chakraborty, M., Igarashi, A., Suksompong, W., Zick, Y.: Weighted envy-freeness in indivisible item allocation. *ACM Trans. Econ. Comput.* **9**(3) (Aug 2021)
12. Chakraborty, M., Segal-Halevi, E., Suksompong, W.: Weighted fairness notions for indivisible items revisited. *ACM Trans. Econ. Comput.* **12**(3) (Sep 2024)
13. Dickerson, J.P., Goldman, J.R., Karp, J., Procaccia, A.D., Sandholm, T.: The computational rise and fall of fairness. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. pp. 1405–1411 (2014)
14. Gross-Humbert, N., Benabbou, N., Beynier, A., Maudet, N.: On the notion of envy among groups of agents in house allocation problems. In: *ECAI 2023 - 26th European Conference on Artificial Intelligence*. pp. 924–931 (2023)
15. Kyropoulou, M., Suksompong, W., Voudouris, A.A.: Almost envy-freeness in group resource allocation. In: *Proceedings of IJCAI’2019*. pp. 400–406 (7 2019)
16. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Breese, J.S., Feigenbaum, J., Seltzer, M.I. (eds.) *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004)*, New York, NY, USA, May 17–20, 2004. pp. 125–131. ACM (2004). <https://doi.org/10.1145/988772.988792>, <https://doi.org/10.1145/988772.988792>
17. Lucet, M., Benabbou, N., Beynier, A., Maudet, N.: Multilevel fair allocation with matroid-rank preferences (2026), <https://arxiv.org/abs/2512.24105>
18. Mallows, C.L.: Non-null ranking models. *Biometrika* **44**, 114–130 (1957)

19. 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 (2023)
20. Schmidt-Kraepelin, U., Suksompong, W., Wijaya, S.: On multi-level apportionment. *Theory and Decision* (2025). <https://doi.org/https://doi.org/10.1007/s11238-025-10106-3>

A. Missing proofs from Section 3

We start by proving Theorem 2.

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 $\mathcal{L}$.*

Proof. The algorithm is called Sequential Maximin - Multilevel Weighted Round Robin, and is a multilevel extension of the SM-IWRR presented in [19]. Pseudocodes can be found in Algorithms 2 and 3. The pseudocode for MWRR can be found in the main paper.

Algorithm 2 Sequential Maximin (from [19])

```

1: Input:  $N$  - a set of agents ;  $\mathcal{G}$  - a set of items ;  $u$  - a common valuation
2: Output:  $A$  - an allocation
3: Initialize  $A(x) = \{\}$  for  $x \in N$ 
4:  $\mathcal{RI} \leftarrow \mathcal{G}$  ▷ Initial set of remaining items
5: while  $\mathcal{RI} \neq \emptyset$  do
6:    $g \leftarrow \arg \max_{g' \in \mathcal{RI}} u(g')$ 
7:    $x \leftarrow \arg \min_{y \in N} u(A(y))$ 
8:    $A(x) \leftarrow A(x) \cup \{g\}$ 
9:    $\mathcal{RI} \leftarrow \mathcal{RI} \setminus \{g\}$ 
10: end while
11: return  $A$ 

```

The proof the allocation is EFX at the leaves is exactly the same than in Theorem 3.5 of [19]. Moreover, we proved in Theorem 3 that the MWRR was M-WEF1 under root-child-common valuations. All-common valuations being a special case of root-child-common valuations, we get the same guarantee. Recall that under identical preferences, the three adaptations of M-WEF1 coincide. Hence, we focus on showing that the returned allocation is M-WEF1. We extend the proof of [19] to apply to the multilevel setting.

After the SM algorithm ran at the leaves, which computed an EFX allocation between the leaves, denote $A \in \mathcal{A}_{\mathcal{L}}^{\mathcal{G}}$ the resulting allocation at the leaves, and π the induced multilevel allocation, i.e. $\pi(x) = A(x)$ for any $x \in \mathcal{L}$, and $\pi(i) = \cup_{x \in \mathcal{L}(i)} A(x)$ for any $i \in \mathcal{I}$. Consider the set of bundles $\{A_1, \dots, A_{\ell}\}$, where $\ell = |\mathcal{L}|$. We relabel it so that $u(A_1) \geq u(A_2) \geq \dots \geq u(A_{\ell})$. Note that we denote the utilities of the leaves by u as they are identical. For any leaf $x \in \mathcal{L}$, we define the representative good value to be $\hat{u}(r_x) = u(A_x) - u(A_{\ell})$, where r_x is the representative good of bundle A_x . We make the following two claims:

Claim (1). For all $x \in \mathcal{L}$, $\hat{u}(r_x)$ is upper-bounded by the value of any good in A_x .

