

Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation

Rupert Freeman
University of Virginia
FreemanR@darden.virginia.edu

Nisarg Shah
University of Toronto
nisarg@cs.toronto.edu

Rohit Vaish
Rensselaer Polytechnic Institute
rohitv@alum.iisc.ac.in

Abstract

We study the problem of allocating indivisible goods among agents with additive valuations. When randomization is allowed, it is possible to achieve compelling notions of fairness such as envy-freeness, which states that no agent should prefer any other agent's allocation to her own. When allocations must be deterministic, achieving exact fairness is impossible but approximate notions such as envy-freeness up to one good can be guaranteed. Our goal in this work is to achieve both simultaneously, by constructing a randomized allocation that is exactly fair ex-ante and approximately fair ex-post. The key question we address is whether ex-ante envy-freeness can be achieved in combination with ex-post envy-freeness up to one good. We settle this positively by designing an efficient algorithm that achieves both properties simultaneously. If we additionally require economic efficiency, we obtain an impossibility result. However, we show that economic efficiency and ex-ante envy-freeness can be simultaneously achieved if we slightly relax our ex-post fairness guarantee. On our way, we characterize the well-known Maximum Nash Welfare allocation rule in terms of a recently introduced fairness guarantee that applies to groups of agents, not just individuals.

1 Introduction

Allocating resources in non-monetary environments is commonplace in social institutions. Estates must be divided among beneficiaries, jointly held assets split between partners in a divorce, tasks assigned to employees, and educational resources distributed among public schools. Versions of this resource allocation problem - ranging from applied to theoretical models - have been studied in various research communities including mathematics, economics, operations research, and computer science. A central concern in these settings is *fairness*: how can these decisions be made so as to not systematically disadvantage individuals or groups of people?

In this work, we consider a fundamental problem in fair resource allocation. A set of m indivisible goods are to be divided among a set of n agents who have different preferences over the goods. An agent's preferences are expressed through a valuation function that assigns a value to every subset of the goods. We restrict ourselves to the case of additive valuation functions, in which each agent has a value for each good and her value for a subset is simply the sum of her values for the goods in the subset. Additive valuations ignore complementary and substitutive effects that may occur in practice, but they are an appealing tradeoff between simplicity and expressiveness. For instance, the Adjusted Winner procedure [19] that is commonly used in dispute resolution assumes additive valuations, as does the popular fair division website Spliddit (<http://spliddit.org>).

A particularly appealing notion of fairness is that of *envy-freeness* (EF) [31], which requires that no agent values the resources allocated to another agent more than the resources allocated to herself. When allocations are deterministic, it is not always possible to achieve envy-freeness; imagine two agents liking a single good, which must be given to one of them, leaving the other envious. The fundamental unfairness present in deterministic allocations motivates the use of randomized mechanisms instead. Indeed, many mechanisms used in practice exploit lotteries to, for example, determine a priority ordering. With the power of randomization, it is easy to achieve *ex-ante* envy-freeness: if a single agent is chosen uniformly at random and then allocated all the goods, no agent envies any other in expectation. However, this allocation induces a large amount of envy *ex-post*, since one agent receives everything and all others receive nothing.

While we cannot eliminate ex-post envy, a significant body of work has focused on defining and achieving approximate notions fairness for deterministic allocations [18]. One compelling notion is *envy-freeness up to one good* (EF1) [22, 39], which requires that the envy of any agent towards another agent can be eliminated by the removal of at most one good from the envied agent’s allocation. A deterministic EF1 allocation can always be achieved. For instance, it is known that the round-robin method — where the agents choose goods one at a time in a repeating fixed order — is guaranteed to output an EF1 allocation. Agents who come later in the ordering may envy those who come early, but only up to a single good. One might hope that if the order of the agents was chosen randomly among the $n!$ possible orderings then ex-ante envy-freeness could additionally be achieved, but this is not the case (we present an example in Section 3). Other methods of achieving deterministic EF1 allocation are known [10, 21, 24, 39], but do not naturally lend themselves to exploiting randomization.

This motivates a natural question. *Can we retain envy-freeness up to one good as an ex-post guarantee, and simultaneously obtain (exact) envy-freeness ex-ante?* In other words, can we always randomize over EF1 allocations such that the resulting randomized allocation is EF? In yet other words, does the polytope of EF allocations intersect with the convex hull of EF1 allocations? We show that the answer to this natural and elegant question is *yes*. More generally, we study various combinations of ex-ante and ex-post fairness and efficiency guarantees, and identify combinations that can (and cannot) be achieved simultaneously. Our constructive results yield efficient algorithms; these improve upon prior algorithms which provide either only ex-ante or only ex-post guarantees, thus paving the way for fairer resource allocation in practice.

1.1 Our Results

Our main result is the development of a novel algorithm, *Recursive Probabilistic Serial*, that produces an ex-ante envy-free (EF) distribution over deterministic allocations that each satisfy envy-freeness up to one good (EF1). Our algorithm is an adaptation of the classic *Probabilistic Serial* (PS) algorithm of Bogomolnaia and Moulin [16] that inherits much of the desirable behavior of PS (in particular, ex-ante envy-freeness), while also allowing a simple and natural decomposition over EF1 allocations. While a naïve version of our algorithm yields a distribution over a possibly-exponential number of deterministic allocations, we show that a support-size polynomial in the number of agents and goods is sufficient, thus yielding an efficient variant.

In addition to ex-ante EF and ex-post EF1, one may want to achieve the economic efficiency notion of Pareto optimality, which states that it should be impossible to find an allocation that improves some agent’s utility without reducing any other agent’s. In Section 4 we show that it is impossible to achieve ex-ante Pareto optimality (that is, with respect to the randomized allocation) in conjunction with ex-ante EF and ex-post EF1. However, in Section 5 we show that strong ex-ante guarantees — in terms of both fairness and economic efficiency — can be achieved if we are willing to compromise on the ex-post guarantee. In particular, we are able to achieve ex-ante *group fairness* (GF) [28], which generalizes both envy-freeness and Pareto optimality, in conjunction with two ex-post fairness properties that are incomparable but are both implied by EF1: *proportionality up to one good* (Prop1) [27] and *envy-freeness up to one good more-and-less* (EF₁⁺) [9]. Our algorithm uses a rounding of the well-known maximum Nash welfare (MNW) allocation, and in Section 6, we provide a novel characterization which shows

| | Prop1 + EF ₁ ¹ | EF1 | EF1 + PO | EF1 + fPO |
|------|--------------------------------------|-----------|----------|-----------|
| Prop | ✓ | ✓ | ? | ✗ (Thm 3) |
| EF | ✓ | ✓ (Thm 2) | ? | ✗ |
| GF | ✓ (Cor 1) | ✗ (Thm 3) | ✗ | ✗ |

Table 1: Summary of our results. Each row corresponds to an ex-ante guarantee while each column corresponds to an ex-post guarantee. A ✓ indicates possibility (and polynomial-time computation), a ✗ indicates impossibility, and a ? indicates an open problem. Note that positive results propagate above and to the left, while negative results propagate below and to the right.

that this is the only allocation rule that can be used to achieve the desired properties. Table 1 summarizes our results.

Finally, in Section 7, we extend our results to the case of bads. In particular, for Recursive Probabilistic Serial, we show that the natural extension only manages to achieve ex-post EF2 (in addition to ex-ante EF), but ex-post EF1 can be recovered through a simple modification.

