Best of Both Worlds: Agents with Entitlements

Martin Hoefer
Goethe University
Frankfurt, Germany
mhoefer@em.uni-frankfurt.de

Marco Schmalhofer
Goethe University
Frankfurt, Germany
schmalhofer@em.uni-frankfurt.de

Giovanna Varricchio
Goethe University
Frankfurt, Germany
varricchio@em.uni-frankfurt.de

ABSTRACT

Fair division of indivisible goods is a central challenge in artificial intelligence. For many prominent fairness criteria including *envy-freeness* (EF) or *proportionality* (PROP), no allocations satisfying these criteria might exist. Two popular remedies to this problem are randomization or relaxation of fairness concepts. A timely research direction is to combine the advantages of both, commonly referred to as *Best of Both Worlds* (BoBW).

We consider fair division *with entitlements*, which allows to adjust notions of fairness to heterogeneous priorities among agents. This is an important generalization to standard fair division models and is not well-understood in terms of BoBW results. Our main result is a lottery for additive valuations and different entitlements that is ex-ante *weighted envy-free* (WEF), as well as ex-post *weighted proportional up to one good* (WPROP1) and *weighted transfer envy-free up to one good* (WEF(1, 1)). It can be computed in strongly polynomial time. We show that this result is tight – ex-ante WEF is incompatible with any stronger ex-post WEF relaxation.

In addition, we extend BoBW results on group fairness to entitlements and explore generalizations of our results to instances with more expressive valuation functions.

KEYWORDS

Fair Division; Best of Both Worlds; Entitled Agents; Random Allocation

ACM Reference Format:
Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. 2023. Best of Both Worlds: Agents with Entitlements. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

Fair division of a set of indivisible goods is a prominent challenge at the intersection of economics and computer science. It has attracted a lot of attention over the last decades due to many applications in both simple and complex real-world scenarios. Formally, we face an allocation problem with finite sets $\mathcal{N}$ of n agents and $\mathcal{G}$ of m goods. Each agent $i \in \mathcal{N}$ has a valuation function $v_i : 2^{\mathcal{G}} \to \mathbb{R}_{\geq 0}$. The goal is to compute a "fair" allocation $\mathcal{A} = (A_1, \dots, A_n)$, i.e., a fair partition of the goods among the agents.

What is fair can certainly be a matter of debate. For this reason, several fairness criteria have been introduced and studied. *Envy-freeness* (EF) is probably one of the most intuitive concepts – it postulates that once goods are allocated no agent strictly prefers goods received by any other agent, i.e., $v_i(A_i) \geq v_i(A_j)$ for all $i, j \in \mathcal{N}$. EF is a comparison-based notion. In contrast, there are also threshold-based ones such as *Proportionality* (PROP): $\mathcal{A}$ is proportional if every agent receives a bundle whose value is at least her proportional share, i.e., $v_i(A_i) \geq v_i(\mathcal{G})/n$ for every $i \in \mathcal{N}$.

Unfortunately, for indivisible goods, neither PROP- nor EF-allocations may exist. Two natural conceptual remedies to this non-existence problem are (1) randomization or (2) relaxation of fairness concepts. Towards (1), a random allocation that is EF in expectation always exists (for every set of valuation functions): Select an agent uniformly at random and give the entire set of goods $\mathcal{G}$ to her. Then, however, every realization in the support is highly unfair – there is always an agent who receives everything, while all others get nothing. Moreover, it is easy to see that such an allocation might not even be Pareto-optimal. Towards (2), a well-known relaxation of EF is *envy-freeness up to one good* (EF1) [10, 22]: Every agent shall value her own bundle at least as much as any other agent's bundle after removing some good from the latter, i.e., for every $i, j \in \mathcal{N}$ there is $g \in \mathcal{G}$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$. Whenever the valuations of the agents are monotone, an EF1 allocation always exists and can be computed in polynomial time [22]. However, different EF1 allocations may advantage different agents. Similarly to EF1, *proportionality up to one good* (PROP1) has also been studied [16].

A timely research direction is to combine the advantages of both randomization and relaxation, commonly referred to as *Best of Both Worlds* (BoBW) results. An important result was obtained by both Aziz [2] and Freeman et al. [19] for additive valuations – a lottery over deterministic allocations that is EF in expectation (ex-ante) and EF1 for every allocation in the support (ex-post). Moreover, the lottery can be computed in polynomial time. Both papers generalize the Probabilistic Serial (PS) rule [9] for the matching case, when there are n agents and $m = n$ goods. PS is ex-ante EF. By the Birkhoff-von Neumann decomposition, it can be represented as a lottery over polynomially many deterministic allocations. Furthermore, any allocation in the support assigns to each agent exactly one good. This implies ex-post EF1. Both [2, 19] generalize the application of the Birkhoff-von Neumann decomposition to instances with arbitrarily many goods.

In our work, we consider a more general framework to allow more flexibility in the definition of fairness. Concepts like EF or PROP imply that all agents are symmetric, i.e., they are ideally treated as equals. In many scenarios, however, there is an inherent asymmetry in the agent population. Alternatively, it can be beneficial for an allocation mechanism to have the option to reward certain agents. We follow the formal framework of *entitlements* [4, 13] that enables increased expressiveness. Formally, each agent $i \in \mathcal{N}$ now has a weight, or priority $w_i > 0$. Fairness notions like EF or PROP are then refined based on these weights (see Section 2 for formal definitions). Generally, we will use a prefix "W" to refer to a fairness concept in the context of entitlements.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

1.1 Our Contribution

We study lotteries having both ex-ante and ex-post guarantees for fair division with additive valuations and different entitlements. We provide a lottery that is ex-ante *weighted stochastic-dominance envy-free* (WSD-EF) and consequently ex-ante WEF. Differently from [2, 19], we make use of a stronger decomposition theorem by Budish et al. [11] and show it is possible to achieve ex-post WPROP1 and WEF$(1, 1)$. The latter implies that in every allocation $\mathcal{A}$ in the support, weighted envy from agent i to j can be eliminated by moving entirely one good from A_j to A_i. All our constructions can be carried out in strongly polynomial time. Perhaps surprisingly, this result is tight – we show that ex-ante WEF is incompatible with any stronger ex-post WEF notion. Therefore, a direct extension of [2, 19] to a lottery with ex-ante WEF and ex-post WEF1 is impossible.

Freeman et al. [19] investigate further combinations of ex-ante and ex-post properties; namely, they provide a lottery that is ex-ante *group fair* (GF) as well as ex-post PROP1 and EF_1^1. In an EF_1^1-allocation $\mathcal{A}$, we can eliminate envy from i to j when we remove one good from A_j and add one good to A_i; differently from EF$(1, 1)$, the good added to A_i is *not required* to come from A_j. We prove that this result can be adapted to hold also with entitlements.

Finally, we expand the scope of BoBW towards more general valuations. For equal entitlements, ex-ante EF and ex-post EF1 is possible in more general cases. For different entitlements, ex-ante WEF and ex-post WEF$(1, 1)$ or WPROP1 are no longer compatible (even for two agents, one additive and one unit-demand). For this reason, we focus on threshold-based guarantees – we show that it is possible to compute in polynomial time a lottery that is ex-ante WPROP and ex-post WPROP1, even for XOS valuations.

Due to space limits, all missing proofs and examples are deferred to [21].

1.2 Related Work

Fair division attracted an enormous amount of attention, and there is a large number of surveys. We refer to a rather recent one by Amanatidis et al. [1] and restrict attention to more directly related works.

Other than envy-freeness [2, 19], the Max-Min-Share (MMS) is studied by Babaioff et al. [6] in the BoBW framework: The authors design a lottery simultaneously achieving ex-ante PROP and ex-post PROP1 + $\frac{1}{2}$-MMS.