Claim (2). Given any node $i \in \mathcal{I}$, and any two of its children $j, k \in \mathcal{C}(i)$. Denote $\hat{\pi}$ the multilevel allocation of representative goods allocated to each node in $\mathcal{N}$, resulting

Algorithm 3 Sequential Maximin - Multilevel Weighted Round Robin (SM-MWRR)

```

1: Input:  $\mathcal{T}$  - a multilevel tree ;  $\mathcal{G}$  - a set of items ;  $u$  - a common valuation
2: Output:  $\pi$  - a multilevel allocation
3: Run the SM algorithm with input  $\mathcal{L}$ ,  $\mathcal{G}$ , and valuation  $u$ , and obtain allocation  $A' \in \mathcal{A}_{\mathcal{L}}^{\mathcal{G}}$ 
4:  $x_{min} \leftarrow \arg \min_{y \in \mathcal{L}(i)} u(A(y))$ 
5:  $A'_{min} \leftarrow A'(x_{min})$ 
6: Initialize a set of representative goods,  $\mathcal{RG} = \{\}$ 
7: for  $x \in \mathcal{L}$  do
8:   Create a representative good  $r_x$  such that  $\hat{u}(r_x) = u(A'(x)) - u(A'_{min})$ 
9:    $\mathcal{RG} \leftarrow \mathcal{RG} \cup \{r_x\}$ 
10: end for
11: Run the MWRR with input  $\mathcal{T}$ , items  $\mathcal{RG}$ , and identical valuation  $\hat{u}$ , and obtain the multilevel allocation  $\hat{\pi}$ 
12: Initialise multilevel allocation  $\pi$  such that  $\pi(1) = \mathcal{G}$  and  $\pi(i) = \{\}$  for any  $i \in \mathcal{N} \setminus \{1\}$ 
13: for  $i \in \mathcal{N} \setminus \{1\}$  do
14:   for  $x \in \mathcal{L}(i)$  do
15:     if  $r_x \in \hat{\pi}(i)$  then
16:        $\pi(i) \leftarrow \pi(i) \cup A'(x)$ 
17:     end if
18:   end for
19: end for
20: return  $\pi$ 

```

from allocation A to the leaves. If $\hat{\pi}$ is a M-WEF1 allocation of representative goods, then it remains M-WEF1 by replacing the representative goods by their corresponding bundles.

Claim 1 holds because the allocation A , hence π , is EFX between the leaves, and therefore, for any leaf $x \in \mathcal{L}$, and any item $g \in A_x$, we have $u(A_x \setminus \{g\}) \leq u(A_\ell)$. Then, as valuations are additive, $u(A_x) - u(\{g\}) \leq u(A_\ell)$, and hence $\hat{u}(r_x) = u(A_x) - u(A_\ell) \leq u(\{g\})$.

We now prove Claim 2. Assume we have a M-WEF1 allocation of the representative goods, and denote it $\hat{\pi}$. For any two children j, k of node $i \in \mathcal{I}$, there must exist a representative good $r_{\max} \in \hat{\pi}(k)$ such that $\hat{u}(r_{\max}) = \max_{y \in \mathcal{L}(k)} \hat{u}(r_y)$, and the following holds:

$$\frac{\sum_{x \in \mathcal{L}(j)} \hat{u}(r_x)}{w_j} \geq \frac{\sum_{y \in \mathcal{L}(k)} \hat{u}(r_y) - \hat{u}(r_{\max})}{w_k}$$

By the definition of a representative good, for any leaf $x \in \mathcal{L}(j)$, we have $\hat{u}(r_x) = u(A_x) - u(A_\ell)$, and since A_ℓ is the least-valued bundle of A , we get

$$\begin{aligned} \frac{\sum_{x \in \mathcal{L}(j)} \hat{u}(r_x)}{w_j} &= \frac{\sum_{x \in \mathcal{L}(j)} u(A_x) - u(A_\ell)}{w_j} \\ &= \frac{\sum_{x \in \mathcal{L}(j)} u(A_x)}{w_j} - u(A_\ell), \quad \text{since } w_j = |\mathcal{L}(j)| \end{aligned}$$