1.2 Related Work

A large body of work in computer science and economics has focused on finding exactly ex-ante fair randomized allocations as well as approximately fair deterministic allocations, and we cite those works as appropriate throughout the paper. Combining the two approaches was recently listed as an “interesting challenge” by Aziz [5]; however, little work has focused on this problem. Two exceptions are Aleksandrov et al. [3] and Budish et al. [23]. Aleksandrov et al. [3] consider randomized allocation mechanisms for an online fair division problem and analyze their ex-ante and ex-post fairness guarantees. The style of their results is very similar to ours, however they restrict attention to binary utilities, which simplifies the problem significantly. Budish et al. [23] study the problem of implementing a general class of random allocation mechanisms subject to ex-post constraints, although the ex-post constraints are not the same as ours. In particular, they do not consider ex-post axiomatic guarantees from the fair division literature as we do.

In the random assignment literature in economics, the idea of constructing a fractional assignment and implementing it as a lottery over pure assignments was introduced by Hylland and Zeckhauser [37]. Later work has studied both ex-ante and ex-post fairness and efficiency guarantees provided by mechanisms in this setting [1, 16, 26, 40], but most of this work studies ordinal utilities and does not consider approximate notions of ex-post fairness. (The standard random assignment setting has n agents, n items, and requires that each agent receive exactly one item. The notion of ex-post fairness that we use in this work, envy-freeness up to one good, is vacuous in this restricted setting.) Gajdos and Tallon [32] study the relationship between ex-ante and ex-post fairness but in their model the randomness comes from nature, not the allocation rule. Other work studies the problem of implementing a fractional outcome over deterministic outcomes subject to (possibly soft) constraints [2, 23], but the constraints allowed by these papers do not fully capture our ex-post fairness notions.

In Section 6, we obtain a characterization of the maximum Nash welfare (MNW) rule — this is equivalent to competitive equilibrium with equal incomes (CEEI) [4, 30], and is also known as the competitive rule or the Walras rule — for fractional allocations using group fairness and a property called replication-invariance. There are several related characterizations in the literature. Assuming monotonic and strictly convex preferences, it is known that CEEI is the only replication-invariant allocation rule that is in the core [29], or group envy-free [47], or proportional and Pareto optimal [45]. While the core, group envy-freeness, or proportionality together with Pareto optimality are all weaker than group fairness (making these characterizations seemingly stronger), to the best of our knowledge, these characterizations only apply for *strictly convex* preferences, and not to additive preferences. We also note that our proof technique is significantly different from the proofs of these prior characterizations because ad-

ditive preferences do not admit some of the nice structural implications that strictly convex preferences do when combined with replication invariance, as we note in Section 6.

2 Preliminaries

For any positive integer $r \in \mathbb{N}$, define $[r] := \{1, \dots, r\}$. Let $N = [n]$ denote a set of *agents*, and M denote a set of *goods* where $m := |M|$.

Fractional and Randomized Allocations. A *fractional allocation* of the items in M to the agents in N is specified by a non-negative $n \times m$ matrix $A \in [0, 1]^{n \times m}$ such that for every item $r \in M$, we have $\sum_{i \in N} A_{i,r} \leq 1$; here, $A_{i,r}$ denotes the fraction of item r assigned to agent i . We say that a fractional allocation A is *complete* if $\sum_{i \in N} A_{i,r} = 1$ for every $r \in M$, and call it *partial* otherwise.

A fractional allocation A is *integral* if $A_{i,r} \in \{0, 1\}$ for every $i \in N$ and $r \in M$. For integral allocations, we will find it convenient to denote the binary vector $A_i = (A_{i,r})_{r \in M}$ as a set $A_i := \{j \in M : A_{i,r} = 1\}$. We will refer to A_i as the *bundle* of items assigned to agent i , and denote the allocation A as an ordered tuple of bundles $A = (A_1, \dots, A_n)$. When we simply say ‘an allocation’, it will mean a fractional allocation, unless otherwise clear from the context. For notational clarity, we use letters X or Y for fractional allocations, and write A or B for integral allocations. Let \mathcal{X} be the set of all (complete or partial) fractional allocations.

A *randomized allocation* is a lottery over integral allocations (we denote them by bold letters for clarity). Formally, a randomized allocation \mathbf{X} is specified by a set of $\ell \in \mathbb{N}$ ordered pairs $\{(p^k, A^k)\}_{k \in [\ell]}$, where, for every $k \in [\ell]$, A^k is an integral allocation implemented with probability $p^k \in [0, 1]$, and $\sum_{k \in [\ell]} p^k = 1$. The *support* of \mathbf{X} is the set of integral allocations $\{A^1, \dots, A^\ell\}$.

A randomized allocation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ is naturally associated with the fractional allocation $X := \sum_{k \in [\ell]} p^k A^k$, where $X_{i,r}$ is the (marginal) probability of agent i receiving good r under \mathbf{X} . In this case, we say that randomized allocation \mathbf{X} *implements* fractional allocation X . There may be many randomized allocations implementing a given fractional allocation. In Sections 3 and 5, we discuss and make use of two known techniques for implementing a given fractional allocation.

Preferences. Each agent $i \in N$ has an *additive* valuation function v_i , where $v_{i,r} \geq 0$ denotes the agent’s utility for fully receiving good $r \in M$. Later, in Section 7, we consider allocation of *bads*, where $v_{i,r} \leq 0$ for each agent $i \in N$ and bad $r \in M$. Note that v_i induces a weak order \succsim_i over goods where $g_r \succsim_i g_s$ if and only if $v_{i,r} \geq v_{i,s}$, and $g_r \succ_i g_s$ if and only if $v_{i,r} > v_{i,s}$. The utility of agent i under an allocation $X \in \mathcal{X}$ is given by $v_i(X_i) = \sum_{r \in M} X_{i,r} \cdot v_{i,r}$. We assume that for each good $r \in M$, there exists at least one agent $i \in N$ with $v_{i,r} > 0$. This is without loss of generality as goods valued zero by everyone can be allocated arbitrarily.

Allocation rule. A fair division instance I is defined by the triple $(N, M, (v_i)_{i \in N})$. We let \mathcal{I} denote the set of all instances. An *allocation rule* $f : \mathcal{I} \rightarrow 2^{\mathcal{X}}$ maps instances to (sets of) allocations.

Definition 1 (Fractional Maximum Nash Welfare Rule). Given an instance $I \in \mathcal{I}$, the fractional Maximum Nash Welfare (MNW) rule returns all fractional allocations that maximize the product of agents’ utilities, i.e., $\text{MNW}(I) := \arg \max_{X \in \mathcal{X}} \prod_{i \in N} v_i(X_i)$. We refer to an allocation $A \in \text{MNW}(I)$ as a fractional MNW allocation. Integral MNW allocations maximize the product of agents’ utilities across all integral allocations; however, there is a subtle tie-breaking involved in that case [24].

We now discuss a number of properties concerning fairness and efficiency of allocations and allocation rules.

2.1 Fairness and Efficiency Properties

For any property $\langle P \rangle$ defined for a fractional allocation, we say that a randomized allocation \mathbf{X} satisfies $\langle P \rangle$ *ex-ante* if the fractional allocation X it implements satisfies $\langle P \rangle$. Similarly, for any property $\langle Q \rangle$ defined for an integral allocation, we say that a randomized allocation \mathbf{X} satisfies $\langle Q \rangle$ *ex-post* if every integral allocation in its support satisfies $\langle Q \rangle$.

We begin with a classic fairness notion called *proportionality*.

Definition 2 (Proportionality (Prop); [44]). An allocation X is *proportional* if for each agent $i \in N$, $v_i(X_i) \geq v_i(\bar{1}^m)/n$, where $v_i(\bar{1}^m)$ is agent i 's utility for receiving all goods fully.

For complete allocations, the following fairness guarantee is logically stronger than proportionality, and has been a subject of substantial research.

Definition 3 (Envy-Freeness (EF); [31]). An allocation X is *envy-free* if for every pair of agents $i, j \in N$, we have $v_i(X_i) \geq v_i(X_j)$.

