Best of Both Worlds Fairness under Entitlements

Haris Aziz
UNSW Sydney
Sydney, Australia
haris.aziz@unsw.edu.au

Aditya Ganguly
UNSW Sydney
Sydney, Australia
a.ganguly@unsw.edu.au

Evi Micha
University of Toronto
Toronto, Canada
emicha@cs.toronto.edu

ABSTRACT

We consider probabilistic allocation of indivisible items to agents with additive valuations and weighted entitlements. We explore how far ex-ante and ex-post fairness properties can be achieved simultaneously. Our first result is that in contrast to the case of same entitlements, well-established adaptations of ex-ante envy-freeness and ex-post envy-freeness up to one item (EF1) to the case of entitlements are not compatible. We then present a polynomial-time algorithm that achieves weighted ex-ante envy-freeness and ex-post weighted envy-freeness up to 1 transfer. The outcome is ex-ante weighted envy-free for all utilities consistent with the underlying ordinal preferences but it is not Pareto optimal. We then present an alternative polynomial-time algorithm that satisfies Pareto optimality (both ex-ante and ex-post), ex-ante weighted envy-freeness and ex-post weighted proportionality up to one item.

1 INTRODUCTION

The problem of fairly allocating resources among people and groups that lay claim to them stretches far back into history. The archetypal example of such a problem appears in numerous accounts from antiquity, where a plot of land has to be divided fairly between two people who may value different parts of the land differently to each other. The solution - 'divide and choose', in which one person cuts the land into two parts that she deems equally valuable and the other chooses their preferred piece - is one of the earliest classical algorithms.

This is of course just one instance of a problem within the broad field of fair division, which, in modern times, has had its origins in the works of Steinhaus , Knaster and Banach (see, e.g., [25]). As contextual motivations have grown over time, it has lead to the development of models and axioms capable of capturing the intricacies of the real-world problems they represent. Suksompong [24], for example, surveys a selection of ways in which general fair division problems can be altered in order to reveal new questions with unique mathematical flavours.

We study fair division settings in which all of the items being allocated are *indivisible* - so agents must either get all or none of an item. We call problems of this nature *allocation problems*. Some real-life examples of our premise include the distribution among heirs of indivisible items from an estate and the assignment of courses to lecturers who wish to teach them. As noted by Budish [11], research into this problem is primarily motivated by outcomes that are *fair* or *efficient*. Allocation problems with indivisible items present a number of inherent difficulties compared to their counterpart where items are infinitely divisible. This issue can be partly alleviated by introducing randomization and leveraging the properties of divisible items by issuing agents items with a certain probability. More recently though, through the works of the likes of Lipton et al. [19], Budish [11], Gourvès et al. [18] and Caragiannis et al. [13], there has been an effort to study fairness in the allocation of indivisible items by first introducing appropriate notions of *approximate* fairness that reflect the indivisible quality of the items. Fundamental questions in this field relate to the existence and computation of allocations satisfying these fairness axioms.

An interesting approach to studying fairness that has gained traction in recent years has been the idea of '*best of both worlds fairness*' [3, 5, 17]. At a high level, this refers to outcomes to allocation problems that have both a randomized (or fractional) component and a deterministic component. Best of both worlds fair outcomes first make use of randomization and offer agents probability shares in the items in a fair way. After this, the random outcomes must be 'implemented' by a lottery over deterministic outcomes that all satisfy certain other fairness axioms.

In all previous papers on best of both worlds fairness, it is assumed that agents have equal entitlements. However, in many scenarios agents may have asymmetric entitlements. It can be argued that agents having different entitlements is an even more common use-case for real-life applications of allocation problems than agents having equal entitlements (see, e.g., [16], and so generalising existing best of both worlds fairness research in this way provides a valuable new insight into how to apply such fairness principles to a host of problems in which agents cannot be assumed to have equal entitlements over items.

In this paper we examine the following central problem.

> *When agents have entitlements, to what extent can best of both worlds fairness can be satisfied?*

Contributions. Our first result is that in sharp contrast to the case of same entitlements, well-established adaptations of ex-ante envy-freeness and ex-post envy-freeness up to one item (EF1) to the case of entitlements are not compatible. The result underscores the additional challenge when handling asymmetric entitlements. We then propose a weak verion of the ex-post envy-freeness properties called weighted envy-freeness up to 1 transfer (WEF1-T). We then present a polynomial-time algorithm that achieves weighted ex-ante envy-freeness and ex-post WEF1T. The outcome is ex-ante weighted envy-free for all utilities consistent with the underlying ordinal preferences but it is not Pareto optimal. We then present an alternative polynomial-time algorithm that satisfies ex-ante weighted envy-freeness and ex-post weighted proportionality up to one item.

2 RELATED WORK

In this paper, we consider randomized allocation of items. A seminal work on random allocation is by Bogomolnaia and Moulin [9] who presented the probabilistic serial algorithm that satisfies envy-freeness with respect to all cardinal utilities consistent with the ordinal preferences. We consider the additional issue of weighted

	Weighted PS lottery Algo	Weighted Max Nash lottery Algo
ex-ante WEF for all consistent utilities	yes	no
ex-ante WEF	yes	no
ex-post WEF1-T	yes	?
ex-post WPROP1	?	yes
ex-ante PO	no	yes
ex-post PO	no	yes

Table 1: Summary of results

entitlements and also best of both world (BoBW) fairness that also pertains to ex-post fairness guarantees.