Moreover, the right-hand side, for the same reason, can be rewritten as :

$$\frac{\sum_{y \in \mathcal{L}(k)} \hat{u}(r_y) - \hat{u}(r_{\max})}{w_k} = \frac{\sum_{y \in \mathcal{L}(k)} u(A_y) - \hat{u}(r_{\max})}{w_k} - u(A_\ell)$$

Since in this setting, $u(\pi(j)) = \sum_{x \in \mathcal{L}(j)} u(A_x)$ and $u(\pi(k)) = \sum_{y \in \mathcal{L}(k)} u(A_y)$, we get

$$\begin{aligned} \frac{u(\pi(j))}{w_j} &= \frac{\sum_{x \in \mathcal{L}(j)} u(A_x)}{w_j} \geq \\ &\quad \frac{\sum_{y \in \mathcal{L}(k)} u(A_y) - \hat{u}(r_{\max})}{w_k} = \frac{u(\pi(k)) - \hat{u}(r_{\max})}{w_k}. \end{aligned}$$

Finally, from Claim 1, we know $\hat{u}(r_{\max}) \leq u(g_{\max})$, where $g_{\max}$ is some good from bundle $\pi(k)$ because by definition of a representative good, if $\hat{u}(r_{\max}) = \max_{y \in \mathcal{L}(k)} \hat{u}(r_y)$, then $u(g_{\max}) \geq u(g)$ for any $g \in \pi(k)$. Hence, we get

$$\frac{u(\pi(j))}{w_j} \geq \frac{u(\pi(k)) - \hat{u}(r_{\max})}{w_k} \geq \frac{u(\pi(k)) - u(g_{\max})}{w_k}$$

In SM-MWRR, we first compute an allocation to the leaves that is EFX between the leaves, then we construct the representative goods sets and values. Then, MWRR is run on the given tree, but with the representative goods to allocate and leaves equipped with the representative good utility function, $\hat{u}$. Since all leaves have the same utility function, we know from Theorem 3 that the returned allocation, denoted $\hat{\pi}$, is M-WEF1 w.r.t. $\hat{u}$. By Claim 2, we know this implies that the allocation w.r.t. the true bundles of items and utility functions u . This concludes the proof. $\square$

We then prove Corollary 1.

Corollary 1. *Under root-child-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 [19] readily extends to our setting and establishes EF1 among the leaves. Their argument proceeds as follows.

Consider a single execution of the algorithm and partition the resulting sequence of allocations into $\lceil \frac{m}{|\mathcal{L}|} \rceil$ rounds, where in each round every leaf receives exactly one item. At any given round, a leaf x prefers her bundle to that of any other leaf y who picks after her in the same round. Moreover, leaf x prefers the item she selects in a given round to the item selected in the next round by any leaf y who picked before her in the current round. Consequently, at the end of the algorithm, leaf x prefers her final bundle to that of any leaf y appearing after her in the picking order, and prefers her bundle to that of any leaf y appearing before her, up to the removal of the first item allocated to y . $\square$

B. Missing proofs from Section 4

We show that $\bar{v}_i(S)$ is indeed a convex combination of the $(u_x)_{x \in \mathcal{L}(i)}$ for any internal node $i \in \mathcal{I}$.

Lemma 3. *For any internal node $i \in \mathcal{I}$ and any bundle $S \subseteq \mathcal{G}$, the agnostic estimated utility function $\bar{v}_i(S)$ is a convex combination of the $(u_x)_{x \in \mathcal{L}(i)}$.*

Proof. First, recall that we have

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

We show that $\sum_{x \in \mathcal{L}(i)} W(x, i) = 1$.

$$\begin{aligned} \sum_{x \in \mathcal{L}(i)} W(x, i) &= \sum_{j \in \mathcal{C}(i)} W(j, i) \sum_{x \in \mathcal{L}(j)} W(x, j) \\ &= \sum_{j \in \mathcal{C}(i)} W(j, i) \dots \sum_{k \in \mathcal{C}(\mathcal{P}(k))} W(k, \mathcal{P}(k)) \underbrace{\sum_{x \in \mathcal{C}(k)} W(x, k)}_{=1} \\ &= \sum_{j \in \mathcal{C}(i)} W(j, i) \dots \underbrace{\sum_{k \in \mathcal{C}(\mathcal{P}(k))} W(k, \mathcal{P}(k))}_{=1} \cdot 1 \\ &= 1 \end{aligned}$$

□

We then provide the full proof of Theorem 4.