Given allocations X and Y , we say that agent i *SD-prefers* X_i to Y_i , written $X_i \succsim_i^{SD} Y_i$, if for every good $g \in M$, we have that $\sum_{g_r \in \{g' \in M: g' \succsim_i g\}} X_{i,r} \geq \sum_{g_r \in \{g' \in M: g' \succsim_i g\}} Y_{i,r}$. Here, "SD" refers to first-order stochastic dominance. It is easy to check that $X_i \succsim_i^{SD} Y_i$ is equivalent to $v'_i(X_i) \geq v'_i(Y_i)$ under every additive valuation v'_i consistent with the ordinal preference relation \succsim_i .

Definition 4 (SD-Envy-Freeness (SD-EF); [16]). An allocation X is *SD-envy-free* if for every pair of agents $i, j \in N$, we have $X_i \succsim_i^{SD} X_j$.

Equivalently, an allocation is SD-envy-free if it is envy-free under any additive valuation functions of the agents consistent with $(\succsim_i)_{i \in N}$. Next, we discuss economic efficiency of allocations.

Definition 5 (Fractional Pareto Optimality (fPO) and Pareto Optimality (PO)). An allocation X is *fractionally Pareto optimal* if there is no fractional allocation Y that Pareto-dominates it, i.e., satisfies $v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in N$ and at least one inequality is strict. An integral allocation A is *Pareto optimal* if there is no integral allocation B that Pareto-dominates it.

For fractional allocations, fPO is traditionally just referred to as Pareto optimality, so we will use fPO and PO interchangeably. However, for integral allocations, fPO is stronger than PO [10].

Proposition 1. *If a randomized allocation is ex-ante PO, then it is also ex-post fPO.*

Proof. Proof. If a randomized allocation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ implementing a fractional allocation X is not ex-post fPO, then for some $k \in [\ell]$, integral allocation A^k must be Pareto dominated by a fractional allocation, say Y . Then, the fractional allocation $X' := p^k \cdot Y + \sum_{r \neq k} p^r \cdot A^r$ Pareto-dominates X , which implies that \mathbf{X} is not ex-ante PO. \square

The following property generalizes various properties mentioned above. Given any set $S \subseteq N$ of agents, we write $\cup_{i \in S} X_i$ to denote the union of the fractional allocations to agents in S , i.e., $\cup_{i \in S} X_i := (\sum_{i \in S} X_{i,r})_{r \in M}$.

Definition 6 (Group Fairness (GF); [28]). An allocation X is *group fair* if for all non-empty subsets of agents $S, T \subseteq N$, there is no fractional allocation Y of $\cup_{i \in T} X_i$ to the agents in S such that $\frac{|S|}{|T|} \cdot v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in S$ and at least one inequality is strict.

Note that imposing the above constraint over restricted (S, T) pairs can recover properties such as proportionality ($|S| = 1, T = N$), envy-freeness ($|S| = |T| = 1$), and fractional Pareto optimality ($S = T = N$).

While the fairness properties mentioned above are compelling, they cannot be guaranteed with integral allocations. For such allocations, it is meaningful to consider approximate fairness notions, some of which are discussed below.

Definition 7 (Proportionality Up To One Good (Prop1); [27]). An integral allocation A is *proportional up to one good* if for every agent $i \in N$, either $v_i(A_i) \geq v_i(\bar{1}^m)/n$ or there exists a good $r \notin A_i$ such that $v_i(A_i) + v_{i,r} \geq v_i(\bar{1}^m)/n$, where $v_i(\bar{1}^m)$ is agent i 's utility for receiving all goods fully.

Definition 8 (Envy-Freeness Up To One Good (EF1); [22, 39]). An integral allocation A is *envy-free up to one good* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $v_i(A_i) \geq v_i(A_j \setminus \{r\})$ for some good $r \in A_j$.

All of our possibility and impossibility results also apply for the following natural strengthening of EF1, which has been formally defined in subsequent work [7].

Definition 9 (SD-Envy-Freeness up to One Good (SD-EF1)). An integral allocation A is *SD-envy-free up to one good* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $A_i \succsim_i^{SD} A_j \setminus \{r\}$ for some good $r \in A_j$.

This is equivalent to agent i not envying agent j up to one good under every additive valuation consistent with the ordinal preference relation \succsim_i .

The next property, called envy-freeness up to one good more-and-less (EF₁¹) [9], is a relaxation of EF1 and enjoys strong algorithmic support in conjunction with PO, for goods (Section 5) as well as bads (Section 7).

Definition 10 (Envy-Freeness Up To One Good More-and-Less (EF₁¹)). An integral allocation A is *envy-free up to one good more-and-less* if for every pair of agents $i, j \in N$ such that $A_j \neq \emptyset$, we have $v_i(A_i \cup \{r_i\}) \geq v_i(A_j \setminus \{r_j\})$ for some goods $r_i \notin A_i$ and $r_j \in A_j$.

These three guarantees are logically related as follows: SD-EF1 \implies EF1 \implies EF₁¹ and SD-EF1 \implies EF1 \implies Prop1.

3 Possibility: Ex-Ante EF + Ex-Post EF1

This section describes our main result that ex-ante envy-freeness can be achieved in conjunction with ex-post envy-freeness up to one good (EF1).

A natural approach towards this question is to start with the round-robin method. Under this method, we fix an agent ordering, and then agents take turns picking one good at a time in a cyclic fashion. At each step, an agent picks her most valuable good that is still available. It is well-known that for any agent ordering, this method produces an integral allocation that is EF1 [24]. Further, it is easy to see that uniformly randomizing the agent ordering — the so-called *randomized round-robin method* — also achieves ex-ante proportionality.

Proposition 2. *With additive valuations, randomized round-robin is ex-ante proportional and ex-post envy-free up to one good.*

Proof. Proof. We already noticed that randomized round-robin is ex-post EF1. For ex-ante Prop, fix an agent i , and suppose, without loss of generality, that $v_{i,1} \geq v_{i,2} \geq \dots \geq v_{i,m}$. For ease of presentation, suppose that $m = cn$ for some positive integer c . When agent i is k -th in the ordering, her value for her allocation is at least $v_{i,k} + v_{i,k+n} + \dots + v_{i,k+(c-1)n} = \sum_{c'=1}^c v_{i,(c'-1)n+k}$ because the c' -th good she chooses must be no less preferred than her $(c' - 1)n + k$ -th most preferred good. Since she appears in every position in the ordering with probability $1/n$, her expected value is at least $\frac{1}{n} \sum_{k=1}^m \sum_{c'=1}^c v_{i,(c'-1)n+k} = \frac{1}{n} \sum_{r=1}^m g_{i,r}$, which matches her proportionality guarantee. \square

However, as we observe below, this approach fails to achieve ex-ante envy-freeness.

| | g_1 | g_2 | g_3 |
|-------|--------------------|--------------------|--------------------|
| a_1 | $1 + \varepsilon$ | 1 | $1 - 4\varepsilon$ |
| a_2 | $1 - 4\varepsilon$ | $1 + \varepsilon$ | 1 |
| a_3 | $1 + \varepsilon$ | $1 - 4\varepsilon$ | 1 |

Example 1. Consider an instance with three agents a_1, a_2, a_3 and three goods g_1, g_2, g_3 as shown below, where $0 < \varepsilon < 0.25$ (we note that a similar instance was used by Aziz and Ye [6, Proposition 13] in the divisible goods (i.e., cake cutting) setting to show that a constrained serial dictatorship mechanism violates envy-freeness):

The allocations produced by various agent orderings are as follows: The order (a_1, a_2, a_3) induces the allocation $A^1 = (\{g_1\}, \{g_2\}, \{g_3\})$, the order (a_1, a_3, a_2) induces the allocation $A^2 = (\{g_1\}, \{g_2\}, \{g_3\})$, the order (a_2, a_1, a_3) induces the allocation $A^3 = (\{g_1\}, \{g_2\}, \{g_3\})$, the order (a_2, a_3, a_1) induces the allocation $A^4 = (\{g_3\}, \{g_2\}, \{g_1\})$, the order (a_3, a_1, a_2) induces the allocation $A^5 = (\{g_2\}, \{g_3\}, \{g_1\})$, and the order (a_3, a_2, a_1) induces the allocation $A^6 = (\{g_3\}, \{g_2\}, \{g_1\})$.