The interest in BoBW fairness was reignited by Freeman et al. [17], who first showed that ex-ante EF and ex-post EF1 BoBW outcomes can be found for any allocation problem instance. Aziz [2] designed an efficient algorithm known as PS-Lottery, which first calls the Probabilistic Serial rule of Bogomolnaia & Moulin [10] in order to generate an envy-free random allocation (from which they get ex-ante EF), and then is able to decompose this random allocation as the convex combination of EF1 deterministic allocations, yielding EF1 ex-post. In fact, since PS-Lottery does not take in agents' cardinal utilities, both the ex-ante and ex-post fairness properties can be strengthened to their stochastic dominant versions. Aziz [1] proposed research directions regarding probabilistic decision making with desirable ex-ante and ex-post properties.

Having established existence of BoBW fair allocations, Freeman et al. [17] show that BoBW fair allocations can be found in conjunction with a strong version of efficiency - that is ex-ante fractional Pareto optimality (fPO) - as long as the ex-post fairness condition is relaxed. Such outcomes are constructed by a novel algorithm known as MNW-Lottery, which can be viewed as a means by which to implement the random allocation generated by the Maximum Nash Welfare rule (see [13] for more details on the surprisingly fair properties of the MNW rule). Note that the Nash welfare is defined to be the product of all agents' (expected) utilities from a (random) allocation.

Although our work mainly intends to study BoBW outcomes from the perspective of envy-based fairness, we briefly mention a recent paper that engages with BoBW fairness from a fair-share guarantee perspective. Babaioff et al. [5] establish the existence of BoBW outcomes that are ex-ante proportional and ex-post guarantee every agent at least half of their maximin faire share (as well as proportionality up to at most one item) with a polynomial-time algorithm.

Chakraborty et al. [14] proposed two extensions of EF1 to the case of weighted entitlements. They proposed (1) weighted envy-freeness up to one item (WEF1) where envy can be eliminated by removing an item from the envied agent's bundle, and (2) weak weighted envy-freeness up to one item (WWEF1) where envy can be eliminated either by removing an item (as in the strong version) or by replicating an item from the envied agent's bundle in the envying agent's bundle. They presented various algorithmic results for both concepts.

We also briefly touch on some alternative notions of approximate fairness for agents with different entitlements that are mentioned elsewhere in the literature. Farhadi et al. [16] also study the model of weighted allocation problems, but approach fairness through the lens of fair-share guarantees rather than envy-based fairness (namely, a weighted version of the maximin fair share). In particular, they show that it is impossible to find an algorithm that can guarantee every agent more than $1/n$ of their weighted maximin share regardless of agents' entitlements. In contrast, the work of Proacaccia and Wang [22] that this study generalises shows that when agents have equal entitlements, allocations can always be found that guarantee all agents at least $2/3$ (a constant proportion) of their MMS. This stark difference underlines the potential difficulty in translating results from the setting of symmetric agents to a setting where agents may have unequal entitlements. In a closely related work, Babaioff et al. [6] raise some concerns about the way in which the weighted maximum fair share axiom of Farhadi et al. [16] aligns with the intuitive understanding behind this property. They instead propose a different notion of fair share known as the *AnyPrice share*, which happens to always be greater than the (unweighted) MMS. It is then shown that a constant fraction of the AnyPrice share can always be guaranteed to all agents, even when they have different entitlements.

Finally in this section, we mention another topic closely related to fairness for agents with different entitlements. Consider an adaptation of a motivating problem from Benabbou et al. [7] in which instead of allocating items to individual agents who all have equal entitlements, items (in this example, public housing in Singapore) are allocated to people belonging to ethnic groups, and each ethnic group should receive items fairly in relation to the size of the group. This highlights the relationship between the problems of fair allocation to agents with different entitlements and fair allocation to agents belonging to groups of various sizes.

3 PRELIMINARIES

Let $[t] = \{1, \ldots, t\}$, for each $t \in \mathbb{N}$. We consider a set $N = [n]$ of n agents and a set $O = \{o_1, \ldots, o_m\}$ of m items. Each agent i is characterized by a weight $w_i \geq 0$ which shows her *entitlement*. Without loss of generality, we assume that $\sum_{i \in N} w_i = 1$. Let $w = (w_1, \ldots, w_n)$. A subset of the items is called a *bundle*.

A *fractional allocation* A is a $(n \times m)$ matrix in which the entry in the i^{th} row and j^{th} column, which with a slight abuse of notation denoted by A_{i,o_j}, is the fraction of good o_j assigned to agent i; for each $o_j \in O$, $\sum_{i \in N} A_{i,o_j} = 1$. A fractional allocation is *integral*, if $A_{i,o_j} \in \{0, 1\}$ for all $i \in N$ and $o_j \in O$. Under an integral allocation A, we denote with A_i the bundle that is assigned to agent i and the allocation A can be characterized by the bundles of the agents, i.e. $A = (A_1, \ldots, A_n)$. When we refer to an allocation, we will mean a fractional allocation, unless otherwise is explicitly specified. For notational clarity, we will use X for fractional allocations and A for integral allocations.

A *randomized allocation* is a lottery over integral allocations. In particularly, a randomized allocation R is determined by k pairs $\{p^j, A^j\}$, where the integral allocation A^j is implemented with probability $p^j > 0$ and $\sum_{j\in[k]} p^j = 1$. We say that such an integral allocation is in the support of the randomized allocation. Moreover, we say that a fractional allocation X *implements* a randomized allocation R, if the marginal probability of agent i receiving item o_j is X_{i,o_j}.

Each agent is endowed with a *valuation* function $u_i : 2^O \to \mathbb{R}_{\geq 0}$. We assume that these valuations are *additive* which means that for each $T \subseteq O$, $u_i(T) = \sum_{g\in T} u_i(\{g\})$. For simplicity, we write $u_i(g)$ instead of $u_i(\{g\})$. We denote with $u = (u_1, \ldots, u_n)$ the utility profile. The utility of an agent i for an allocation X, which with a slight abuse of notation is denoted by $u_i(X)$, is given by $u_i(X) = \sum_{o_j\in O} X_{i,o_j} \cdot u_i(o_j)$. The utility function u_i induces a preference ranking $\succsim_i$ such that $o \succsim_i o'$ if and only if $u_i(o) \geq u_i(o')$. We call $\succsim = (\succsim_1, \ldots, \succsim_n)$ as the preference profile. An instance in our setting is given by $(N, O, w, \succsim)$.