Theorem 4. *MWRR always returns a $M[\text{pess}]$ -WEF1 allocation.*

Proof. We show that MWRR computes an allocation $\pi \in \Pi$ such that $\pi|_{\mathcal{C}(p)}$ is pessimistic WEF1 for any $p \in \mathcal{I}$.

First, notice that for any p such that $h(p) = 1$, the proof follows immediately from the proof of [11] that their "least weight-adjusted frequent picker" procedure is WEF1. Indeed, at such node p the children in $\mathcal{C}(p)$ are leaves and hence can compare directly their bundle with that on their sibling. Hence, MWRR acts exactly the same way than their algorithm. The only difference is in the tie-breaking scheme, which does not intervene in their proof.

Hence, our objective is to prove that the result holds for any $p \in \mathcal{I}$ such that $h(p) \geq 2$. First, notice that the first $|\mathcal{C}(p)|$ picks are a simple round-robin, and each child receives one item. Hence, this first round is obviously pessimistic WEF1 as each child has at most one item at this point. Then, we focus on showing that the allocation computed remains WEF1 after the first pick.

Lemma 4 (From [11]). *Consider an internal node $p \in \mathcal{I}$ and one of its children $i \in \mathcal{C}(p)$ selected by MWRR at some iteration t , and suppose it is not i 's first pick. Let t_i and t_j be the numbers of times child i and some other child $j \in \mathcal{C}(p)$ appear in the prefix of iteration t (not including t itself). Then, $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$.*

Proof. Since i was picked at iteration t , it must be that $i \in \arg \min_{i' \in \mathcal{C}(p)} \frac{t_{i'}}{w_{i'}}$. Thus, for any $j \in \mathcal{C}(p)$, we have $\frac{t_i}{w_i} \leq \frac{t_j}{w_j}$. Hence, $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$. □

Lemma 5. Consider MWRR at some iteration t picked an internal node $i \in \mathcal{I}$, and $x \in \mathcal{L}(i)$ was eventually selected to pick an item among the remaining items, denoted $\mathcal{RG}(t)$. Assume x chooses $g \in \mathcal{RG}(t)$, then

$$u_x(g) \geq \min_{y \in \mathcal{L}(i)} u_y(g'), \quad \forall g' \in \mathcal{RG}(t)$$

Proof. Let $x \in \mathcal{L}(i)$ pick item $g \in \mathcal{RG}(t)$ among the remaining items at iteration t . Then, by construction, we have

$$g \in \arg \max_{g' \in \mathcal{RG}(t)} u_x(g')$$

Suppose by contradiction that $\exists g' \in \mathcal{RG}(t)$ such that

$$u_x(g) < \min_{y \in \mathcal{L}(i)} u_y(g')$$

Since $g \in \arg \max_{g' \in \mathcal{RG}(t)} u_x(g')$, it must be in particular that

$$u_x(g') < \min_{y \in \mathcal{L}(i)} u_y(g')$$

However since $x \in \mathcal{L}(i)$, it yields a contradiction. $\square$

We then proceed to show that, for any two children $i, j \in \mathcal{C}(p)$, MWRR computes an allocation such that j is pessimistic weighted envy-free up to the first chosen item by i . To show this, we keep the same proof of [11], and use our Lemma 5 at some point in their proof to be able to use the same argument adapted to our setting.

Lemma 6 (Mostly from [11]). Suppose that, for every iteration t in which agent i picks an item (i.e. one of its leaves eventually does) after her first pick, we have t_i and t_j for some other agent $j \in \mathcal{C}(p)$ satisfies $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$. Then, in the partial multilevel allocation up to and including i 's latest pick, agent j is pessimistic weighted envy-free towards i up to the first item picked.

Proof. Let $\gamma := \frac{w_j}{w_i}$. Consider any iteration t in which agent i is chosen after her first pick. Let agent j 's minimum values over its leaves in $\mathcal{L}(j)$ for the items allocated to agent i in the latter's second, third, ..., $(t_i + 1)$ st picks (the last one occurring at current iteration t be $\beta_1, \beta_2, \dots, \beta_{t_i}$ respectively. For instance, we have $\beta_1 = \min_{y \in \mathcal{L}(j)} u_y(g^2)$ where g^2 is the item selected at agent i 's second pick. Notice that β 's are different than those of [11].