Let $A := \frac{1}{6}(A^1 + \dots + A^6)$ denote the fractional allocation implemented by randomized round-robin. Notice that $v_1(A_1) = 3(1 + \varepsilon) + 1 + 2(1 - 4\varepsilon) = 6 - 5\varepsilon$ and $v_1(A_2) = 5 + (1 - 4\varepsilon) = 6 - 4\varepsilon$, implying that a_1 envies a_2 .

Let us revisit our motivating example from the introduction, where the same set of goods need to be allocated repeatedly over time. As we discussed, a seemingly promising solution is to use round-robin every time, and cycle through all $n!$ possible agent orderings to eliminate the “advantage” that any agent has over another in a specific ordering. In the long run, the allocation produced would resemble the fractional allocation returned by randomized round-robin. While Proposition 2 shows that this would guarantee each agent her proportional share in the long run, Example 1 points out that envy may not vanish, even in the long run.

Instead of starting from a method that guarantees ex-post EF1 and using it to achieve ex-ante EF, let us do the opposite: Start from a fractional EF allocation and implement it using integral EF1 allocations. *Probabilistic serial* is a well-studied algorithm that produces a fractional envy-free allocation [16]. This algorithm uses a *serial eating* protocol wherein all agents simultaneously start eating their respective favorite goods at the same constant speed. Once a good is completely consumed by a subset of agents, each of those agents proceeds to eating her favorite available good at the same speed. The algorithm terminates when all goods have been eaten, and the fraction of each good consumed by an agent is allocated to her. A useful property of this algorithm is that it only uses the ordinal preferences of agents over goods, and computes an allocation that is ex-ante envy-free for *any* additive utilities consistent with the ordinal preferences; thus, it is *SD-envy-free*.

In Section 3.1, we present a variant of this procedure, which we call *recursive probabilistic serial*, that provides the desired fairness guarantees, i.e., ex-ante EF and ex-post EF1. While this procedure efficiently produces a sample integral allocation, computing the full randomized allocation it returns may take exponential time. Later, in Section 3.2, we describe a modification of this procedure that runs in strongly polynomial time.

It is worth noting that in a follow-up to our work, Aziz [7] has shown that the fractional allocation produced by probabilistic serial can also be implemented using an ex-post EF1 randomized allocation. The key idea in his proof is to consider multiple “copies” of each agent—one for each unit of time—such that the i -th copy consumes the same good(s) as consumed by the agent between the time steps $t = i - 1$ and $t = i$. The desired implementation is then obtained by applying Birkhoff-von Neumann theorem to the “expanded” allocation matrix whose columns comprise of the goods, and rows comprise of all copies of each agent.

Algorithm 1 Ex-ante EF + ex-post EF1

```
1:  $B \leftarrow (\emptyset, \dots, \emptyset)$  ▷ Initialize the allocation
2:  $M^0 \leftarrow M$  ▷ Initialize the set of available goods
3: for  $t = 1$  to  $\lceil \frac{m}{n} \rceil$  do
4:    $X^t \leftarrow \text{EATING}(M^t, 1)$  ▷ Run the eating protocol for one unit of time
5:   BvN:  $X^t = \sum_{k=1}^{\ell} w^{t,k} A^{t,k}$  ▷ Decomposition from Theorem 1
6:    $B^t \leftarrow A^{t,k}$  with probability  $w^{t,k}$  ▷ Sampling step
7:    $M^{t+1} \leftarrow M^t \setminus \bigcup_{i \in N} B_i^t$  ▷ Fix the assignments according to the sampled allocation
8:    $B_i \leftarrow B_i \cup_{t=1}^{\lceil \frac{m}{n} \rceil} B_i^t$  for all  $i \in N$  ▷ Update integral allocation
9: end for
10: return  $B$ 
```

3.1 Recursive Probabilistic Serial

The eating protocol used by the probabilistic serial algorithm can be formally defined as follows: Given as input a subset of goods $M' \subseteq M$, the algorithm allows each agent to consume its favorite available good in M' out of those that have not yet been fully consumed at a rate of one good per unit time. Let $X^t := \text{EATING}(M', t)$ denote the (partial) allocation obtained when the algorithm is run for t units of time. Note that $\text{EATING}(M, \lceil \frac{m}{n} \rceil)$ is equivalent to the outcome of the probabilistic serial algorithm [16].

We will make use of the following classic theorem [15, 49], stated in a general form.

Theorem 1 (Birkhoff-von Neumann). *Let X be a real-valued $n \times m$ matrix such that*

1. $X_{i,r} \in [0, 1]$ for every $i \in [n]$ and $r \in [m]$,
2. $\sum_{r=1}^m X_{i,r} \leq 1$ for every $i \in [n]$, and
3. $\sum_{i=1}^n X_{i,r} \leq 1$ for every $r \in [m]$.

Then, in strongly polynomial time, one can compute integral matrices A^1, A^2, \dots, A^q and weights $w^1, w^2, \dots, w^q \in [0, 1]$ such that

1. $\sum_{k=1}^q w^k = 1$ and $X = \sum_{k=1}^q w^k A^k$,
2. for each integral matrix A^k , we have $A_{i,r}^k \in \{0, 1\}$ for every $i \in [n]$ and $r \in [m]$, and
3. for every $k \in [q]$, we have
 - (a) $\lfloor \sum_{r=1}^m X_{i,r} \rfloor \leq \sum_{r=1}^m A_{i,r}^k \leq \lceil \sum_{r=1}^m X_{i,r} \rceil$ for every $i \in [n]$, and
 - (b) $\lfloor \sum_{i=1}^n X_{i,r} \rfloor \leq \sum_{i=1}^n A_{i,r}^k \leq \lceil \sum_{i=1}^n X_{i,r} \rceil$ for every $r \in [m]$.

We will now describe our algorithm. The algorithm proceeds in stages. In each stage, the EATING procedure is run for *one* unit of time, after which every agent has consumed a total mass of one item (although an individual item may be fractionally consumed). Then, Theorem 1 is used to decompose this fractional (partial) allocation into integral allocations. By implication (3a) of Theorem 1, each agent receives exactly one item in each integral allocation in the support. We interpret the weights w^1, \dots, w^q as probabilities, and sample an integral allocation from this distribution before proceeding to the next stage. The algorithm then fixes the assignments of the goods to the agents in accordance with the aforementioned integral partial allocation, and repeats the eating procedure with the reduced set of goods. In the final stage of the algorithm, fewer than n goods might remain, in which case some agents do not receive any good under the integral allocation. The full procedure is described in Algorithm 1. Note that Algorithm 1 is a randomized algorithm that returns an integral allocation in every execution. One could, however, also consider the implied fractional allocation, which is obtained by taking expectation over the randomness in Line 6. Algorithm 3 in Section 9.1 provides the pseudocode for a recursive formulation of the algorithm. Example 2 illustrates the execution of our algorithm.

which establishes EF1. \square

Theorem 2. *The randomized allocation implemented by Algorithm 1 is ex-ante envy-free and ex-post EF1.*

Proof. Proof. The ex-post EF1 guarantee follows readily from Lemma 2. To prove ex-ante EF, notice that each agent eats her most preferred available good at each point in time during the EATING algorithm. Therefore, for any iteration t , the partial fractional allocation X^t computed by the EATING(\cdot) procedure in iteration t is envy-free.