3.1 Fairness Concepts

We make an important distinction between a fairness property holding ex-ante and ex-post. For any property $\langle P\rangle$ defined for an allocation, we say that a randomized allocation R satisfies $\langle P\rangle$ *ex-ante* if the allocation X that implements it satisfies $\langle P\rangle$. For any property $\langle Q\rangle$ defined for an integral allocation, we say that a randomized allocation R satisfies $\langle Q\rangle$ *ex-post* if every integral allocation in its support satisfies $\langle Q\rangle$.

Our main goal is to find allocations that are weighted envy-free.

Definition 3.1 (Weighted Envy-Freeness (WEF). An allocation X is *weighted envy-free (WEF)* if for all agents i, j,

$$\frac{u_i(X)}{w_i} \geq \frac{u_j(X)}{w_j}.$$

Since there are instances in which no integral allocation is weighted envy-free, ex-post weighted envy-freeness is not always achievable. Therefore, relaxations of this have been studied. While when all the agents have the same weight, envy-freenes up to one item is the most natural approximation of envy-freenes, notions of approximate envy-based fairness are slightly less obvious to adapt to the weighted problem. We primarily adopt two definitions that introduced by Chakraborty et al. [14].

Definition 3.2 (Weighted Envy-freeness Up To One Item (WEF1)). An integral allocation A is (strongly) weighted envy-free up to one item (WEF1) if for all agents i, j, either

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

or there exists an item $o \in A_j$ such that

$$\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j \setminus o)}{w_j}.$$

A weaker approximation of weighted envy-freeness is the following.

Definition 3.3 (Weakly Weighted Envy-freeness Up To One Item (WWEF1)). An integral allocation A is weakly weighted envy-free up to one item (WWEF1) if for all agents i, j, either

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

or there exists an item $o \in A_j$ such that either

$$\frac{u_i(A_i)}{w_i} \geq \frac{u_i(A_j \setminus o)}{w_j}$$

or

$$\frac{u_i(A_i \cup o)}{w_i} \geq \frac{u_i(A_j)}{w_j}.$$

We also propose a new axiom that further relaxes WWEF1.

Definition 3.4 (Weighted Envy-freeness Up To One Transfer (WEF1-T)). An integral allocation A is weighted envy-free up to one transfer (WEF1-T) if for all agents i, j, either

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

or there exists an item $o \in A_j$ such that

$$\frac{u_i(A_i \cup o)}{w_i} \geq \frac{u_i(A_j \setminus o)}{w_j}.$$

It is easy to see that WEF1-T is weaker than WWEF1 and in turn, WEF1.

We are also interested in allocations that are proportionally fair.

Definition 3.5 (Weighted Proportional Up To One Item (WPROP1)). An integral allocation A is *weighted proportional up to one item* (WPROP1) if for each agent i, either

$$u_i(A_i) \geq w_i \cdot u_i(O)$$

or there exists an item $o \in O \setminus A_i$ such that

$$u_i(A_i \cup \{o\}) \geq w_i \cdot u_i(O).$$

Lastly, we define the efficiency of an allocation through the following definition.

*Definition 3.6 (*Fractional Pareto Optimality (fPO) and Pareto Optimality (PO)*).* An allocation X is *fractional Pareto optimal* if there does not exists another allocation X' that Pareto dominates X, i.e. $u_i(X') \geq u_i(X)$ for all $i \in N$ and $u_i(X') > u_i(X)$ for some $i \in N$. An integral allocation A is *Pareto optimal*, if there does not exist another integral allocation A' that Pareto dominates A, i.e. $u_i(A') \geq u_i(A)$ for all $i \in N$ and $u_i(A') > u_i(A)$ for some $i \in N$.

4 INCOMPATIBILITY BETWEEN EX-ANTE WEF AND EX-POST WWEF1

In the best of both worlds literature, the central result is that when there are no entitlements, then ex-ante EF (or equivalent WEF) and ex-post EF1 (or equivalent WEF1) can be achieved simultaneously. In this section, we surprisingly show that this compatibility does not hold when we move to asymmetric entitlements and in particular we show an even stronger impossibility result: there exists no randomized allocation that satisfies ex-ante WEF and ex-post WWEF1.

Theorem 4.1. *Even for the case of two agents, there exists no randomized allocation that satisfies ex-ante WEF and ex-post WWEF1.*

Proof. Consider an allocation problem with two agents such that $w_1 = 0.6$ and $w_2 = 0.4$ and two items o_1 and o_2 such that $u_i(o_1) = u_i(o_2)$ for each agent $i \in [2]$.

Suppose for a contradiction that there is a WWEF1 allocation that gives both items to agent 1. Clearly, agent 2 has weighted envy for agent 1's allocation (whereas agent 1 has no weighted envy towards agent 2), so it must be that either removing one of agent 1's items or copying one of agent 1's items and giving it to agent 2 eliminates her envy. In the first case, if an item is removed from agent 1, then agent 2 is clearly still envious as she has no items. In the second case, due to symmetry suppose without loss of generality that agent 2 receives o_1. Then, her weight-adjusted utility is $u_2(o_1)/0.4 = 2.5 \cdot u_2(o_1)$, while her perception of agent 1's weight adjusted utility is $2 \cdot u_2(o_1)/0.6 > 3 \cdot u_2(o_1)$. Hence, agent 2 is still envious of agent 1, which is a contradiction. So there is no WWEF1 allocation that gives both items to agent 1.