When agents are endowed with ordinal preferences rather than cardinal valuation functions, stochastic-dominance envy-freeness is the most prominent fairness notion for lotteries. It was first considered by Bogomolnaia and Moulin [9] and later systematically studied by Aziz et al. [3].

An orthogonal direction is pursued by Caragiannis et al. [12] by introducing interim EF, a trade-off between ex-ante and ex-post EF.

For fair division with entitlements, the literature has focused on characterizing picking sequences guaranteeing fairness properties [13–15], the problem of maximizing Nash social welfare [20, 23], and introducing appropriate shares [5, 18].

2 PRELIMINARIES

A fair division instance $\mathcal{I}$ is given by a triple $(\mathcal{N}, \mathcal{G}, \{v_i\}_{i\in\mathcal{N}})$, where $\mathcal{N}$ is a set of n agents and $\mathcal{G}$ is a set of m indivisible goods. Every agent $i \in \mathcal{N}$ has a valuation function $v_i : 2^{\mathcal{G}} \to \mathbb{R}_{\geq 0}$, where $v_i(A)$ represents the value, or utility, of i for the bundle $A \subseteq \mathcal{G}$. We assume that valuations are monotone ($v(A) \leq v(B)$ for $A \subseteq B$) and normalized ($v(\emptyset) = 0$). For each $i \in \mathcal{N}$ and $g \in \mathcal{G}$, $v_i(g) \geq 0$ represents the value i assigns to the good g. A valuation function v_i is said to be *additive* if $v_i(A) = \sum_{g\in A} v_i(g)$.

In what follows, ties are broken according to a fixed ordering of $\mathcal{G}$. This serves to avoid technical and tedious tie-breaking issues.

Entitlements. We study fair division with entitlements. Each agent $i \in \mathcal{N}$ is endowed with an *entitlement* or *weight* $w_i > 0$. For convenience, we assume w.l.o.g. $\sum_{i\in\mathcal{N}} w_i = 1$. We say that agents have *equal entitlements* if $w_i = \frac{1}{n}$, for all $i \in \mathcal{N}$, and refer to this as the *unweighted setting*.

We now provide a simple example of a fair division instance with entitled agent; this example will be used in the rest of the paper to explain our approach.

Example 1 (A fair division instance with entitlements). *We outline an instance $\mathcal{I}^*$ given by $(\mathcal{N}, \mathcal{G}, \{v_i\}_{i\in\mathcal{N}})$ and entitlements w. The agents are $\mathcal{N} = \{1, 2, 3\}$, the goods $\mathcal{G} = \{g_1, g_2, g_3, g_4\}$, and $w_1 = \frac{1}{2}$, $w_2 = \frac{1}{3}$ and $w_3 = \frac{1}{6}$ is the entitlement of agent 1, 2 and 3, respectively. The valuation functions are additive with values:*

g	g_1	g_2	g_3	g_4
$v_1(g)$	8	8	5	2
$v_2(g)$	3	5	4	1
$v_3(g)$	4	7	6	2

Throughout the paper, whenever we use $\mathcal{I}^*$, we mean the instance we just described. ■

2.1 Weighted Fairness Notions

An allocation $\mathcal{A} = (A_1, \ldots, A_n)$ is a partition of $\mathcal{G}$ among the agents, where $A_i \cap A_j = \emptyset$, for each $i \neq j$, and $\bigcup_{i\in\mathcal{N}} A_i = \mathcal{G}$. An allocation $\mathcal{A}$ is *weighted proportional* (WPROP) if, for each i, $v_i(A_i) \geq w_i \cdot v_i(\mathcal{G})$ and *weighted envy-free* (WEF) if, for each i, j,

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j)}{w_j}.$$

Since goods are indivisible such allocations may not always exist, and relaxed versions have been defined. An allocation $\mathcal{A}$ is *weighted proportional up to one good* (WPROP1) if for each $i \in \mathcal{N}$ there exists $g \in \mathcal{G}$ such that $v_i(A_i \cup \{g\}) \geq w_i \cdot v_i(\mathcal{G})$. Note that, for additive valuations, WEF $\Rightarrow$ WPROP but, differently from equal entitlements, WEF1 $\not\Rightarrow$ WPROP1. Concerning envy-freeness, we have already discussed EF and EF1 in the introduction. We here work with a broader definition that generalizes these notions.

Definition 1 (WEF(x, y)). *For $x, y \in [0, 1]$, an allocation $\mathcal{A}$ is called* WEF(x, y) *if for each $i, j \in \mathcal{N}$ either $A_j = \emptyset$ or there exists $g \in A_j$ such that*

$$\frac{v_i(A_i) + y \cdot v_i(g)}{w_i} \geq \frac{v_i(A_j) - x \cdot v_i(g)}{w_j}.$$

The definition of WEF(x, y), introduced in [15], is meaningful mostly for additive valuations. For general valuations, the idea of

$\mathsf{WEF}(1,1)$ can be expressed by $w_j \cdot v_i(A_i \cup \{g\}) \geq w_i \cdot v_i(A_j \setminus \{g\})$; analogously for $\mathsf{WEF}(0,1)$ and $\mathsf{WEF}(1,0)$. Conceptually, $\mathsf{WEF}(1,0)$ coincides with a notion of *weighted envy-freeness up to one good* (WEF1). $\mathsf{WEF}(1,1)$ has also been called *weighted transfer envy-freeness up to one good* [13]. In $\mathsf{WEF}(1,1)$ the good g added to A_i must come from A_j: Assuming that g may come from any other bundle leads to the following (weaker) notion introduced in [7].

Definition 2 (WEF^1_1). *An allocation $\mathcal{A}$ is called* weighted envy-free up to one good more and less (WEF^1_1) *if for each $i, j \in \mathcal{N}$ either $A_j = \emptyset$ or there exist $g_i, g_j \in \mathcal{G}$ such that $w_j \cdot v_i(A_i \cup \{g_i\}) \geq w_i \cdot v_i(A_j \setminus \{g_j\})$.*

We move on to fairness concepts for fractional allocations. A *fractional allocation* $X = (x_{ig})_{i \in \mathcal{N}, g \in \mathcal{G}} \in [0,1]^{n \times m}$ specifies the fraction of good g that agent i receives. We assume fractional allocations are complete, i.e., $\sum_{i \in \mathcal{N}} x_{ig} = 1$ for every $g \in \mathcal{G}$.

Group fairness was first introduced in [17] and extended to fractional allocations by [19]. Towards extending group fairness for fractional allocations to weighted agents, consider a subset of agents $S \subseteq \mathcal{N}$. We define $w_S = \sum_{i \in S} w_i$ as the weight of the set, and $\cup_{j \in S} X_j := \left(\sum_{j \in S} x_{jg}\right)_{g \in \mathcal{G}}$ as the total fractions of each good $g \in \mathcal{G}$ assigned to the agents of S.

Definition 3 (WGF). *A fractional allocation X is* weighted group fair (WGF) *if for all non-empty subsets of agents $S, T \subseteq \mathcal{N}$, there is no fractional allocation X' of $\cup_{j \in T} X_j$ to the agents in S such that $w_S \cdot v_i(X'_i) \geq w_T \cdot v_i(X_i)$, for all $i \in S$ and at least one inequality is strict.*

Similarly to the unweighted setting, weighted group fairness implies other (weighted) envy and efficiency notions, for example, WEF (if $|S| = |T| = 1$), WPROP (if $|S| = 1, T = \mathcal{N}$), and Pareto-optimality (if $S = T = \mathcal{N}$).

We finally focus on stochastic dominance, a standard fairness notion for random allocations. For convenience, we here define it using fractional allocations. Given any $i \in \mathcal{N}$, let us denote by X_i and X'_i the fractional bundles of agent i in the allocations X and X', respectively. Agent i SD prefers X_i to X'_i, written $X_i \succeq_i^{\mathsf{SD}} X'_i$, if for any $g^* \in \mathcal{G}$

$$\sum_{g : v_i(g) \geq v_i(g^*)} x_{ig} \geq \sum_{g : v_i(g) \geq v_i(g^*)} x'_{ig} \ ,$$

where x_{ig} and x'_{ig} represents the fraction of g that i owns in the two fractional bundles.