Fix a pair of agents $i, j \in N$. We want to show that $\mathbb{E}[v_i(B_i)] \geq \mathbb{E}[v_i(B_j)]$. Note that $B_i = \bigcup_{t=1}^{\lceil m/n \rceil} B_i^t$ and $B_j = \bigcup_{t=1}^{\lceil m/n \rceil} B_j^t$. Due to additive preferences and linearity of expectation, it is sufficient to show that for every t , $\mathbb{E}[v_i(B_i^t)] \geq \mathbb{E}[v_i(B_j^t)]$. Note that conditioned on B^1, \dots, B^{t-1} , the expected value of any agent for B_i^t is the value of that agent for X_i^t . Hence, for each possible value of B^1, \dots, B^{t-1} , envy-freeness of X^t implies that

$$\begin{aligned} \mathbb{E}[v_i(B_i^t)|B^1, \dots, B^{t-1}] &= \mathbb{E}[v_i(X_i^t)|B^1, \dots, B^{t-1}] \geq \mathbb{E}[v_i(X_j^t)|B^1, \dots, B^{t-1}] \\ &= \mathbb{E}[v_i(B_j^t)|B^1, \dots, B^{t-1}], \end{aligned}$$

as desired. \square

Note that like probabilistic serial, Recursive Probabilistic Serial (Algorithm 1) takes only the ordinal preferences of agents over goods as input; thus, it is in fact ex-ante SD-EF and ex-post SD-EF1.

3.2 A Polynomial-Time Algorithm for Ex-Ante EF and Ex-Post EF1

Note that recursive probabilistic serial (Algorithm 1) efficiently produces a sample integral EF1 allocation, but computing the full randomized allocation constructed by it would require exponential time. This is because at each iteration, the algorithm “branches out” into a polynomial number of subinstances (equal to the size of the support in the BvN decomposition) and this happens polynomially many times. Also, for the same reason, the randomized allocation it constructs can, in the worst case, have an exponential number of integral allocations in the support.

A natural approach towards reducing the support size (in the hope of reducing the overall running time) is to use the well-known Carathéodory’s theorem from convex analysis, which states that if a point $x \in \mathbb{R}^d$ lies in the convex hull of a set P , then x can be expressed as a convex combination of at most $d + 1$ vertices of P . Since any fractional allocation X is a point in an $n \times m$ -dimensional space (i.e., $X \in \mathbb{R}^{n \times m}$), by Carathéodory’s theorem, if it is a convex combination of integral allocations in the support $\{A^1, \dots, A^\ell\}$, then it can also be written as a convex combination of at most $mn + 1$ allocations from this set. While this establishes the *existence* of a convex combination with small support, it does not automatically provide an efficient *algorithm* for converting a convex combination with large support to one with small support. Indeed, reducing the support size to at most $mn + 1$ once it has already grown exponentially may not be possible efficiently. Instead, in our modification of recursive probabilistic serial, we ensure that the support size remains $mn + 1$ after every step of “branching out”. Note that every step of “branching out” can increase the support size from $mn + 1$ to $(mn + 1) \cdot p(m, n)$, where $p(m, n)$ is a polynomial that bounds the size of the support produced by the BvN decomposition. Thus, we need a way to reduce the support size from at most $(mn + 1) \cdot p(m, n)$ to at most $mn + 1$. This is where the following beautiful result comes in handy.

Proposition 3 (Grötschel et al. 36). *There is a polynomial-time algorithm that, given a polytope \mathcal{P} with a strong separation oracle and a point $y \in \mathcal{P}$, returns a representation of y as a convex combination of $\dim(\mathcal{P}) + 1$ vertices of \mathcal{P} .*

We refer the reader to the aforementioned paper for formal definitions of a polytope, its dimension, and a strong separation oracle. Let us describe in a bit more detail how we use this result. It is instructive to revisit the first two iterations of Algorithm 1. Notice that at the end of the first iteration ($t = 1$), the fractional partial allocation X^1 consists of a polynomial number of integral allocations in its support, i.e., $X^1 = \sum_{k=1}^{\ell} w^{1,k} A^{1,k}$, where $\ell = p(m, n)$ for some polynomial p . Consider the polytope \mathcal{P}^1 formed by the convex hull of the integral allocations $A^{1,1}, \dots, A^{1,\ell}$ as follows:

$$\mathcal{P}^1 := \{z : z = \sum_k \alpha_k A^{1,k} \text{ where } \alpha_k \geq 0 \text{ for all } k \in [\ell] \text{ and } \sum_k \alpha_k = 1\}.$$

Observe that \mathcal{P}^1 admits a polynomial-time strong separation oracle. That is, one can decide in polynomial time whether a given point $y \in \mathbb{R}^{n \times m}$ belongs to \mathcal{P}^1 , or return a separating hyperplane if it doesn't (this can be done via linear programming). This implies that an $(mn+1)$ -sized decomposition of X^1 , as guaranteed by Carathéodory's theorem, can be computed in polynomial time.

Having computed an $(mn + 1)$ -sized decomposition of X^1 , we now obtain at most $mn + 1$ subinstances in each of which we re-run our algorithm. Notice that doing so results in a probability distribution over $(mn + 1) \cdot p(m, n)$ deterministic allocations at the end of the second iteration. We can now re-invoke the aforementioned algorithmic version of Carathéodory's theorem to once again obtain a decomposition over at most $mn + 1$ of these integral allocations in polynomial time. It is easy to see that this argument in fact holds for all time steps via induction. Algorithm 4 in Section 9.1 summarizes the polynomial-time implementation of our algorithm.

Notice that the randomized allocation thus produced has support that is a subset of the support of the randomized allocation produced by recursive probabilistic serial. Hence, the ex-post EF1 guarantee of this modified procedure follows immediately from Lemma 1. Finally, it is also easy to verify that the randomized allocation is also ex-ante EF by essentially the same argument as in the proof of Theorem 2: Each individual step of the eating procedure produces an envy-free allocation, so even if we change the convex combination over such allocations, it remains envy-free.

Again, like Recursive Probabilistic Serial (Algorithm 1), this algorithm takes only the ordinal preferences of agents over goods as input; thus, it is in fact ex-ante SD-EF and ex-post SD-EF1.

4 Impossibility: Ex-Ante Prop + Ex-Post EF1 + Ex-Post fPO

In the previous section, we showed that ex-ante EF and ex-post EF1 can be achieved simultaneously in polynomial time. The obvious next question, then, is whether we can achieve stronger guarantees. One property that our algorithms from Section 3 lack is efficiency. We defined three efficiency notions in Section 2 that are related through the following logical implications: ex-ante PO \Rightarrow ex-post fPO \Rightarrow ex-post PO.

Let us consider adding the weakest of them: ex-post PO. Unfortunately, we were not able to settle whether ex-ante EF (or even ex-ante Prop) is compatible with ex-post EF1 and ex-post PO. The difficulty is that we do not understand the structure of integral EF1+PO allocations; we elaborate on this issue in Section 8 and lay out several more open questions that stem from this.

Let us move on and consider imposing a slightly stronger efficiency notion: ex-post fPO. We note that integral EF1+fPO allocations are known to always exist [10]. Hence, the question of whether we can randomize over such allocations to achieve a desirable ex-ante fairness guarantee is meaningful. However, in this case, we show that even achieving ex-ante proportionality is impossible along with ex-post EF1 and ex-post fPO.