Similarly, we can show that no WWEF1 allocation exists that gives both items to agent 2. Suppose for a contradiction that there is a WWEF1 allocation that gives both items to agent 2. In this case, agent 1 has weighted envy for agent 2's allocation, so it must be that either removing one of agent 2's items or copying one of agent 2's items and giving it to agent 1 eliminates her envy. If an item is to be removed from agent 2, then agent 1 is clearly still envious as she has no items. Alternatively, if agent 1 receives, without loss of generality due to symmetry, item o_1, her weight-adjusted utility is $u_1(o_1)/0.6 = \frac{5}{3} \cdot u_2(o_1)$, while her perception of agent 2's weight adjusted utility is $2 \cdot u_2(o_1)/0.4 = 5 \cdot u_2(o_1)$. Hence, agent 1 is still envious of agent 2, which is a contradiction. So there is no WWEF1 allocation that gives both items to agent 2.

So, the only WWEF1 allocation must give one item to each agent. At the same time, any WEF fractional allocation must give agent 1 a greater share of the items than agent 2 for agent 1 not to feel weighted envy towards agent 2, since agent 1 has a larger entitlement. Thus, such a random allocation can never be decomposed as a convex combination of WWEF1 deterministic allocations. □

Since WWEF1 is weaker than WEF1, we get the following Corollary.

Corollary 4.2. *Even for the case of two agents, there exists no randomized allocation that satisfies ex-ante WEF and ex-post WEF1.*

5 WEIGHTED PS-LOTTERY ALGORITHM

In the previous section, we show that randomized allocations that are ex-ante WEF and ex-post WWEF1 are not guaranteed to exist. Here, we show that if we weaken WWEF1 to WEF1T, then we get a possibility result. More formally, we show that any instance admits a randomized allocation that is ex-ante WEF and ex-post WEF1T. To prove this result we use Birkhoff's decomposition theorem [8, 20].

Theorem 5.1 (Birkhoff-von Neumann). *For any fractional allocation X, there exists a polynomial time algorithm that computes a decomposition $X = \sum_{k=1}^{K} \lambda^k \cdot A^k$, where each A^k is an integral allocation and $\sum_{k=1}^{K} \lambda^k = 1$.*

In Algorithm 1, we present a natural extension of the PS-Lottery algorithm [1] to the setting of agents having possibly different entitlements. The high level idea of the algorithm is as follows. For each agent i, we create c_i copies of the agent, where $c_i = \lceil w_i \cdot m \rceil$. Obviously, when all the entitlements are the same, then $c_i = c_{i'}$ for each $i, i' \in N$, as in the original algorithm. Then, we add sufficiently many dummy items, denoted by the set D, to ensure that $|O'| = |O \cup D| = \sum_{i \in N} c_i$. Each agent prefers any no dummy item than any dummy item, i.e. $o \succsim_i d$ for each $o \in O$ and each $d \in D$. Then, we run a weighted version of the Probabilistic Serial algorithm on the set of agent representatives and the set of objects including dummies. Each agent effectively eats her most preferred remaining object at a rate proportional to her entitlement until it is totally consumed, and then moves onto her next most preferred object, and so on, until all the objects are consumed by all the agents.

Rather than have the agents eat the objects directly in the Probabilistic Serial procedure, we instead have agent i's c_i representatives eat the objects on her behalf. That is, each representative i_j of agent i eats at rate $w_i m$ during the time interval $[(j-1)/(w_i m), j/(w_i m)]$. The details of this eating are implemented by the representatives themselves, which results in each representative eating exactly one item. Note that the objects eaten by the first representative of an agent will be weakly more preferred to that agent than the objects eaten by the second representative, and so on.

Then, as in the original algorithm we run Birkhoff decomposition algorithm to find a distribution over integral allocations for the agent representatives, and then each agent is assigned the parts that its copies ate. The innovative part of this algorithm is how c_i's are calculated. However, the main challenge is to prove that this algorithm satisfies the desired properties.

We now introduce some lemmas that captures an invariant property of any integral allocation that is produced in the outcome of Algorithm 1.

Lemma 5.2. *If Algorithm 1 returns the randomized allocation $\{(\lambda^k, A^k) : k = 1, \ldots, K\}$ which is implemented by the fractional allocation $X = \sum_{k=1}^{K} \lambda^k \cdot A^k$, then any integral allocation A^k is the outcome of some* sequential allocation *procedure, i.e. there exists a sequence of the agents such that if at each turn, the current agent chooses her most preferred item of the items remaining, the integral allocation produced is equivalent to A^k.*

Proof. Let Y be as it is defined in Algorithm 1, and B^k be any integral allocation implemented in Y, Then, A^k is the integral allocation obtained by combining together the allocations of representatives of the same agent from B^k.

We claim that (assuming agents have strict preferences over the objects) the following sequential allocation procedure generates the integral allocation A^k. Consider the probabilistic serial process from Algorithm 1, and for each representative i_j, let t_{i_j} denote the time at which i_j began to eat the unique item that she was ultimately allocated in B^k. Note that since B^k is consistent with Y, all of the t_{i_j} must exist - i.e. if i_j was allocated an item in B^k, it must mean that she ate at least part of that item in Y.

Order the t_{i_j}'s from first to last, settling ties arbitrarily, and then from this, induce a sequence of agents by replacing each t_{i_j} with the corresponding i_j. Call this sequence of agent representatives $a = \{a_1, a_2, \ldots, a_r\}$. Consider the allocation produced by allowing the agents to sequentially choose their most preferred remaining item at each step. We aim to prove inductively that after any number

Algorithm 1 Weighted PS-Lottery Algorithm

Input: A weighted allocation problem instance $(N, O, \succsim, w)$ where n is the number of agents, m is the number of items and the weight vector w is normalised such that $\sum_{i \in N} w_i = 1$.