If g^* is the first item picked by agent i (i.e. by one of its leaves) and π^t is the partial multilevel allocation up to and including iteration t , then clearly $\bigvee_j (\pi^t(i) \setminus \{g^*\}) = \sum_{x=1}^{t_i} \beta_{t_i}$.

Let the number of times agent j appears in the prefix of agent i 's second pick be τ_1 ; that between agent i 's second and third picks be τ_2 ; ...; that between agent i 's t_i^{th} and $(t_i + 1)^{\text{th}}$ picks be τ_{t_i} .

Let agent j 's values for the items she herself picked during phase $x \in [t_i]$ be $\alpha_1^x, \alpha_2^x, \dots, \alpha_{\tau_x}^x$ (i.e. between agent i 's x^{th} and $(x + 1)^{\text{th}}$ picks, agent j 's received τ_x items). Notice that here when we say "agent j 's values for the items", we actually mean that for each item, we call α_k^x the utility of the leaf that picked the item at iteration k during phase x . Then, we have $v_j(\pi^t) = \sum_{x=1}^{t_i} \sum_{y=1}^{\tau_x} \alpha_y^x$.

Let, for $r \in [t_i]$, $\sum_{x=1}^r \tau_x$ and r be the numbers of times agents j and i appear in the prefix of the latter's $(r+1)^{\text{th}}$ pick respectively. The condition of the lemma that $\frac{t_j}{t_i} \geq \frac{w_j}{w_i}$ yields

$$\sum_{x=1}^r \tau_x \geq r\gamma \quad (1)$$

Note that while $\tau_1 \geq \gamma > 0$ (because each agent is selected once in the first $\mathcal{C}(p)$ picks), it can be that $\tau_x = 0$ for $x \in \{2, 3, \dots, t_i\}$. It corresponds to the scenario where agent i picked more than once without agent j picking in between.

Moreover, we know from Lemma 5 that every time agent j was chosen, the item eventually picked by the selected leaf in $\mathcal{L}(j)$ has a value greater or equal than the worst utility over all leaves in $\mathcal{L}(j)$ for any of the remaining items at this iteration, including those eventually picked by agent i later. Hence, if $\tau_x > 0$ for some $x \in [t_i]$, we have

$$\alpha_y^x \geq \max\{\beta_x, \beta_{x+1}, \dots, \beta_{t_i}\}, \quad \forall y \in [\tau_x]$$

Note that this is where Lemma 5 is necessary to carry on with the same proof than [11]. Summing over all y 's, we get

$$\sum_{y=1}^{\tau_x} \alpha_y^x \geq \tau_x \max\{\beta_x, \beta_{x+1}, \dots, \beta_{t_i}\} \quad (2)$$

Note that Inequality (2) holds trivially for $\tau_x = 0$ since both sides are zero. Hence, it holds for any $x \in [t_i]$.

Now, we claim the following for each $r \in [t_i]$

$$\sum_{x=1}^r \sum_{y=1}^{\tau_x} \alpha_y^x \geq \gamma \sum_{x=1}^r \beta_x + \left(\sum_{x=1}^r \tau_x - r\gamma \right) \max\{\beta_r, \beta_{r+1}, \dots, \beta_{t_i}\}$$

To prove the claim, we proceed by induction on r . For the base case $r = 1$, we obtain from Inequality (2)

$$\begin{aligned} \sum_{y=1}^{\tau_1} \alpha_y^1 &\geq \tau_1 \max\{\beta_1, \beta_2, \dots, \beta_{t_i}\} \\ &\geq \gamma \beta_1 + (\tau_1 - \gamma) \max\{\beta_1, \beta_2, \dots, \beta_{t_i}\} \end{aligned}$$

where second line follows from $\gamma \max\{\beta_1, \beta_2, \dots, \beta_{t_i}\} \geq \gamma \beta_1$.