Theorem 3. *There exists an instance with additive valuations in which no randomized allocation is simultaneously ex-ante proportional, ex-post envy-free up to one good, and ex-post fractionally Pareto optimal.*

Proof. Proof. We present an instance in which the unique integral allocation satisfying EF1+fPO violates proportionality. Specifically, consider an instance with two goods (g_1, g_2) and two agents (a_1, a_2) whose additive valuations are as follows: $v_{1,1} = 1, v_{1,2} = 2$ and $v_{2,1} = 1, v_{2,2} = 3$. This instance has exactly two integral EF1 allocations: $A := (\{g_1\}, \{g_2\})$ and $B := (\{g_2\}, \{g_1\})$. It is easy to check that B is not fPO, since it is Pareto dominated by a fractional allocation X that assigns g_1 completely to a_1 , and splits g_2 equally between the two agents. Indeed, $v_1(X_1) = v_1(g_1) + 0.5 \cdot v_1(g_2) = 2 \geq v_1(B_1)$ and $v_2(X_2) = 0.5 \cdot v_2(g_2) = 1.5 > v_2(B_2)$. To see why A is fPO, notice that it assigns each good to an agent that has the highest valuation for it. Therefore, A maximizes the utilitarian social welfare (i.e., sum of agents’ utilities) and is therefore fPO.

Finally, observe that A violates proportionality since $v_1(A_1) = v_1(g_1) < \frac{1}{2}v_1(g_1 \cup g_2)$. Therefore, any randomized allocation that is ex-post EF1+fPO must be supported entirely on the integral allocation A , in which case it violates ex-ante Prop, as desired. \square

For example, since ex-ante envy-freeness and ex-ante group fairness imply ex-ante proportionality, they are also incompatible with ex-post EF1+fPO. Similarly, since ex-ante PO implies ex-post fPO (Proposition 1), we also get that ex-ante Prop+PO and ex-ante EF+PO are incompatible with ex-post EF1.

We mentioned earlier that achieving ex-ante Prop or ex-ante EF along with ex-post EF1+PO is an open question. Two prominent methods for finding an integral EF1+PO allocation are the integral MNW rule [24] and the market-based rule of Barman et al. [10]. An interesting implication of Theorem 3 is that we cannot hope to achieve even ex-ante proportionality by randomizing over allocations returned by either method. The latter method is guaranteed to return an integral EF1+fPO allocation, so Theorem 3 directly applies. And the MNW rule, while only guaranteed to return an integral EF1+PO allocation, uniquely returns allocation A in the example presented in the proof of Theorem 3, which violates proportionality.

5 Possibility: Ex-Ante GF + Ex-Post Prop1 + Ex-Post EF₁¹

Given the impossibility of ex-ante Prop and ex-post EF1+fPO derived in the previous section, it is evident that retaining efficiency would require relaxing at least one of the fairness guarantees. Ex-ante Prop is already a very weak fairness guarantee. So, in this section, we focus on relaxing ex-post EF1. There are two relaxations that have been proposed in the literature — namely, Prop1 and EF₁¹. We show that both of these can be achieved simultaneously, and in fact, this can be done while strengthening ex-ante Prop and ex-post fPO to ex-ante GF. Recall that ex-ante GF implies not only ex-ante Prop, but also ex-ante PO, which, by Proposition 1, implies ex-post fPO for any implementation of it. In other words, our goal is to implement a fractional GF allocation using integral Prop1+EF₁¹ allocations.

Luckily, we know that a fractional GF allocation always exists. Conitzer et al. [28] argued that a fractional MNW allocation satisfies group fairness. They did not provide a formal proof; for completeness, we provide a proof of this result for the more general cake-cutting setting in Section 10.1. Further, we know that a fractional MNW allocation can be computed in strongly polynomial time [41, 48].

Hence, we ask whether a fractional MNW allocation can be implemented using integral Prop1+ EF₁¹ allocations. Our starting point is a result by Budish et al. [23] that allows implementing any fractional allocation using integral allocations that are very “close” to it in agent utilities. Specifically, they prove the next result deriving and using an extension of the classic Birkhoff-von Neumann theorem [15, 49].

Proposition 4 (Utility Guarantee; Theorem 9 of Budish et al. [23]). *Given any fractional allocation X , one can compute, in strongly polynomial time, a randomized allocation implementing X whose support consists of integral allocations A^1, \dots, A^ℓ such that for every $k \in [\ell]$ and every agent $i \in N$,*

$$|v_i(X_i) - v_i(A_i^k)| \leq \max\{v_i(g) - v_i(g') : 0 < X_{i,g}, X_{i,g'} < 1 \text{ and } g, g' \in M\}.$$

Notice that the upper bound established in Proposition 4 on how much agent i 's utility under an integral allocation A^k in the support can differ from her utility under the fractional allocation X depends only on the fractional allocation X . In contrast, the fairness guarantees we want to establish for the integral allocations in the support — Prop1 and EF_1^1 — consider what happens when we add a good to the bundle of agent i that agent i is *not already allocated*; in other words, we need a stronger guarantee for integral allocations in the support which depends on which goods the agent is (or is not) allocated ex-post.

It turns out that the method proposed by Budish et al. [23] already provides such a guarantee, and their proof can be adapted to establish a more nuanced bound. Specifically, we show that if the agent's ex-ante utility $v_i(X_i)$ exceeds her ex-post utility $v_i(A_i^k)$, then the gap is at most the maximum value the agent has for any good that she *lost* in the integral allocation (i.e., any good g such that $0 < X_{i,g} < 1$ and $A_{i,g}^k = 0$). Similarly, if the ex-post utility exceeds the ex-ante utility, then the gap is at most the maximum value the agent has for any good that she *gained* in the integral allocation (i.e., any good g such that $0 < X_{i,g} < 1$ and $A_{i,g}^k = 1$). We later show that this subtle improvement helps us establish the desired ex-post fairness guarantees.

Lemma 3 (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, \dots, A^ℓ such that for every $k \in [\ell]$ 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 $v_i(A_i^k) - v_i(g_i^+) < v_i(X_i)$.*

The proof of Lemma 3 is presented in Section 10.3.

We now show how Lemma 3 can be used to achieve our desired ex-post fairness guarantees of Prop1 and EF_1^1 .

Theorem 4. *There is a strongly polynomial-time algorithm that, given any fractional proportional (Prop) allocation as input, computes an implementation of it using integral allocations that are proportional up to one good (Prop1). If, in addition, the input is a fractional MNW allocation, then the integral allocations in the support also satisfy envy-freeness up to one good more-and-less (EF_1^1).*

Proof. Proof. Let X be a fractional allocation, and let A^1, \dots, A^ℓ be integral allocations in the support of an implementation of X produced by Lemma 3.

Suppose X satisfies proportionality. We want to show that for each $k \in [\ell]$, A^k is Prop1. Since X is proportional, for every $i \in N$, $v_i(X_i) \geq v_i(\bar{I}^m)/n$, where $v_i(\bar{I}^m)$ is agent i 's utility for receiving all goods fully. Fix $k \in [\ell]$. By Lemma 3, we have that for every agent $i \in N$, either $v_i(A_i^k) \geq v_i(X_i) \geq v_i(\bar{I}^m)/n$, or there exists a good $g \notin A_i^k$ such that $v_i(A_i^k) + v_i(g) > v_i(X_i) \geq v_i(\bar{I}^m)/n$. Therefore, A^k is Prop1.