$c_i \leftarrow \lceil w_i m \rceil$ for each $i \in N$.

$N' \leftarrow \{i_1, \ldots, i_{c_i} : i \in N\}$. For each agent $i \in N$, the set $\{i_1, \ldots, i_{c_i}\}$ is to be thought of as the c_i *representatives* or clones of i.

Define $D = \{d_1, \ldots, d_{\sum_{i \in N} c_i - m}\}$ to be a set of $\sum_{i \in N} c_i - m$ dummy items.

$O' \leftarrow O \cup D$ so that $|O'| = \sum_{i \in N} c_i$.

Set the preference profile $\succsim'$ on O' so that for any $o, o' \in O$ and for all $i \in N$ and $j \in \{1, \ldots, c_i\}$, $o \succsim'_{i_j} o'$ if and only if $o \succsim_i o'$. For all $o \in O$, $d \in D$ and for all $i \in N$ and $j \in \{1, \ldots, c_i\}$, $o \succsim'_{i_j} d$.

Run the Probabilistic Serial algorithm of Bogomolnaia and Moulin [9] on the set of agents N' and the set of items O', where each representative i_j eats only during the time interval $[(j-1)/(w_i m), j/(w_i m)]$ at rate $w_i m$.

Construct a fractional allocation Y of items in O' to agents in N', by assigning to each representative i_j the items that eat during the time interval $[(j-1)/(w_i m), j/(w_i m)]$.

For the (bistochastic) matrix corresponding to Y, compute a Birkhoff decomposition $Y = \sum_{k=1}^{K} \lambda^k \cdot B^k$.

Convert $Y = \sum_{k=1}^{K} \lambda^k \cdot B^k$ into $X = \sum_{k=1}^{K} \lambda^k \cdot A^k$ where all dummy items are ignored and each agent gets the allocation of her representatives .

return Allocation X and its decomposition $\sum_{k=1}^{K} \lambda^k \cdot A^k$.

of agent representatives in the sequence, the items picked by the representatives match those picked by the same representatives in B^k. For the base case (i.e. after 0 representatives from the picking sequence a have chosen), clearly, the items chosen match those chosen in B^k.

For the inductive hypothesis, suppose that after p agent representatives from the picking sequence a have selected their most preferred remaining object, the objects chosen by those representatives match the objects chosen by those same representatives in B^k. We want to show that the $(p+1)^{\text{th}}$ representative (call him x) chooses the same object as x chooses in B^k (call this o').

By our inductive hypothesis, before x chooses his most preferred item, o' must still be available, as o' cannot have been chosen by another representative for their object to match what they received in B^k. We this need to show that o' is x's most preferred remaining item.

We claim that the items that remain available for x to choose must be a subset of those that have not been fully consumed at time t_x in the construction of Y. Suppose that o' is an object that is available for x to choose, and suppose for a contradiction that at time t_x, o' has been fully consumed in Y. Then, there must be an agent representative $y \neq x$ who was allocated o' in B^k, and in Y, y must therefore have begun eating o' before t_x - i.e. $t_y < t_x$. But this contradicts our assumption that o' is available for x to choose, since y preceded x in the picking sequence and by our inductive hypothesis, must have picked o' in order to match with her choice in B^k.

Also note that of the objects available to x (in Y) at time t_x, o' must be his most preferred since he begins eating o' at this time. Therefore, o' must also be x's most preferred remaining object after p representatives have chosen. So x would choose o' and now the first $p+1$ representative's choices coincide with their allocated objects in Y.

Hence, by induction, we can conclude that B^k is equivalent to an integral allocation that is the product of sequential allocation. Then, we conclude that we can get A^k by a sequential allocation by replacing each clone by her representative. □

We proceed with the following necessary lemma.

Lemma 5.3. *If Algorithm 1 returns the randomized allocation $\{(\lambda^k, A^k) : k = 1, \ldots, K\}$ which is implemented by the fractional allocation $X = \sum_{k=1}^{K} \lambda^k \cdot A^k$, then for any integral allocation A^k, there is a picking sequence that delivers A^k such that for any agents i, j, the number of turns t_i and t_j taken by i and j, respectively, at any time after both i and j have had one pick and immediately before i is about to pick, satisfies the inequality*

$$t_j \geq \left\lfloor \frac{t_i w_j}{w_i} \right\rfloor.$$

Proof. Begin by observing that during the Probabilistic Serial process of Algorithm 1 all agent representatives will consume non-dummy items before dummy items. Therefore, at time $t = 1$, the total number of objects that will have been consumed will be $\sum_{i \in N} w_i m = m$, i.e. all of the non-dummy objects will have been totally consumed, and all non-dummy objects will be completely available. A^k must be the outcome of some sequential allocation procedure by Lemma 5.2, so there is a picking sequence that yields this integral allocation. Suppose at some point in the picking sequence (say after x turns), that agent i is to pick next, and suppose she is to pick her r^{th} item. Then, in the randomized allocation Y, this item must have been an item that was consumed partially by the r^{th} representative of i (say at time t). Recall that the r^{th} clone of i eats during the time interval $[(r-1)/(w_i m), r/(w_i m)]$. Thus, it must be that $t \in [(r-1)/(w_i m), r/(w_i m)]$. Now, after x turns in the picking sequence, we want a lower bound on the number of items that have been fully consumed by any other agent $j \neq i$. This number is at least however many representatives of j have finished eating by time t in the Probabilistic Serial process. Let $\ell = \left\lfloor \frac{(r-1) w_j}{w_i} \right\rfloor$ and consider the time t' at which j_ℓ finishes eating in the Probabilistic Serial process. We have that

$$\begin{aligned} t' &= \ell/(w_j m) \\ &= \left\lfloor \frac{(r-1) w_j}{w_i} \right\rfloor \Big/ (w_j m) \\ &\leq \frac{(r-1) w_j}{w_i} \Big/ (w_j m) \\ &= \frac{(r-1)}{w_i m} \leq t. \end{aligned}$$

So ℓ is a lower bound for the number of objects allocated to j after x turns in the picking sequence. Thus, we have that $t_j \geq \ell = \left\lfloor \frac{(r-1)w_j}{w_i} \right\rfloor = \left\lfloor \frac{t_i w_j}{w_i} \right\rfloor$, as required. □

We claim that Algorithm 1 produces a lottery over deterministic outcomes that all satisfy Weighted Envy-Freeness up to one Transfer (WEF-1T).