We say $X_i \succ_i^{\mathsf{SD}} X'_i$, if $X_i \succeq_i^{\mathsf{SD}} X'_i$ and not $X'_i \succeq_i^{\mathsf{SD}} X_i$. Notice that $\{g \mid v_i(g) \geq v_i(g^*)\}$ is the set of goods that i likes at least as much as g^*. Although we defined it by means of v_i, this set only depends on the relative ordering of the goods and not on the valuation v_i.

Definition 4 (SD-EF and WSD-EF). *A random allocation X is* SD-envy-free (SD-EF) *if for all $i, j \in \mathcal{N}, X_i \succeq_i^{\mathsf{SD}} X_j$. Similarly, we say X is* WSD-envy-free (WSD-EF) *if for all $i, j \in \mathcal{N}, w_j \cdot X_i \succeq_i^{\mathsf{SD}} w_i \cdot X_j$.*

2.2 Deterministic Algorithms and Picking Sequences

For additive valuations, a straightforward round-robin algorithm yields an EF1 allocation. Clearly, when agents have different entitlements, the round-robin algorithm might no longer provide a WEF1 allocation. Different entitlements impose different priorities among agents, which has resulted in the consideration of picking sequences.

A *picking sequence* for n agents and m goods is a sequence $\pi = (i_1, \ldots, i_m)$, where $i_h \in \mathcal{N}$, for $h = 1, \ldots, m$. An allocation $\mathcal{A}$ is the result of the picking sequence π if it is the output of the following procedure: Initially every bundle is empty; then, at time step h, i_h inserts in her bundle her most preferred good among the available ones. Once a good is selected, it is removed from the set of available goods.

For our purposes, we will rely on the following characterization for $\mathsf{WEF}(x,y)$ (in the context of additive valuations).

Proposition 1. *Let t_i, t_j be the number of picks of agents i, j, respectively, in a prefix of π. A picking sequence π is $\mathsf{WEF}(x,y)$ if and only if for every prefix of π and every pair of agents i, j, we have $\frac{t_i + y}{w_i} \geq \frac{t_j - x}{w_j}$.*

Chakraborty et al. [15] prove this proposition using the assumption $x + y = 1$, since $\mathsf{WEF}(x,y)$ allocations might not exist for $x + y < 1$. The proof can be easily extended to show the statement for all $x, y \in [0,1]$. Note further that round-robin is not the only picking sequence achieving EF1 for equal entitlements. Any picking sequence that is recursively balanced (RB), i.e., $|t_i - t_j| \leq 1$ in any prefix of π, results in an EF1 allocation [2].

2.3 Random Allocations

A random allocation is a probability distribution $\mathcal{L}$ over deterministic allocations. We mostly focus on additive valuations, so we conveniently use a representation as matrix X of marginal assignment probabilities for each good to each agent (i.e., a complete fractional allocation as defined above). We denote by $X^{\mathcal{L}}$ the fractional allocation corresponding to a lottery $\mathcal{L}$. Notice that different lotteries might produce the very same fractional allocation.

Throughout the paper, we denote by X (resp. Y) fractional (resp. integral) allocations in matrix form. Further, X_i (resp. Y_i) denotes a fractional (resp. integral) bundle of i in X (resp. Y). We will write $v_i(X_i)$ to denote the expected utility of an agent. Clearly, in case of additive valuations we have $v_i(X_i) = \sum_{g \in \mathcal{G}} v_i(g) \cdot x_{ig}$.

2.4 Decomposing Fractional Matrices

A decomposition of a fractional allocation X is a convex combination of integral (deterministic) allocations, i.e., $X = \lambda_1 Y^1 + \cdots + \lambda_k Y^k$, where $\sum_{h=1}^k \lambda_h = 1$, $y_{ig}^h \in \{0,1\}$, and $\sum_{i \in \mathcal{N}} y_{ig}^h = 1$, for each $i \in \mathcal{N}$, $g \in \mathcal{G}$ and $h \in [k] = \{1, \ldots, k\}$.

A *constraint structure* $\mathcal{H}$ consists of a collection of subsets $S \subseteq \mathcal{N} \times \mathcal{G}$. Every $S \in \mathcal{H}$ comes with a lower and upper quota denoted by $\underline{q}_S$ and $\overline{q}_S$, respectively. Quotas are integer numbers stored in $\mathbf{q} = \{(\underline{q}_S, \overline{q}_S) | S \in \mathcal{H}\}$.

An $n \times m$ matrix Y is feasible under $\mathbf{q}$ if for each $S \in \mathcal{H}$

$$\underline{q}_S \leq \sum_{(i,g) \in S} y_{ig} \leq \overline{q}_S \ .$$

A constraint structure $\mathcal{H}$ is a *hierarchy* if, for every $S, S' \in \mathcal{H}$, either $S \cap S' = \emptyset$ or one is contained in the other. $\mathcal{H}$ is a *bihierarchy* if it can be partitioned into $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$, such that $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$ and both $\mathcal{H}_1$ and $\mathcal{H}_2$ are hierarchies.

Budish et al. [11] generalize the well-known decomposition theorem by Birkhoff and von Neumann:

Theorem 1. *Given any fractional allocation X, a bihierarchy $\mathcal{H}$ and corresponding quotas $\mathbf{q}$, if X is feasible under $\mathbf{q}$, then, there exists a polynomial decomposition into integral matrices. Every matrix in the decomposition is feasible under $\mathbf{q}$. Further, the decomposition can be obtained in strongly polynomial time.*

In the rest of this paper, given a fractional allocation X and a bihierarchy $\mathcal{H}$, we define the quotas in $\mathbf{q}$ as follows: for every $S \in \mathcal{H}$ we set $\underline{q}_S = \lfloor x_S \rfloor$ and $\overline{q}_S = \lceil x_S \rceil$, where $x_S = \sum_{(i,g) \in S} x_{ig}$. The decomposition obtained with these quotas and bihierarchy $\mathcal{H}$ will be called the $\mathcal{H}$-decomposition.

Utility Guarantee Bihierarchy. We next define an extremely useful bihierarchy. For a deeper understanding, we refer the reader to [11, 19].

We set $\mathcal{H}_1 = \{C_g \mid g \in \mathcal{G}\}$, where $C_g = \{(i,g) \mid i \in \mathcal{N}\}$ represents the column corresponding to good $g \in \mathcal{G}$.

Roughly speaking, the hierarchy $\mathcal{H}_1$ ensures that, in any allocation of the decomposition, every good is integrally assigned (and therefore the allocation is complete).

For agent $i \in \mathcal{N}$, we consider the goods in non-increasing order of i's valuation, i.e., $v_i(g_1) \geq \ldots \geq v_i(g_m)$. Recall that ties are broken according to a predefined ordering of $\mathcal{G}$. We set $\mathcal{S}_i = \{\{g_1\}, \{g_1, g_2\}, \ldots, \{g_1, \ldots, g_m\}\}$. In other words, for every $h \in [m]$, $\mathcal{S}_i$ contains a set of the h most preferred goods of i. We write (i, S) to denote $\{(i,g) \mid g \in S\}$, and set

$$\mathcal{H}_2 = \{(i,S) \mid i \in \mathcal{N}, S \in \mathcal{S}_i\} \cup \{(i,g) | i \in \mathcal{N}, g \in \mathcal{G}\} \quad . \qquad (1)$$

The second set of constraints implies that if $x_{ig} = 0$ (resp. $x_{ig} = 1$) then $y_{ig} = 0$ (resp. $y_{ig} = 1$), for any Y in the decomposition. Note that (for convenience later on) we slightly abuse notation for $\mathcal{H}_2$ as it is not a set of sets of (row, col)-pairs.