Next, suppose that X maximizes the Nash social welfare among all fractional allocations. Since a fractional MNW allocation is certainly proportional [47], the aforementioned argument still applies for ex-post Prop1. We show that in this case, A^k is also EF_1^1 for each $k \in [\ell]$. Note that since X is a fractional MNW allocation, the following condition is satisfied for any pair of agents $i, j \in N$ and any good $g \in M$ (the condition that transferring an arbitrarily small fraction of good g from agent i to agent j does not increase Nash welfare reduces to this condition):

$$X_{i,g} > 0 \implies \frac{v_i(g)}{v_i(X_i)} \geq \frac{v_j(g)}{v_j(X_j)}. \quad (1)$$

Fix a pair of distinct agents $i, j \in N$. By Lemma 3, either $v_i(A_i^k) \leq v_i(X_i)$, or there exists $g_i^+ \in A_i^k$ with $X_{i, g_i^+} < 1$ such that $v_i(A_i^k \setminus \{g_i^+\}) < v_i(X_i)$. Similarly, either $v_j(A_j^k) \geq v_j(X_j)$, or there exists $g_j^- \notin A_j^k$ with $X_{j, g_j^-} > 0$ such that $v_j(A_j^k \cup \{g_j^-\}) > v_j(X_j)$. To simplify the analysis, let us assume that the second condition holds in both cases. (If $v_i(A_i^k) \leq v_i(X_i)$ (or $v_j(A_j^k) \geq v_j(X_j)$), we can treat g_i^+ (or g_j^-) as a dummy good with $v_i(g_i^+) = 0$ (or $v_j(g_j^-) = 0$).

By summing the right-hand side inequality in Equation (1) over all $g \in A_i^k \setminus \{g_i^+\}$, we get

$$\frac{v_j(A_i^k \setminus \{g_i^+\})}{v_j(X_j)} \leq \frac{v_i(A_i^k \setminus \{g_i^+\})}{v_i(X_i)} < 1.$$

Thus, $v_j(A_i^k \setminus \{g_i^+\}) < v_j(X_j) < v_j(A_j^k \cup \{g_j^-\})$, implying that A^k satisfies EF_1^1 , as desired. \square

Notice that the proof of Theorem 4 establishes a stronger version of Prop1 wherein an agent not receiving her proportional share gets *strictly* more than her proportional share by receiving one additional good. Similarly, it also establishes a stronger version of EF_1^1 wherein an agent envying another agent would strictly prefer her own allocation over the other agent's allocation after adding one missing good to her bundle and removing one good from the other agent's bundle.

Barman and Krishnamurthy [9] recently established that integral Prop1+ EF_1^1 +fPO allocations exist and can be computed in strongly polynomial time. They rely on special-purpose techniques for rounding a fractional MNW allocation. By contrast, Theorem 4 uses a standard technique to round a fractional MNW allocation, computes not just *one* integral Prop1+ EF_1^1 +fPO allocation but rather an implementation of the fractional MNW allocation over such integral allocations, and can be applied to *any* fractional Prop+PO allocation to implement it using integral Prop1+fPO allocations. (Recall that since a fractional MNW allocation is Pareto optimal, any allocation in the support of an implementation of it must be fPO by Proposition 1.)

The reason we use a fractional MNW allocation in Theorem 4 is because it is known to be group fair (GF). This observation, the fact that a fractional MNW allocation can be computed in strongly polynomial time [41, 48], and Theorem 4 immediately yield the main result of this section.

Corollary 1. *There exists a randomized allocation that is ex-ante group fair (GF), ex-post proportional up to one good (Prop1), and ex-post envy-free up to one good more-and-less (EF_1^1). Further, it can be computed in strongly polynomial time.*

6 Characterizing MNW

To prove Corollary 1, we utilized a rounding of a fractional MNW allocation because it is known to be group fair. One may wonder whether it is possible to use any other fractional allocation as the basis for this result. In this section, we answer this negatively by showing a fundamental connection between group fairness and the MNW allocation rule. In particular, we show that MNW, or refinements thereof, is the only allocation rule that satisfies group fairness together with *replication invariance*, which states that the set of allocations returned by an allocation rule should not substantively change if the input is replicated.

6.1 Replication Invariance

Replication invariance has also been well-explored in economic literature on resource allocation [29, 45, 46, 47]. Intuitively, it says that if we take an instance I and duplicate it to form k copies, the output on the duplicated instance contains the duplicate of each possible allocation on the original instance. Recall that an allocation rule outputs a set of (tied) allocations given an instance of the fair division problem. Formally, given an instance $I = (N, M, (v_i)_{i \in N})$ and a positive integer $k \in \mathbb{N}$, let kI denote an instance in which each agent and good is replicated

k times. For an allocation A on I , let kA denote the allocation on kI in which every copy of agent $i \in N$ receives allocation A_i .

Definition 11 (Replication Invariance). An allocation rule f satisfies *replication invariance* if, for every instance $I \in \mathcal{I}$, and every allocation $A \in f(I)$, $kA \in f(kI)$.

We remarked in Section 1.2 that replication invariance works quite well with strictly convex preferences, which is the domain in which many of the early characterization results of the MNW rule hold. As just one example, note that under strictly convex utilities, any fractional envy-free and Pareto optimal allocation must produce identical allocations to two agents with identical valuations, as observed by Varian [47]. To see this, suppose for contradiction that agents i and j have identical valuation v but receive different allocations A_i and A_j . Then, envy-freeness implies that $v(A_i) = v(A_j)$. Now, by strong convexity, $v((A_i + A_j)/2) > v(A_i) = v(A_j)$. Hence, allocation A violates fractional Pareto optimality because giving each of agents i and j the allocation $(A_i + A_j)/2$ would make them both strictly happier. This property comes in handy when dealing with replication because we know that all replicas of an agent must receive identical allocations. In contrast, this property does not hold for additive utilities. For example, if two agents equally like two goods, giving one good fully to one agent and the other good fully to the other agent violates this property and is yet envy-free and Pareto optimal.

6.2 Group Fairness and Replication Invariance Characterize MNW

We have remarked that MNW satisfies group fairness. We begin this section by showing that the fractional MNW rule (which returns the set of all fractional MNW allocations) also satisfies replication invariance.

Proposition 5. *The fractional MNW rule satisfies replication invariance.*

The proof of Proposition 5 is presented in Section 11.1.

We next show that MNW is in fact characterized by group fairness and replication invariance, in the sense that any allocation rule that sometimes produces an allocation that does not maximize the product of valuations necessarily violates either group fairness or replication invariance.

Theorem 5. *Any allocation rule that satisfies group fairness and replication invariance is a refinement of the fractional MNW rule.*

Before proving Theorem 5, we show a technical lemma that will be helpful.

Lemma 4. *For every instance I , and every Pareto optimal allocation $A \notin \text{MNW}(I)$, there exists a pair of agents i, j and $X \subseteq A_i$ with $v_i(X) \in (0, v_i(A_i))$ and $v_j(X) > 0$, such that*

$$v_j(X) > v_j(A_j) \cdot \frac{v_i(X)}{v_i(A_i \setminus X)}. \quad (2)$$

The proof of Lemma 4 is presented in Section 11.2.

Proof. Proof of Theorem 5. Let f be an allocation rule, and suppose that there exists an instance I for which there exists an allocation $A \in f(I)$ but $A \notin \text{MNW}(I)$. If A is not Pareto optimal then it violates group fairness. So suppose A is Pareto optimal. We will show that if it satisfies replication invariance, then it must violate group fairness.

By Lemma 4, there exist agents i, j and $X \subseteq A_i$ with $v_i(X) \in (0, v_i(A_i))$ and $v_j(X) > 0$ such that

$$v_j(X) > v_j(A_j) \cdot \frac{v_i(X)}{v_i(A_i \setminus X)}.$$

Further, suppose that $\frac{v_i(A_i \setminus X)}{v_i(X)}$ is an integer. This is WLOG because we can just take the set of goods X and scale them down until we get an integer. Equation (2) is preserved under this scaling due to additive valuations.