Theorem 5.4. *If Algorithm 1 returns X and its decomposition $\sum_{k=1}^{K} \lambda^k \cdot A^k$, then any integral allocation A^k is WEF1-T.*

Proof. From Lemma 5.3, we know that there exists a picking sequence P that gives A^k. It is sufficient to prove that if the agents pick according to P, for any agents i and j, at no point immediately after i picks an item j does experience envy towards i up to any more than one item transfer from i to j. Suppose that i is about to pick her $(r+1)^{\text{th}}$ item, denoted by o_i^*. Let t_i and t_j denote the number of turns taken by i and j so far. Also, let agent j's valuations of i's second, third, …, t_i^{th} picks be denoted by $\beta_1, \ldots, \beta_{t_i}$, and for each $x \in \{1, \ldots, t_i\}$, signify the number of items chosen by j between i's x^{th} and $(x+1)^{\text{th}}$ picks by τ_x. Then, let j's valuations of each of the items chosen in period x be given by $\alpha_1^x, \ldots, \alpha_{\tau_x}^x$.

From Lemma 5.3, we get that

$$t_j \geq \frac{t_i w_j}{w_i} - 1. \tag{1}$$

Note also that the items chosen by j in period x must naturally be more valuable to j then any items chosen by i afterwards. That is, for each $x \in \{1, \ldots, t_i\}$,

$$\sum_{y=1}^{\tau_x} \alpha_y^x \geq \tau_x \max\{\beta_x, \beta_{x+1}, \ldots, \}. \tag{2}$$

Now, let $\gamma = w_j/w_i$. We make the inductive claim that

$$\sum_{x=1}^{r}\sum_{y=1}^{\tau_x} \alpha_y^x \geq \gamma \sum_{x=1}^{r} \beta_x + \left(\sum_{x=1}^{r} \tau_x - (r\gamma - 1)\right) \max\{\beta_r, \ldots, \beta_{t_i}\} - \max\{\beta_0, \beta_1, \ldots, \beta_{t_i}\}.$$

To check that this is true, first note that the base case holds from a simple application of (2) to get that

$$\sum_{y=1}^{\tau_1} \alpha_y^1 \geq \tau_1 \max\{\beta_1, \beta_2, \ldots, \beta_{t_i}\} \geq \gamma\beta_1 + (\tau_1 - \gamma + 1)\max\{\beta_1, \ldots, \beta_{t_i}\} - \max\{\beta_1, \ldots, \beta_{t_i}\}.$$

Now, for the inductive step, utilising (1), we have

$$\begin{aligned}
\sum_{x=1}^{r}\sum_{y=1}^{\tau_x} \alpha_y^x &= \sum_{x=1}^{r-1}\sum_{y=1}^{\tau_x} \alpha_y^x + \sum_{y=1}^{\tau_r} \alpha_y^r \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - [(r-1)\gamma - 1]\right) \max\{\beta_{r-1}, \ldots, \beta_{t_i}\} \\
&\quad - \max\{\beta_0, \ldots, \beta_{t_i}\} + \sum_{y=1}^{\tau_r} \alpha_y^r \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - [(r-1)\gamma - 1]\right) \max\{\beta_{r-1}, \ldots, \beta_{t_i}\} \\
&\quad - \max\{\beta_0, \ldots, \beta_{t_i}\} + \tau_r \max\{\beta_r, \ldots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r-1} \tau_x - [(r-1)\gamma - 1]\right) \max\{\beta_r, \ldots, \beta_{t_i}\} \\
&\quad - \max\{\beta_0, \ldots, \beta_{t_i}\} + \tau_r \max\{\beta_r, \ldots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \left(\sum_{x=1}^{r} \tau_x - [(r-1)\gamma - 1]\right) \max\{\beta_r, \ldots, \beta_{t_i}\} \\
&\quad - \max\{\beta_0, \ldots, \beta_{t_i}\} \\
&= \gamma \sum_{x=1}^{r-1} \beta_x + \gamma \max\{\beta_r, \ldots, \beta_{t_i}\} + \left(\sum_{x=1}^{r} \tau_x - (r\gamma - 1)\right) \\
&\quad - \max\{\beta_r, \ldots, \beta_{t_i}\} - \max\{\beta_0, \ldots, \beta_{t_i}\} \\
&\geq \gamma \sum_{x=1}^{r} \beta_x + \left(\sum_{x=1}^{r} \tau_x - (r\gamma - 1)\right) \max\{\beta_r, \ldots, \beta_{t_i}\} \\
&\quad - \max\{\beta_0, \ldots, \beta_{t_i}\}.
\end{aligned}$$