Finally, the *utility guarantee bihierarchy* is given by $\mathcal{H}^{\text{UG}} = \mathcal{H}_1 \cup \mathcal{H}_2$. Clearly, both $\mathcal{H}_1$ and $\mathcal{H}_2$ are hierarchies.

This bihierarchy was fundamental in [11] to prove a main result. We here state it in a slightly stronger version (see [19] for the proof).

Corollary 1 (Utility Guarantee ± one Good). *Suppose we are given a fractional allocation X, and additive valuation functions v_i. Then for any matrix Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X the following hold:*

1. *if $v_i(Y_i) < v_i(X_i)$, then $\exists\, g \notin Y_i$ with $x_{ig} > 0$ such that $v_i(Y_i) + v_i(g) > v_i(X_i)$;*
2. *if $v_i(Y_i) > v_i(X_i)$, then $\exists\, g \in Y_i$ with $x_{ig} < 1$ such that $v_i(Y_i) - v_i(g) < v_i(X_i)$.*

In other words, Corollary 1 ensures that, in any deterministic allocation in the $\mathcal{H}^{\text{UG}}$-decomposition, the valuation of any agent i differs from $v_i(X_i)$ by at most the value of one good. Moreover, such a good must have a positive probability of occurring in i's bundle.

3 ADDITIVE VALUATIONS WITH ENTITLEMENTS

In this section, we present a lottery for additive valuations simultaneously achieving ex-ante WSD-EF (and hence ex-ante WEF) and ex-post WEF$(1,1)$ + WPROP1. In contrast to equal entitlements, we show a weaker ex-post guarantee. However, we prove this is necessary as no stronger envy notion is compatible with ex-ante WEF.

We also generalize a result of Freeman et al. [19] to entitlements: Similarly to the unweighted setting, we design a lottery that is ex-ante WGF and ex-post WEF_1^1 + WPROP1.

3.1 Ex-ante WSD-EF and Ex-post WEF$(1,1)$ + WPROP1

The main contribution of this subsection is to prove the following:

Theorem 2. *For entitlements and additive valuations, we can compute in strongly polynomial time a lottery that is ex-ante* WSD-EF *and ex-post* WPROP1 + WEF$(1,1)$.

Let us start by introducing our main algorithm DifferentSpeedsEating (DSE), which is inspired by Eating for equal entitlements in [2]. Agents continuously eat their most preferred available good at speed equal to their entitlement. Every agent starts eating her most preferred good; as soon as a good has been completely eaten it is removed from the set of available goods. Each agent that was eating this good continues eating her most preferred remaining one. The procedure terminates when no good remains. See Algorithm 1 for a formal description. Observe that by precomputing the times at which goods are removed, we can implement the algorithm in strongly polynomial time.

Algorithm 1: DifferentSpeedsEating

Input: An instance $\mathcal{I} = (\mathcal{N}, \mathcal{G}, \{v_i\}_{i \in \mathcal{N}})$ and the entitlements $w_1, \ldots, w_n$
Output: A fractional allocation X

$X \leftarrow \mathbf{0}_{n \times m}$ `// current fractional allocation`
$\mathbf{z} \leftarrow \mathbf{1}_m$ `// remaining supply of each good`
while $\mathcal{G} \neq \emptyset$ **do**
  $\mathbf{s} \leftarrow \mathbf{0}_m$ `// eating speed on each item`
  for $i \in \mathcal{N}$ **do**
    $g^i \leftarrow \arg\max_{g \in \mathcal{G}} v_i(g)$ `// most favored item`
    $\mathbf{s}(g^i) \leftarrow \mathbf{s}(g^i) + w_i$ `// sum speeds on item`
  for $g \in \mathcal{G}$ **do**
    $\mathbf{t}(g) \leftarrow \frac{\mathbf{z}(g)}{\mathbf{s}(g)}$ `// compute finishing times`
  $t \leftarrow \min_{g \in \mathcal{G}} \mathbf{t}(g)$ `// time when first item finished`
  for $i \in \mathcal{N}$ **do**
    $x \leftarrow t \cdot w_i$ `// amount of item eaten by i`
    $x_{ig^i} \leftarrow x_{ig^i} + x$ `// eat fraction of g^i`
    $\mathbf{z}(g^i) \leftarrow \mathbf{z}(g^i) - x$ `// reduce supply of g^i`
  $\mathcal{G} \leftarrow \mathcal{G} \setminus \{g \in \mathcal{G} \mid \mathbf{t}(g) \leq \mathbf{t}(g') \text{ for all } g' \in \mathcal{G}\}$ `// remove finished items`
return X

We denote by X^{DSE} the output of DSE. The key properties are summarized in the following lemma.

Lemma 1 (Properties of DSE). *The following holds:*

1. $\sum_{g \in \mathcal{G}} x_{ig}^{DSE} = w_i \cdot m$ *for each* $i \in \mathcal{N}$;

(2) *the time needed for agent i to eat one unit of goods is $\frac{1}{w_i}$;*
(3) *overall, one unit of goods is consumed in one unit of time, and therefore, DSE runs for m time units.*

Let us observe the behavior of DSE on $\mathcal{I}^*$.

Example 2 (DSE at work). *Letting $i = 2, 3$, agents' priorities in $\mathcal{I}^*$ for the goods are the following:*

$$g_1 \succ_1 g_2 \succ_1 g_3 \succ_1 g_4, \qquad g_2 \succ_i g_3 \succ_i g_1 \succ_i g_4 .$$

Notice that, for agent 1, goods g_1 and g_2 are identical and ties are broken in favor of the good coming first in the ordering $g_1, \ldots, g_4$. Agents 2 and 3 have same priorities, thus, they will always be eating the same good at the same time.

During a run of DSE, whenever a good gets entirely eaten up, the behavior of agents who were eating this good changes.

At the beginning, $x_{ig} = 0$, for all $i \in \mathcal{N}$ and $g \in \mathcal{G}$. Agent 1 starts eating g_1 while agents 2 and 3 good g_2. Notice that agents 2 and 3 together have the same speed as agent 1. At $t = 2$, g_1 and g_2 get fully consumed and $x_{1g_1} = 1$, $x_{2g_2} = \frac{2}{3}$ and $x_{3g_2} = \frac{1}{3}$, respectively. Agent 1 will start eating g_3 as well as agents 2 and 3. All the agents together have speed equal to 1. Notice that agent 1 would prefer good g_2, however, it has been consumed entirely by agents 2 and 3. At $t = 3$, g_3 is now fully consumed. We have $x_{1g_3} = \frac{1}{2}$, $x_{2g_3} = \frac{1}{3}$ and $x_{3g_3} = \frac{1}{6}$, respectively. All the agents are now starting to eat g_4. At $t = 4$, all goods are fully consumed and DSE returns the fractional allocation

$$X^{DSE} = \begin{pmatrix} 1 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & \frac{2}{3} & \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{3} & \frac{1}{6} & \frac{1}{6} \end{pmatrix} .$$

■

Our first result is that the output of DSE is WSD-EF.

Proposition 2. *X^{DSE} is WSD-EF.*

Proof. For convenience, we use $X = X^{\text{DSE}}$. Let us consider an agent $i \in \mathcal{N}$. Note that the goods $g_1, \ldots, g_m$ are ordered in the same manner as in DSE for agent i, since we always break ties according to a predefined ordering of $\mathcal{G}$. Now consider another agent $j \in \mathcal{N}$. Using the notation $G_k = \{g_1, \ldots, g_k\}$ for the first k goods in i's ordering, we show

$$w_j \cdot \sum_{g \in G_k} x_{ig} \geq w_i \cdot \sum_{g \in G_k} x_{jg} , \tag{2}$$

for every $k \in [m]$, and WSD-EF follows for agent i.