For the inductive step, assume the claim holds for $r - 1$. We now prove it for r .

$$\begin{aligned}
\sum_{x=1}^r \sum_{y=1}^{\tau_x} \alpha_y^x &= \sum_{x=1}^{r-1} \sum_{y=1}^{\tau_x} \alpha_y^x + \sum_{y=1}^{\tau_r} \alpha_y^r \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_{r-1}, \beta_r, \dots, \beta_{t_i}\} + \sum_{y=1}^{\tau_r} \alpha_y^r \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_{r-1}, \beta_r, \dots, \beta_{t_i}\} + \tau_r \max\{\beta_r, \dots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \right) \max\{\beta_r, \dots, \beta_{t_i}\} + \tau_r \max\{\beta_r, \dots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^r \tau_x - (r-1)\gamma \right) \max\{\beta_r, \dots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \gamma \max\{\beta_r, \dots, \beta_{t_i}\} + \left(\sum_{x=1}^r \tau_x - r\gamma \right) \max\{\beta_r, \dots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \gamma \beta_r + \left(\sum_{x=1}^r \tau_x - r\gamma \right) \max\{\beta_r, \dots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^r \beta_x + \left(\sum_{x=1}^r \tau_x - r\gamma \right) \max\{\beta_r, \dots, \beta_{t_i}\}
\end{aligned}$$

where second line follows from the inductive hypothesis. The third line comes from Inequality (2). The forth line follows from $\sum_{x=1}^{r-1} \tau_x - (r-1)\gamma \geq 0$ from Inequality (1) and $\max\{\beta_{r-1}, \beta_r, \dots, \beta_{t_i}\} \geq \max\{\beta_r, \dots, \beta_{t_i}\}$.

This completes the induction.

Now, let $r = t_i$, we get

$$\begin{aligned}
\sum_{x=1}^{t_i} \sum_{y=1}^{\tau_{t_x}} \alpha_y^x &\geq \gamma \sum_{x=1}^{t_i} \beta_x + \left(\sum_{x=1}^{t_i} \tau_x - t_i\gamma \right) \beta_{t_i} \\
&\geq \gamma \sum_{x=1}^{t_i} \beta_x
\end{aligned}$$

where second line follows from Inequality (1).

Hence, as $v_j(\pi) = \sum_{x=1}^{t_i} \sum_{y=1}^{\tau_x} \alpha_y^x$ and $\check{v}_j(\pi(i) \setminus \{g^*\}) = \sum_{x=1}^{t_i} \beta_x$, we get

$$\begin{aligned}
v_j(\pi) &\geq \gamma \cdot \check{v}_j(\pi(i) \setminus \{g^*\}) \\
&\Leftrightarrow \frac{v_j(\pi)}{w_j} \geq \frac{\check{v}_j(\pi(i) \setminus \{g^*\})}{w_i}
\end{aligned}$$

Hence, agent j is pessimistic weighted envy-free towards i up to the first item agent i picked. $\square$

Hence, since Lemma 6 holds for any $p \in \mathcal{I}$ and any of its children, it concludes the proof. $\square$

C. Additional experimental results

We provide hereafter further details on the experimental results we obtain.

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.

Preference generation. In order to design a robust experimental protocol, we use four different methods to generate the preferences of our leaves.

Mallows. We adapt the Mallows model [18], commonly used in voting theory, to a multilevel setting. Although it generates ordinal preferences, it remains useful for fair allocation when we want similar rankings over items across agents. The model relies on a central ranking and a dispersion parameter $\phi \in [0, 1]$: when $\phi = 0$, rankings match the central ranking exactly, while $\phi = 1$ yields uniformly random rankings (Impartial Culture).

In the multilevel version, each internal node i is assigned its own dispersion parameter ϕ_i . Starting from a uniformly sampled root ranking, we recursively generate central rankings for each node: the children of node i receive rankings drawn from a Mallows model centered at i 's ranking with parameter ϕ_i . This process continues down the tree until all leaves are assigned rankings, which we convert into (non-normalized) utilities using the `pref_voting` library.

This framework enables fine control over preference correlations across the hierarchy: low ϕ_i enforces similarity among siblings, while higher values introduce greater diversity.

Resampling-Dirichlet. The second method we use was introduced in [9]. Our multilevel adaptation works as follows: we first generate randomly a central approval vector V , where $m * p$ items are approved, with $p \in \{0.3, 0.6, 0.9\}$ a probability. Then, starting at the root, we generate approval vector for each of its children – if item g is approved in V , then child i approves it with probability p , and if g is not approved, then i approves it with probability $\phi \in \{0.2, 0.8\}$. We repeat this procedure until reaching the leaves.