Thus, taking $r = t_i$, we have that

$$\begin{aligned}
&\sum_{x=1}^{t_i}\sum_{y=1}^{\tau_x} \alpha_y^x \\
&\quad \geq \gamma \sum_{x=1}^{t_i} \beta_x + \left(\sum_{x=1}^{t_i} \tau_x - (r \cdot t_i - 1)\right) \beta_{t_i} - \max\{\beta_0, \ldots, \beta_{t_i}\} \\
&\quad = \gamma \sum_{x=1}^{t_i} \beta_x + \left(t_j - (w_j/w_i \cdot t_i - 1)\right) \beta_{t_i} - \max\{\beta_0, \ldots, \beta_{t_i}\} \\
&\quad \geq \gamma \sum_{x=1}^{t_i} \beta_x - \max\{\beta_0, \ldots, \beta_{t_i}\}
\end{aligned} \tag{3}$$

where the third inequality follows from Equation (1). Hence, we get that

$$\begin{aligned}
&\sum_{x=1}^{t_i}\sum_{y=1}^{\tau_x} \alpha_y^x + \max\{\beta_0, \ldots, \beta_{t_i}, u_j(o_i^*)\} \\
&\qquad \geq \sum_{x=1}^{t_i}\sum_{y=1}^{\tau_x} \alpha_y^x + \max\{\beta_0, \ldots, \beta_{t_i}\} \\
&\qquad \geq \gamma \sum_{x=1}^{t_i} \beta_x \\
&\qquad \geq \gamma \left(\sum_{x=1}^{t_i} \beta_x + u_j(o_i^*) - \max\{\beta_0, \ldots, \beta_{t_i}, u_j(o_i^*)\}\right).
\end{aligned}$$

where the third inequality follows from Equation (3). Now, the statement follows by noticing that the utility that j has for i's bundle after i chooses o_i^* is equal to

$$\sum_{x=1}^{t_i} \beta_x + u_j(o_i^*),$$

while the utility of j for her own bundle at this point is equal to

$$\sum_{x=1}^{t_i}\sum_{y=1}^{\tau_x} \alpha_y^x.$$

□

The Weighted PS-Lottery Algorithm has desirable fairness properties. It is not just ex-ante WEF but also ex-ante WEF for utilities consistent with the underlying ordinal preferences. However, the Weighted PS-Lottery Algorithm does not achieve Pareto optimality ex-post or ex-ante. In the next section, we present an alternative rule that satisfies ex-ante Pareto optimality.

6 WEIGHTED MAXIMUM NASH WELFARE LOTTERY ALGORITHM

In this section, we present an algorithm called Weighted Maximum Nash Welfare Lottery Algorithm that is a natural adaptation of an algorithm presented by Freeman et al. [17]. The algorithm first computes a fractional allocation X that maximizes the weighted Nash welfare: $X = \arg\max_p \prod_{i\in N} u_i(p)^{w_i}$. This carefully decomposes the given random allocation into a probability distribution over integral allocations using an algorithm due to Budish et al. [12].

Algorithm 2 Weighted Maximum Nash Welfare Lottery Algorithm

Input: $I = (N, M, v)$.

1: $X \leftarrow \arg\max_p \prod_{i\in N} u_i(p)^{w_i}$ Fischer market allocation for (N, O, u) with budgets of agents $(w_1, \ldots, w_n)$ (using an algorithm of Orlin [21].)

2: For any $i \in N$ and any $k \in [m]$, let $Q_{i,k} \coloneqq \sum_{t=1}^{k} X_{i,g_{i,t}}$ be the total fractional amount of the k most preferred goods assigned to agent i under X.

3: Consider the following set of bihierarchical constraints on a generic fractional allocation Y:

$$\begin{aligned} &\mathcal{H}_1 : \lfloor Q_{i,k} \rfloor \le \sum_{t=1}^{k} Y_{i,g_{i,t}} \le \lceil Q_{i,k} \rceil, \ \forall i \in N \text{ and } \forall k \in [m], \\ &\mathcal{H}_2 : \sum_{i\in N} Y_{i,g} = 1, \ \forall g \in M. \end{aligned} \tag{4}$$

4: Use the algorithm of Budish et al. [12] to find the randomized allocation $\sum_{k=1}^{K} \lambda^k A^k$ implementing the fractional allocation X that satisfies the same constraints as (4).

5: **return** Allocation X for instance I and its decomposition $\sum_{k=1}^{K} \lambda^k A^k$.

Lemma 6.1 (Utility Guarantee ++). *Given a fractional allocation X, one can compute, in strongly polynomial time, a randomized allocation implementing X whose support consists of integral allocations $A^1, \ldots, A^K$ such that for every $k \in [K]$ and every agent $i \in N$, the following hold:*

(1) If $v_i(A_i^k) < v_i(X_i)$, then $\exists\, g_i^- \notin A_i^k$ with $X_{i,g_i^-} > 0$ such that $v_i(A_i^k) + v_i(g_i^-) > v_i(X_i)$.

(2) If $v_i(A_i^k) > v_i(X_i)$, then $\exists\, g_i^+ \in A_i^k$ with $X_{i,g_i^+} < 1$ such that $v_i(A_i^k) - v_i(g_i^+) < v_i(X_i)$.

Freeman et al. [17] proved that the decomposition method as outlined in Step 3 and 4 of Algorithm 2 gives the guarantee in Lemma 6.1.

Theorem 6.2. *Algorithm 2 is a strongly polynomial-time algorithm that gives an outcome that is ex-ante WEF, and ex-post WPROP1, ex-ante Pareto optimal, and ex-post Pareto optimal.*

Proof. Consider X the ex-ante outcome of the WMNW rule: $X \in \arg\max_p \prod_{i\in N} u_i(p)^{w_i}$. Since the rule coincides with the Fischer Market Rule which gives an outcome that is competitive equilibrium with gives budgets equivalently entitlements, it implies that each agent gets a bundle that gives the agent maximum utility within his budget. Hence, the outcome satisfies WEF.[1] Hence, X satisfies ex-ante weighted proportionality as well.