Let t_k be the time when i stops eating g_k during the run of DSE. We set $t_k = t_{k-1}$ if good g_k has been completely consumed before time t_{k-1} by others. This means that, by the time t_k, no good in G_k remains available. On the one hand, until time t_k, agent i could only consume goods in G_k, implying $w_i \cdot t_k = \sum_{g \in G_k} x_{ig}$. On the other hand, every good in G_k has been fully consumed by that time, i.e., $w_j \cdot t_k \geq \sum_{g \in G_k} x_{jg}$, for every $j \in \mathcal{N}$. Combining these two properties proves Equation (2) and, hence, the theorem. □

It is known that SD-EF implies ex-ante EF for additive valuations; it remains true for different entitlements.

Proposition 3. *Given a fractional allocation X, if X is ex-ante* WSD-EF*, then X is ex-ante* WEF.

Proposition 2 and Proposition 3 show that the outcome of DSE satisfies the ex-ante properties stated in Theorem 2 . We next show that X^{DSE} can be decomposed into a lottery with good ex-post properties. To this end, we use Theorem 1 with the bihierarchy $\mathcal{H}^{\text{UG}}$.

Example 3 (The $\mathcal{H}^{\text{UG}}$-decomposition). *The $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is a convex combination $\lambda_1 Y^1 + \cdots + \lambda_k Y^k$, for some integer k. Every allocation Y^h is deterministic and its properties are determined by the bihierarchy $\mathcal{H}^{\text{UG}}$. In the following, we use Y to refer to a generic deterministic allocation in the decomposition.*

Recall that $\mathcal{H}^{\text{UG}} = \mathcal{H}_1 \cup \mathcal{H}_2$. The hierarchy $\mathcal{H}_1$ deals only with columns and ensures that any Y in the decomposition is complete.

Let us now consider $\mathcal{H}_2$ defined in 1.

Note that only one agent appears in any pair of $\mathcal{H}_2$. Hence, we discuss the implications of Theorem 1 agent by agent.

Let us consider agent 1. The pair $(1, S)$ belongs to $\mathcal{H}_2$ if and only if $S \in \{\{g_1\}, \{g_1, g_2\}, \{g_1, g_2, g_3\}, \{g_1, g_2, g_3, g_4\}\} \cup \{\{g_2\}, \{g_3\}, \{g_4\}\}$.

The feasibility conditions imply:

$$y_{1g_1} = 1, \qquad y_{1g_1} + y_{1g_2} = 1,$$
$$1 \leq y_{1g_1} + y_{1g_2} + y_{1g_3} \leq 2, \qquad y_{1g_1} + y_{1g_2} + y_{1g_3} + y_{1g_4} = 2 ,$$

and

$$y_{1g_2} = 0, \qquad 0 \leq y_{1g_3} \leq 1, \qquad 0 \leq y_{1g_4} \leq 1 .$$

In other words, in any deterministic allocation Y, agent 1 always receives 2 goods. In particular, she always gets g_1 but never g_2. Moreover, she gets either g_3 or g_4, but not both of them.

Similarly, by imposing the corresponding feasibility conditions on agent 2 and 3 we can deduce: i) the bundle of agent 2 is of size either 1 or 2, never contains g_1 and must contain one good between g_2 and g_3, and possibly contains g_4; ii) the bundle of agent 3 contains at most one of g_2, g_3, g_4 and never contains g_1.

Finally, we provide a concrete $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} for $\mathcal{I}^*$ with the aforementioned properties:

$$X^{DSE} = \frac{1}{6} \cdot \begin{pmatrix} 1\,0\,0\,1 \\ 0\,1\,0\,0 \\ 0\,0\,1\,0 \end{pmatrix} + \frac{1}{6} \cdot \begin{pmatrix} 1\,0\,1\,0 \\ 0\,1\,0\,0 \\ 0\,0\,0\,1 \end{pmatrix} + \frac{1}{3} \cdot \begin{pmatrix} 1\,0\,0\,1 \\ 0\,0\,1\,0 \\ 0\,1\,0\,0 \end{pmatrix} + \frac{1}{3} \cdot \begin{pmatrix} 1\,0\,1\,0 \\ 0\,1\,0\,1 \\ 0\,0\,0\,0 \end{pmatrix} .$$

■

Theorem 3. *Every deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is* WEF(1, 1).

To show the theorem we need some preliminary notions.

Eating Time. We define the *eating time* $t(g)$ of a good g as the point in time when it has been entirely consumed (during a run of DSE). Whenever an agent starts eating a good g, she can start eating another good only after the eating time of g.

Goods Eaten by i at Time t. Recall that DSE runs for m units of time. Every agent i exactly eats a total mass of $w_i \cdot m$ of $\mathcal{G}$ during DSE. Let $g_1, \ldots, g_m$ be the ordering of goods according to v_i. We define $\mathsf{Eaten}(i, t) = \{g_1, \ldots, g_\ell\} = G_\ell$, where g_ℓ is either a good that agent i just finished to consume (i.e., t is the eating time of g_ℓ and agent i was consuming it) or agent i at time t is eating the good $g_{\ell+1}$, which has not been finished yet. Consequently, by time t, agent i may have contributed only to the consumption of goods in G_ℓ. In particular, all goods in G_ℓ have been entirely consumed (by i and/or others), since otherwise i would not start eating $g_{\ell+1}$.

Recall that w_i is the speed of i. At time $t = \frac{k}{w_i}$ agent i ate a total mass k of goods. With the next lemma, we show that the $\mathcal{H}^{\text{UG}}$-decomposition guarantees agent i deterministically receives at most k goods from the ones eaten by time $\frac{k}{w_i}$.

Lemma 2. *Given any deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE}, for every $i \in \mathcal{N}$ and $k = 1, \ldots, \lfloor w_i \cdot m \rfloor$,*

$$|Y_i \cap \mathsf{Eaten}(i, \tfrac{k}{w_i})| \le k \ .$$

Furthermore, $\lfloor w_i \cdot m \rfloor \le |Y_i| \le \lceil w_i \cdot m \rceil$.

Proof. By definition, $\mathsf{Eaten}(i, \frac{k}{w_i}) = G_\ell$, the ℓ most preferred goods of i, for some ℓ. Thus, by the time $\frac{k}{w_i}$, agent i only ate goods in G_ℓ and possibly is currently eating the next less preferred good. Moreover, goods are eaten by i in the same ordering we used to build the collection $\mathcal{S}_i$ in the definition of $\mathcal{H}^{\text{UG}}$ implying $(i, G_\ell) \in \mathcal{H}^{\text{UG}}$. Since $|Y_i \cap \mathsf{Eaten}(i, \frac{k}{w_i})| = \sum_{g \in G_\ell} y_{ig}$, the $\mathcal{H}^{\text{UG}}$-decomposition properties imply $\sum_{g \in G_\ell} y_{ig} \le \lceil \sum_{g \in G_\ell} x_{ig} \rceil$. This last is upper-bounded by k because of these two simple observations: g_ℓ is fully consumed by the time $\frac{k}{w_i}$, and at that time agent i ate k units of goods. The first claim follows.

The second claim immediately follows by the $\mathcal{H}^{\text{UG}}$-decomposition properties, since $(i, \mathcal{G}) \in \mathcal{H}^{\text{UG}}$. □

Given any deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition, consider agent i and sort the goods in Y_i in a non-increasing manner with respect to v_i: $Y_i = \{g^i_1, \ldots, g^i_{h_i}\}$ and $v_i(g^i_1) \ge \cdots \ge v_i(g^i_{h_i})$. By Lemma 2, we see $h_i = \lfloor w_i \cdot m \rfloor$ or $h_i = \lceil w_i \cdot m \rceil$.