For each leaf x , we construct a utility vector u_x . Let App be the ordered list of items approved by x . For a parameter $t \geq 0$, we assign utilities $u_x(g) = 10^{-10}$ for non-approved items, and $u_x(g) = \frac{2}{(j/|App|+0.01)^t}$ for the j -th approved item. If no item is approved, one is selected uniformly at random. Finally, we sample a normalized utility vector from a Dirichlet distribution parametrized by this vector.

Cost utilities. The third method, proposed in [6], assumes a utility vector describing the public value of each item (drawn uniformly at random). Each agent can either approve or disapprove an item. If she approves it, then her utility for the item is the public utility; otherwise, she has no utility for it. We adapt it to the multilevel setting by computing from the root to the leaves, an approval vector for each node i as follows: for any item $g \in \mathcal{G}$, if $\mathcal{P}(i)$ approves it, then i approves with probability $p \in \{0.3, 0.6, 0.9\}$, while if $\mathcal{P}(i)$ disapproves it, then i approves it with probability $\phi \in \{0.2, 0.8\}$.

Correlated utilities. The last method, proposed in [13], assumes again a central utility vector V , drawn uniformly at random. The utility of any agent x for any item g is then drawn from a normal distribution whose mean is $V(g)$, i.e. $u_x(g) \sim \mathcal{N}(V(g), \sigma_x)$ with $\sigma_x \in \{0.2, 0.8\}$ the standard deviation associated with agent x . Our adaptation to the multilevel setting follows the same logic: we draw uniformly at random a central vector. Starting at the root, we sample some utility vector for the children of the root, i.e. for each $i \in \mathcal{C}(1)$ and each item $g \in \mathcal{G}$, $u_i(g) \sim \mathcal{N}(V(g), \sigma_1)$. We repeat this process until reaching the leaves.

Results. In the main body of the paper, we reported our experimental results and observed that, among the 192,000 generated instances, fewer than 50 led MWRR to return a non-M[agno]-WEF1 allocation. In this section, we provide the detailed tables corresponding to those instances. We observe that, across all preference generation methods, non-M[agno]-WEF1 allocations occur only for instances defined on comb trees with random weights (See Tables 3, 4, 5, and 6). Note, however, that such unfair allocations are not limited to this class of instances; they also arise in other settings (see, for example, the proof of Theorem 6).

Table 3: Proportion of non-M[agno]-WEF1 allocations (95% CI) - Mallows - RW - Comb tree. Columns: (φ_1, φ_2) .

N	M	(2, 2)	(2, 8)	(8, 2)	(8, 8)
15	15	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
15	30	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	31	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	62	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.01]
63	63	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
63	126	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	127	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	254	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]

Table 4: Proportion of non-M[agno]-WEF1 allocations (95% CI) - resampling dirichlet - RW - Comb tree. Columns: (p, φ) .

N	M	(3, 2)	(3, 8)	(6, 2)	(6, 8)	(9, 2)	(9, 8)
15	15	[0.00, 0.05]	[0.00, 0.01]	[0.00, 0.02]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]
15	30	[0.00, 0.02]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	31	[0.00, 0.02]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]
31	62	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
63	63	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]
63	126	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	127	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	254	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]

Table 5: Proportion of non-M[agno]-WEF1 allocations (95% CI) - Cost utilities - RW - Comb tree. Columns: (p, φ) .

N	M	(3, 2)	(3, 8)	(6, 2)	(6, 8)	(9, 2)	(9, 8)
15	15	[0.00, 0.04]	[0.00, 0.01]	[0.00, 0.01]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]
15	30	[0.00, 0.03]	[0.00, 0.01]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	31	[0.00, 0.04]	[0.00, 0.01]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	62	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
63	63	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
63	126	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	127	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	254	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]

Table 6: Proportion of non-M[agno]-WEF1 allocations (95% CI) - Correlated preferences - RW - Comb tree. Columns: (φ_1, φ_2) .

N	M	(2, 2)	(2, 8)	(8, 2)	(8, 8)
15	15	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
15	30	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	31	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
31	62	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
63	63	[0.00, 0.00]	[0.00, 0.01]	[0.00, 0.00]	[0.00, 0.00]
63	126	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	127	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]
127	254	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]	[0.00, 0.00]