Next we, prove that any ex-ante WPROP allocation can be implemented by a lottery over integral allocations all of which are WPROP1. Let X be a fractional allocation, and let $A^1, \ldots, A^K$ be integral allocations in the support of an implementation of X produced by Lemma 6.1. Suppose X satisfies WPROP. We want to show that for each $k \in [K]$, A^k is WPROP1. Since X is WPROP, for every $i \in N$, $v_i(X_i) \ge v_i(O)w_i$, where $v_i(O)$ is agent i's utility for receiving all goods fully. Fix $k \in [K]$. By Lemma 6.1, we have that for every agent $i \in N$, either $v_i(A_i^k) \ge v_i(X_i) \ge v_i(O)w_i$, or there exists a good $g \notin A_i^k$ such that $v_i(A_i^k) + v_i(g) > v_i(X_i) \ge v_i(O)w_i$. Therefore, A^k is WPROP1.

Since X maximizes the weighted Nash social welfare, it is ex-ante Pareto optimal. It follows that it each integral allocation that is used in the decomposition of X is Pareto optimal as well.

□

7 CONCLUSIONS

In this paper, we combined two research directions in fair division: (1) handling weighted entitlements and (2) best of both world fairness. Our main contribution is presenting two algorithms that have desirable ex-post and ex-ante fairness properties. Whereas the weighted PS lottery algorithm satisfies ex-ante weighted WEF with respect to all consistent cardinal utilities, it is not Pareto optimal. On the other hand, the Weighted Maximum Nash Welfare Lottery Algorithm satisfies Pareto optimality but does not satisfy ex-ante weighted WEF with respect to all consistent cardinal utilities.

We are unable to establish whether there exists an algorithm that satisfies all the following three properties: ex-ante WEF, ex-post WEF1-T, and ex-post WPROP1. Understanding whether these three properties are satisfied by some rule remains an open problem. We also suggest that the problem of best of both worlds fairness with regards to the *group fairness* axiom (an envy-based notion) of Conitzer et al. [15] would be an engaging future direction. For more information on recent developments to do with group fairness (albeit in the latter reference where items can be goods or chores), we refer the reader to Scarlett et al. [23] and Aziz & Rey [4].

REFERENCES

[1] H. Aziz. 2019. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. In *The Future of Economic Design*, J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker (Eds.). Springer-Verlag.

[2] H. Aziz. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*. Springer, 341–355.

[3] H. Aziz, R. Freeman, N. Shah, and R. Vaish. 2020. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. (2020). Preprint.

[4] H. Aziz and S. Rey. 2021. Almost Group Envy-free Allocation of Indivisible Goods and Chores. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*. 39–45.

[1] The same argument can also be extended to prove group versions of weighted envy-freeness.

[5] M. Babaioff, T. Ezra, and U Feige. 2021. Best-of-Both-Worlds Fair-Share Allocations. (2021). Preprint avalailable at https://arxiv.org/abs/2102.04909.

[6] M. Babaioff, T. Ezra, and U Feige. 2021. Fair-Share Allocations for Agents with Arbitrary Entitlements. (2021). https://arxiv.org/abs/2103.04304.

[7] N. Benabbou, M. Chakraborty, X. Ho, J. Sliwinski, and Y. Zick. 2018. Diversity Constraints in Public Housing Allocation. In *International Conference on Autonomous Agents and Multiagent Systems*. 973–981.

[8] G. Birkhoff. 1946. Three observations on linear algebra. *Univ. Nac. Tacuman, Rev. Ser. A* 5 (1946), 147–151.

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

[10] A. Bogomolnaia and H Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (October 2001).

[11] E. Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (December 2011), 1061–1103.

[12] E. Budish, Y. Che, F. Kojima, and P. Milgrom. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *The American Economic Review* 103, 2 (April 2013), 585–623.

[13] I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM-EC Conference*. 305–322.

[14] M. Chakraborty, A. Igarashi, W. Suksompong, and Y. Zick. 2021. Weighted Envy-Freeness in Indivisible Item Allocation. *ACM Transactions on Economics and Computation* 9, 18 (September 2021), 1–39.

[15] V. Conitzer, F. Freeman, and N. Shah. 2019. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the AAAI Conference on Artificial Intelligence*. https://doi.org/10.1609/aaai.v33i01.33011853, 1853–1860.

[16] A. Farhadi, M. Ghodsi, M.T. Hajiaghayi, S. Lahaie, D. Pennock, M. Seddighin, S. Seddighin, and H. Yami. 2017. Fair Allocation of Indivisible Goods to Asymmetric Agents. *Journal of Artificial Intelligence Research* 64, 1 (2017), 1–20.

[17] R. Freeman, N. Shah, and R. Vaish. 2020. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. In *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia (Eds.). ACM, 21–22. https://doi.org/10.1145/3391403.3399537

[18] L. Gourvès, J. Monnot, and L. Tlilane. 2014. Near fairness in matroids. In *Proceedings of the Twenty-first European Conference on Artificial Intelligence.* 393–398.

[19] R.J. Lipton, E. Markakis, Mossel E., and A. Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*. https://doi.org/10.1145/988772.988792, 125–131.

[20] J. Von Neumann. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games* 2, 0 (1953), 5–12.

[21] J.B. Orlin. 2010. Improved Algorithms for Computing Fisher's Market Clearing Prices. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*. 291–300.

[22] A. D. Procaccia and J. Wang. 2014. Fair enough: Guaranteeing approximate maximin shares. *J. ACM* 65, 2 (April 2014), 1–27.

[23] J. Scarlett, N. Teh, and Y. Zick. 2021. For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods. In *IJCAI 2021 Workshop on AI for Social Good*. 21–22.

[24] H. Steinhaus. 1949. Sur la division pragmatique. *Econometrica* 17, (supplement) (1949), 315–319.

[25] W. Suksompong. 2021. Constraints in Fair Division. *ACM SIGecom Exchanges* 19, 2 (November 2021), 46–61.