Stopping Time. Given any deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X, for each $i \in \mathcal{N}$ and $k \in [h_i]$, we define the *stopping time* by $s(g^i_k) = \min\{t(g^i_k), \frac{k}{w_i}\}$. Note that $s(g^i_k)$, differently from the stopping time $t(g^i_k)$ which solely depends on the run of DSE, also depends on Y. Indeed, in Y_i good g^i_k is the k-th most preferred good. However, if the eating time is greater than $\frac{k}{w_i}$, this good might appear as $(k+1)$-th most preferred good in another deterministic allocation of the decomposition. For convenience, we omit Y in the notation since we only discuss stopping times of single allocations. Let us show a couple of useful properties of stopping times.

Lemma 3. *Given any deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE}, let g^i_k be the k-th most preferred good in Y_i, it holds $s(g^i_k) \in \left(\frac{k-1}{w_i}, \frac{k}{w_i}\right]$.*

Proof. By definition, $s(g^i_k) = \min\{t(g^i_k), \frac{k}{w_i}\} \le \frac{k}{w_i}$. For contradiction, suppose $t(g^i_k) \le \frac{k-1}{w_i}$. Then, $g^i_k \in Y_i \cap \mathsf{Eaten}\left(i, \frac{k-1}{w_i}\right)$. Notice that $t(g^i_1) \le \cdots \le t(g^i_k)$, by definition of DSE, and therefore $g^i_h \in Y_i \cap \mathsf{Eaten}\left(i, \frac{k-1}{w_i}\right)$, for each $h = 1, \ldots, k$.

In conclusion, $\left|Y_i \cap \mathsf{Eaten}\left(i, \frac{k-1}{w_i}\right)\right| \ge k$ which is a contradiction by Lemma 2, and hence $t(g^i_k) > \frac{k-1}{w_i}$. □

Notice that for the eating time $t(g^i_k)$ the same lower bound holds; however, we can only upper bound $t(g^i_k)$ by $\frac{k+1}{w_i}$. This difference will be crucial in the proof of Theorem 3 and motivates the definition of stopping times.

Lemma 4. *Given any deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE}, let g^i_k be the k-th most preferred good in Y_i. For every good g coming earlier in i's ordering of goods, it holds that $s(g) < s(g^i_k)$.*

Proof. The claim follows by the definition of stopping time and the properties of DSE. Indeed, by the definition of stopping time $s(g) \le t(g)$, and $t(g) < \min\{t(g^i_k), \frac{k}{w_i}\} = s(g^i_k)$. The second inequality holds because at time $s(g^i_k)$ agent i is eating or finishes to eat g^i_k, and g must have been eaten before i starts eating g^i_k. Further, the inequality is strict since agent i ate a positive fraction of g^i_k (that is, $x_{ig^i_k} > 0$); otherwise, since $(i, g^i_k) \in \mathcal{H}_2$, $x_{ig^i_k} = 0$ would imply $y_{ig^i_k} = 0$ and, hence, $g^i_k \notin Y_i$. □

We are now ready to show Theorem 3.

Proof of Theorem 3. Let Y be any deterministic allocation in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE}. The proof proceeds as follows: We first generate a picking sequence π, then show that Y is the output of such a picking sequence, and finally prove that π satisfies Proposition 1, for $x = y = 1$. This shows that Y is $\mathsf{WEF}(1,1)$.

Defining π. We sort the goods $\mathcal{G}$ in a non-decreasing order of stopping times $s_1, \ldots, s_m$ (defined according to Y). If $g \in Y_i$ is the h-th good in this ordering, then, $\pi(h) = i$.

Y is the result of π. Assume i is the h-th agent in π. Assume that $\pi(h) = i$ is the k-th occurrence of i in π. We show that for each $h \in [m]$, the most preferred available good for i is exactly g^i_k. Let us proceed by induction on h.

For $h = 1$, clearly, $k = 1$. By Lemma 4, g^i_1 must be the most preferred good of i, otherwise we contradict the fact that $s_1 = s(g^i_1)$ is the minimum stopping time. At this point no good has been assigned, so i selects g^i_1.

Assume the statement is true until the h-th component of π. We show it is true for $h+1 \le m$. Suppose a good g coming before g^i_k, in i's ordering, is still available. By Lemma 4, there exists h' s.t. $s_{h'} = s(g) < s(g^i_k)$ with $h' \le h$. By the inductive hypothesis, g must have been assigned to $\pi(h')$. On the other hand, g^i_k is still available, otherwise there exists $h' \le h$, such that $\pi(h')$ picked g^i_k during the h'-th round – a contradiction with the inductive hypothesis.

π satisfies Proposition 1. We now show that π satisfies $\mathsf{WEF}(1,1)$. Consider any prefix of π and any pair of agents i, j. Let us denote by t_i (resp. t_j) the number of picks of agent i (resp. j) in the considered prefix. Let s_j and s_i be the stopping time of the good selected by j at her t_j-th pick and the stopping time of the good selected by i at her (t_i+1)-th pick, respectively. If i has no (t_i+1)-th pick, we set $s_i = m \le \frac{t_i+1}{w_i}$. Within the considered prefix of π, agent j already made its t_j-th pick but i didn't make its (t_i+1)-th pick. Now by definition of π, $s_j \le s_i$. By Lemma 3, $s_j > \frac{t_j-1}{w_j}$ and $s_i \le \frac{t_i+1}{w_i}$. We finally get $\frac{t_j-1}{w_j} < \frac{t_i+1}{w_i}$. This shows that the hypothesis of Proposition 1 is fulfilled for $x = y = 1$. □

Note that if we had chosen eating rather than stopping times for the picking sequence, we could only deduce $\frac{t_j-1}{w_j} < \frac{t_i+2}{w_i}$ which is not sufficient to show $\mathsf{WEF}(1,1)$.

As X^{DSE} is (ex-ante) WEF, it is also WPROP. By ex-ante WPROP and Corollary 1, the following is implied.

Proposition 4. *Every deterministic allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is* WPROP1.

Proof. The fractional allocation X^{DSE} is WEF, and hence WPROP. Therefore, $v_i(X_i) \geq w_i \cdot v_i(\mathcal{G})$. By Corollary 1, for any Y in the $\mathcal{H}^{\text{UG}}$-decomposition, $v_i(Y_i) + v_i(g) > v_i(X_i)$, for some $g \in \mathcal{G} \setminus Y_i$. This implies $v_i(Y_i \cup \{g\}) \geq w_i \cdot v_i(\mathcal{G})$, and WPROP1 follows. □

In conclusion, we proved that the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is a lottery achieving ex-ante WSD-EF, and therefore ex-ante WEF, and ex-post $\mathsf{WEF}(1,1)$ + WPROP1. As a consequence of Theorem 1, our lottery has polynomial support and the computation requires strongly polynomial time.

While our guarantee is weaker than the ex-post EF1 for equal entitlements, we show that our lottery is, in a sense, best possible in terms of ex-post guarantees. Indeed, we prove that no stronger ex-post envy notion is compatible with ex-ante WEF.

Proposition 5. *For every pair $x, y \in [0,1]$ such that $x + y < 2$, ex-ante* WEF *is incompatible with ex-post* $\mathsf{WEF}(x,y)$.

Proof. Consider a fair division instance $\mathcal{I} = (\mathcal{N}, \mathcal{G}, \{v_i\}_{i \in \mathcal{N}})$, where $\mathcal{N} = \{1,2\}$ and $\mathcal{G} = \{g_1, g_2\}$. Moreover, $v_i(g_1) = v_i(g_2) = 1$, for $i = 1,2$. Let us set $w_1 \in \left(\frac{y}{2+y-x}, \frac{1}{2}\right)$ and $w_2 = 1 - w_1$. Observe that $\frac{y}{2+y-x} < \frac{1}{2}$, since $x + y < 2$. In any ex-ante WEF allocation agent 1 receives in expectation less than one good. This means, the allocation $Y = (Y_1, Y_2) = (\emptyset, \mathcal{G})$ is in the support of any ex-ante WEF lottery. Therefore, since $w_1 + w_2 = 1$, for each $g \in Y_2$

$$\begin{aligned} w_1 \cdot (v_1(Y_2) - x \cdot v_1(g)) &= w_1 \cdot (2 - x) \\ &> \frac{y}{2 + y - x} \cdot (2 - x) > w_2 \cdot y \\ &= w_2 \cdot (v_1(Y_1) + y \cdot v_1(g)) \ . \end{aligned}$$

This proves Y is not $\mathsf{WEF}(x,y)$. □

Remark. For equal entitlements our approach also provides ex-ante EF and ex-post EF1. The ex-ante property follows directly since $w_i = 1/n$. For ex-post EF1, similarly to [2], it is possible to show that any allocation Y in the $\mathcal{H}^{\text{UG}}$-decomposition of the X^{DSE} is the result of an RB picking sequence. In particular, this holds for the picking sequence defined in the proof of Theorem 3.

3.2 Ex-ante WGF and Ex-post WEF_1^1 + WPROP1

In this subsection, we generalize Theorem 4/Corollary 1 of Freeman et al. [19] to entitlements. We follow the general argument and incorporate some technical extensions to allow for different agent weights.

Theorem 4. *For entitlements and additive valuations, we can compute in strongly polynomial time a lottery that is ex-ante* WGF *and ex-post* WEF_1^1 + WPROP1.

Note that ex-ante WGF implies ex-ante WEF. Moreover, ex-ante WGF implies ex-ante Pareto optimality.

Remark. One might wonder whether the ex-post guarantee in Theorem 4 could be replaced with $\mathsf{WEF}(x,y)$ for some parameters $x, y \in [0,1]$. There are instances where this is impossible, even in the unweighted setting. Consider the following example: There are three agents 1, 2, 3, three light goods, and one heavy good. Agents 1 and 2 have the same valuation function, they value the heavy good at 6, and each light good at 1. Agent 3 values the light goods at 1 and the heavy good at 0.

Now consider a fractional group fair allocation X. Observe that all light goods need to be allocated completely to agent 3 in X. If one of the first two agents (say, agent 1) gets a fraction $\epsilon > 0$ of light goods, the group fairness condition is violated for $S = \{2,3\}$ and $T = \{1,3\}$: If one reallocates the fraction ϵ of light goods from agent 1 to agent 3, then the utility agent 3 strictly improves and the utility of agent 2 remains unchanged.

Now consider any allocation Y in the support of a lottery implementing X. Y needs to give all light goods to agent 3, so at least one of the first two agents gets no good at all. This agent is then envious to agent 3, and transferring one good from agent 3 to agent 1 cannot remove this envy.

4 EXTENSIONS TO GENERAL VALUATIONS

In this section, we explore to which extent our techniques apply to more general valuations. Let us start by introducing the classes that will be discussed.

Classes of Valuations. A valuation function v_i is *k-unit-demand* if $v_i(A)$ is given by the sum of the k most valuable goods in A for i. A valuation function is *multi-demand*, iff it is k-unit-demand for some $k \in \mathbb{N}$. If $k = 1$ we talk about *unit-demand.*

A valuation function v_i is *cancelable* if

$$v_i(S \cup \{g\}) > v_i(T \cup \{g\}) \Rightarrow v_i(S) > v_i(T) \ ,$$

for all $S, T \subseteq \mathcal{G}$ and $\forall g \in \mathcal{G} \setminus (S \cup T)$. Cancelable valuations generalize several classes studied in the literature, e.g., additive, weakly-additive, budget-additive, product, and unit-demand, see [8].

A valuation function v_i is *XOS* if there is a family of additive set functions $\mathcal{F}_i$ such that $v_i(A) = \max_{f \in \mathcal{F}_i} f(A)$. XOS generalize additive and submodular valuations.

Both ex-ante WEF and ex-ante WSD-EF allocations exist for all valuations. In particular, for ex-ante WEF we use the trivial uniform random assignment from the introduction. For ex-ante WSD-EF it is sufficient to invoke DSE only using agents' priorities over single goods. Observe that for general valuations it is no longer true that ex-ante WSD-EF implies ex-ante WEF. We next show that ex-ante WEF and either ex-post WPROP1 or ex-post $\mathsf{WEF}(1,1)$ is no longer guaranteed.

Theorem 5. *For general valuations, ex-ante* WEF *is not compatible with* WPROP1 *or* $\mathsf{WEF}(1,1)$.

Proof. Let us consider a fair division instance with two agents and four goods. Suppose agent 1 has an entitlement of $\frac{2}{3}$, and has value 1 for any bundle (except the empty bundle, for which she has value 0). Agent 2 has entitlement $\frac{1}{3}$ and value k for any bundle of size k, for $k = 0, \ldots, 4$.

Let us denote by p_k the probability that 1 receives k goods. Clearly, if the allocation is ex-post WPROP1 or ex-post $\mathsf{WEF}(1,1)$, then $p_0 = 0$ and $\sum_{k=1}^{4} p_k = 1$. Since the allocation is complete, p_k is also the probability that agent 2 receives $4 - k$ goods. Agent 1 is ex-ante WEF if and only if

$$\frac{1}{3} \cdot (p_1 + p_2 + p_3 + p_4) \geq \frac{2}{3} \cdot (p_0 + p_1 + p_2 + p_3)$$

which implies $1 - p_0 \geq 2 \cdot (1 - p_4)$ and, hence, $2p_4 \geq 1 + p_0$. Thus, $p_4 \geq \frac{1}{2}$. Therefore, the allocation where agent 1 receives every good and 2 no good occurs with positive probability. However, such an allocation is neither WPROP1 nor WEF$(1, 1)$ (not even WEF$_1^1$) for agent 2. □

Notice that both agents in the above example value 1 each good, and hence agent 1 is unit-demand and agent 2 is additive.

4.1 XOS Valuations

For an agent i with XOS valuation, our algorithm only makes use of the additive function f_i such that $v_i(\mathcal{G}) = \sum_{g \in \mathcal{G}} f_i(g)$, i.e., the additive function for the grand bundle. Therefore, either we assume f_i to be known or we have access to an XOS-oracle (using which f_i can be obtained with a single query). Given a query with a set $A \subseteq \mathcal{G}$, the XOS-oracle returns a function $f \in \mathcal{F}_i$ that maximizes $f(A)$.

Let X be the fractional allocation with $x_{ig} = w_i$, for each $i \in \mathcal{N}$ and $g \in \mathcal{G}$.

Proposition 6. *X is ex-ante* WPROP.

Proof. Let $\lambda_1 Y^1 + \cdots + \lambda_k Y^k$ be any decomposition of X. For any allocation $Y \in \{Y^\ell\}_{\ell \in [k]}$, $v_i(Y_i) = \max_{f \in \mathcal{F}_i} f(Y_i) \geq f_i(Y_i)$, since $f_i \in \mathcal{F}_i$. Hence, the expected utility of agent i in the lottery is

$$\begin{aligned}\sum_{h=1}^{k} \lambda_h v_i(Y_i^h) &\geq \sum_{h=1}^{k} \lambda_h f_i(Y_i^h) = \sum_{h=1}^{k} \sum_{g \in Y_i^h} \lambda_h f_i(g) \\ &= \sum_{g \in \mathcal{G}} \sum_{h : g \in Y_i^h} \lambda_h f_i(g) = \sum_{g \in \mathcal{G}} f_i(g) \sum_{h : g \in Y_i^h} \lambda_h \\ &= \sum_{g \in \mathcal{G}} x_{ig} f_i(g) = w_i \sum_{g \in \mathcal{G}} f_i(g) = w_i \cdot v_i(\mathcal{G}) \ .\end{aligned}$$

□

In order to apply Theorem 1, we need to set up an appropriate additive function. For the next result, we assume that agent i has additive valuation f_i, for each $i \in \mathcal{N}$.

Proposition 7. *The $\mathcal{H}^{\text{UG}}$-decomposition of X is ex-post* WPROP1.

Proof. Given any allocation Y of the decomposition, by definition of XOS, Corollary 1 and Proposition 6, we see that

$$v_i(Y_i \cup \{g\}) \geq f_i(Y_i \cup \{g\}) = f_i(Y_i) + f_i(g) > f_i(X_i) = w_i \cdot v_i(\mathcal{G}) \ .$$

□

4.2 Equal Entitlements

Here we briefly discuss to which extent we can guarantee BoBW results for equally entitled agents and general valuations. In particular, we explore valuation functions for which Eating together with the $\mathcal{H}^{\text{UG}}$-decomposition can be used to guarantee ex-ante EF and ex-post EF1.

Both Eating and the definition of the $\mathcal{H}^{\text{UG}}$ bihierarchy only depend on the ranking of each agent for singleton bundles of goods. Therefore, we can determine a fractional allocation X^{DSE} with Eating and compute its $\mathcal{H}^{\text{UG}}$-decomposition for any class of valuation functions. This also yields the desired properties for broader classes of valuations.

Theorem 6. *For equal entitlements, if an agent $i \in \mathcal{N}$ has a k-unit-demand valuation for some $k \in \mathbb{N}$, then the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is ex-ante* EF *and ex-post* EF1 *for i.*

We point out that we must rely on the $\mathcal{H}^{\text{UG}}$-decomposition to prove ex-ante EF. In fact, when valuations are multi-demand, not all lotteries implementing X^{DSE} are ex-ante EF.

Being additive valuations a special case of multi-demand valuations (it is sufficient to set $k = m$), we have the following:

Corollary 2. *For equal entitlements and any combination of additive and multi-demand valuations, the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is ex-ante* EF *and ex-post* EF1.

Another interesting observation is that the concept of SD-EF depends only on the ranking of single goods provided by the agents. Thus, the output of Eating is always an ex-ante SD-EF allocation, regardless of the valuation function. By proving that any RB picking sequence gives an EF1 allocation for cancelable valuations (see full version), we obtain the following:

Theorem 7. *For equal entitlements and cancelable valuations, the $\mathcal{H}^{\text{UG}}$-decomposition of X^{DSE} is ex-ante* SD-EF *and ex-post* EF1.

Unfortunately, we were not able to prove that the lottery is ex-ante EF (only ex-ante SD-EF) for cancelable valuations, and this remains an interesting open question.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we obtain best of both worlds results for fair division with entitlements. Our results for additive valuations paint a rather complete picture. We present a lottery that can be computed in strongly polynomial time and guarantees ex-ante WEF and ex-post WEF$(1, 1)$ + WPROP1. This is tight in the sense that any stronger notion of WEF(x, y) is incompatible with ex-ante WEF. We also present a lottery that is ex-ante WGF (and therefore ex-ante Pareto optimal) and ex-post WEF$_1^1$ + WPROP1. Again, ex-ante WGF is incompatible with stronger ex-post notions.

We also explore how some of our results can be extended to more general valuation functions. These insights represent an interesting first step, but many important open problems remain. As a prominent one, to the best of our knowledge, it is open for which classes of valuation functions ex-ante EF is always compatible with ex-post EF1 in the unweighted setting. In addition, providing tight guarantees with entitlements and combinations of other fairness concepts (such as, e.g., variants of the Max-Min-Share) is an interesting direction for future work.

Finally, in our work, we have not put particular attention to Pareto optimality. This is motivated by the impossibility result of [2, 19]: i) ex-ante Prop, ex-post EF1, and ex-post fractional Pareto optimality are incompatible and ii) ex-ante SD-EF, ex-post EF1, and ex-post Pareto optimality are incompatible even for 2 agents. Reducing the ex-post EF1 guarantee in favor of Pareto optimality is indeed another interesting research direction.

ACKNOWLEDGMENTS

Martin Hoefer and Giovanna Varricchio were supported by DFG grant Ho 3831/5-1.

REFERENCES

[1] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A. Voudouris. 2022. Fair Division of Indivisible Goods: A Survey. In *Proc. 31st Int. Joint Conf. Artif. Intell. (IJCAI)*. 5385–5393.

[2] Haris Aziz. 2020. Simultaneously achieving ex-ante and ex-post fairness. In *Proc. 16th Conf. Web and Internet Econ. (WINE)*. 341–355.

[3] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. 2014. Fair assignment of indivisible objects under ordinal preferences. In *Proc. 13th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1305–1312.

[4] Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy. 2020. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Oper. Res. Lett.* 48, 5 (2020), 573–578.

[5] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2021. Fair-Share Allocations for Agents with Arbitrary Entitlements. In *Proc. 22nd Conf. Econ. Comput. (EC)*. 127–127.

[6] Moshe Babaioff, Tomer Ezra, and Uriel Feige. 2022. On Best-of-Both-Worlds Fair-Share Allocations. In *Proc. 18th Conf. Web and Internet Econ. (WINE)*. 237–255.

[7] Siddharth Barman and Sanath Kumar Krishnamurthy. 2019. On the Proximity of Markets with Integral Equilibria. In *Proc. 33rd Conf. Artif. Intell. (AAAI)*. 1748–1755.

[8] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. 2022. Almost full EFX exists for four agents. In *Proc. 36th Conf. Artif. Intell. (AAAI)*. 4826–4833.

[9] Anna Bogomolnaia and Herve Moulin. 2001. A new solution to the random assignment problem. *J. Econ. Theory* 100 (2001), 295–328.

[10] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *J. Political Econ.* 119, 6 (2011), 1061–1103.

[11] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. Designing random allocation mechanisms: Theory and applications. *Amer. Econ. Rev.* 103, 2 (2013), 585–623.

[12] Ioannis Caragiannis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2021. On Interim Envy-Free Allocation Lotteries. In *Proc. 22nd Conf. Econ. Comput. (EC)*. 264–284.

[13] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Trans. Econ. Comput.* 9, 3 (2021), 1–39.

[14] Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. 2021. Picking sequences and monotonicity in weighted fair division. *Artif. Intell.* 301 (2021), 103578.

[15] Mithun Chakraborty, Erel Segal-Halevi, and Warut Suksompong. 2022. Weighted Fairness Notions for Indivisible Items Revisited. In *Proc. 36th Conf. Artif. Intell. (AAAI)*. 4949–4956.

[16] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. 2017. Fair Public Decision Making. In *Proc. 18th Conf. Econ. Comput. (EC)*. 629–646.

[17] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. 2019. Group Fairness for the Allocation of Indivisible Goods. In *Proc. 33rd Conf. Artif. Intell. (AAAI)*. 1853–1860.

[18] Alireza Farhadi, Mohammad Taghi Hajiaghayi, Mohammad Ghodsi, Sébastien Lahaie, David M. Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2017. Fair Allocation of Indivisible Goods to Asymmetric Agents. In *Proc. 16th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*. 1535–1537.

[19] Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2020. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proc. 21st Conf. Econ. Comput. (EC)*. 21–22.

[20] Jugal Garg, Pooja Kulkarni, and Rucha Kulkarni. 2020. Approximating Nash Social Welfare under Submodular Valuations through (Un)Matchings. In *Proc. 31st Symp. Discret. Algorithms (SODA)*. 2673–2687.

[21] Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. 2022. Best of Both Worlds: Agents with Entitlements. https://arxiv.org/abs/2209.03908

[22] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proc. 5th Conf. Econ. Comput. (EC)*. 125–131.

[23] Warut Suksompong and Nicholas Teh. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences* 117 (2022), 101–108.