Let $k = \frac{v_i(A_i \setminus X)}{v_i(X)} \in \mathbb{N}$ and consider the replicated instance kI . Denote the copies of agent i in instance kI by i_1, \dots, i_k . By replication invariance, $kA \in f(kI)$. We show that kA does not satisfy group fairness. Let S consist of all k copies of i and a single copy of j , and T consist of all k copies of i . For every $\ell \in [k]$, let $B_{i_\ell} = A_{i_\ell} \setminus X_\ell$, where X_ℓ is a copy of X , and let $B_j = X_1 \cup \dots \cup X_k$. For every copy i_ℓ of agent i , we have

$$\frac{|S|}{|T|} \cdot v_{i_\ell}(B_{i_\ell}) = \frac{|S|}{|T|} \cdot v_{i_\ell}(A_{i_\ell} \setminus X_k) = \frac{\frac{v_i(A_i \setminus X)}{v_i(X)} + 1}{\frac{v_i(A_i \setminus X)}{v_i(X)}} v_i(A_i \setminus X) = v_i(A_i \setminus X) + v_i(X) = v_i(A_i)$$

and for the single copy of j , we have

$$\begin{aligned} \frac{|S|}{|T|} \cdot v_j(B_j) &= \frac{|S|}{|T|} \cdot k \cdot v_j(X) = \frac{\frac{v_i(A_i \setminus X)}{v_i(X)} + 1}{\frac{v_i(A_i \setminus X)}{v_i(X)}} \cdot \frac{v_i(A_i \setminus X)}{v_i(X)} \cdot v_j(X) \\ &= \frac{v_i(A_i)}{v_i(X)} \cdot v_j(X) > \frac{v_i(A_i \setminus X)}{v_i(X)} v_j(X) > v_j(A_j), \end{aligned}$$

where the final inequality follows from Equation (2). Group fairness is violated. \square

We explore this characterization further in the appendix. In particular, we investigate the implications of relaxing or strengthening group fairness. We show that Theorem 5 ceases to hold even under a relatively mild relaxation of group fairness. We also show that if group fairness is strengthened to *fractional group fairness*, which extends group fairness to allow groups that consist of fractions of agents, then Theorem 5 holds even without requiring replication invariance.

7 The Case of Bads

In this section, we consider the case where we are dividing *bads* instead of goods. We show that our positive results from Sections 3 and 5 still apply in this setting.

Formally, let M be a set of m bads. Fractional, integral, and randomized allocations are defined as in the case of goods. However, the agents now have non-positive valuation functions, i.e., $v_{i,r} \leq 0$ for all agents $i \in N$ and bads $r \in M$. The definitions of proportionality, envy-freeness, (fractional) Pareto optimality, and group fairness are unchanged from the goods case. However, the definitions of approximate fairness need to be modified. Crucially, when assessing whether an allocation is approximately fair from the perspective of agent i , we often *remove a bad* from (rather than add a good to) i 's bundle and/or *add a bad* to (rather than remove a good from) some other agent j 's bundle. This causes the relaxations to differ from their goods counterparts significantly.

Definition 12 (Proportionality Up To One Bad (Prop1)). An integral allocation A is *proportional up to one bad* if for every agent $i \in N$, either $v_i(A_i) \geq v_i(\bar{I}^m)/n$ or there exists a bad $r \in A_i$ such that $v_i(A_i \setminus \{r\}) \geq v_i(\bar{I}^m)/n$, where $v_i(\bar{I}^m)$ is agent i 's valuation for receiving all bads fully.

Definition 13 (Envy-Freeness Up To k Bads (EF k)). An integral allocation A is *envy-free up to k bads* if for every pair of agents $i, j \in N$, there exists $S_i \subseteq A_i$ with $|S_i| \leq k$ such that $v_i(A_i \setminus S_i) \geq v_i(A_j)$.

Definition 14 (Envy-Freeness Up To One Bad More-and-Less (EF $_1^1$)). An integral allocation A is *envy-free up to one bad more-and-less* if for every pair of agents $i, j \in N$ such that $A_i \neq \emptyset$, we have $v_i(A_i \setminus \{r_i\}) \geq v_i(A_j \cup \{r_j\})$ for some bads $r_i \in A_i$ and $r_j \notin A_j$.

Finally, for the case of bads, there is no known equivalent of the MNW allocation rule. For example, it is known that maximizing or minimizing the product of valuations (or the product of absolute valuations) leads to dramatically unfair outcomes. However, fractional CEEI

allocations are still known to exist; we refer an interested reader to the work of Bogomolnaia et al. [17]. It is easy to show, as in the case of goods, that these allocations must be group fair (GF). We do not need to introduce the definition of CEEI for bads; we only need the following property, observed by Bogomolnaia et al. [17, Lemma 6]. If X is a fractional CEEI allocation of bads, then $v_i(X_i) < 0$ for all agents $i \in N$ (this assumes that each agent has a strictly negative value for at least one bad; if an agent has zero value for all bads, we can effectively remove her from the instance). Further, for all agents $i, j \in N$ and bads $r \in M$,

$$X_{i,r} > 0 \implies \frac{v_{i,r}}{v_i(X_i)} \leq \frac{v_{j,r}}{v_j(X_j)}. \quad (3)$$

Note that the quantities on both sides of the inequality are non-negative. This property is similar to that of CEEI allocations for goods that we used in the proof of Theorem 8.

7.1 Ex-Ante EF + Ex-Post EF1 for Bads

Note that our algorithms from Section 3 — recursive probabilistic serial (RPS) and its polynomial-time variant — only use ordinal preferences of agents over goods. Hence, these algorithms naturally extend to the case of bads: Given an agent’s valuation function over bads, we can construct an ordinal preference over bads in which the bads are sorted in a non-increasing order of the agent’s valuation for them.

One might wonder how the guarantees that our algorithms provide for goods translate to the case of bads. While the ex-ante EF guarantee carries over, unfortunately the ex-post EF1 guarantee enjoyed by our algorithms in the case of goods does not directly translate to the case of bads. To see this, consider an instance with two agents 1, 2 and three bads b_1, b_2, b_3 such that $v_i(b_r) = -r$ for all $i \in [2]$ and $r \in [3]$. That is, all agents strictly prefer bad b_1 to b_2 and b_2 to b_3 . In the first step, the agents start eating b_1 simultaneously, eat half of b_1 each, then move onto b_2 , and eat half of b_2 each. Consider an integral allocation in which b_1 is allocated to agent 1 and b_2 is allocated to agent 2. In the second step, each agent eats half of b_3 . At this point, our algorithm will, with probability 1/2, allocate b_3 to agent 2. This will result in an integral allocation A in the support where $A_1 = \{b_1\}$, and $A_2 = \{b_2, b_3\}$. It is easy to check that agent 2 continues to envy agent 1 even if we remove any one bad from agent 2’s bundle. Removing two bads from agent 2’s bundle removes the envy, however. Hence, this allocation is EF2, but not EF1. In fact, we show that both RPS and its polynomial-time variant are always ex-post EF2 for bads.

Proposition 6. *For additive valuation functions over bads, recursive probabilistic serial (RPS) and its polynomial-time variant from Section 3.2 produce randomized allocations that are ex-ante envy-free (EF) and ex-post envy-free up to two bads (EF2).*

The proof of Proposition 6 is presented in Section 12.1.

Once again, note that because these algorithms only take the ordinal preferences as input, they in fact guarantee ex-ante SD-EF + ex-post SD-EF2 (which can be defined similarly to SD-EF1 using the SD preference relation).

It is natural to ask whether this is the best we can do. Note that in the proof of Proposition 6, an “up to two bads” relaxation was required because the envying agent may receive an extra bad compared to the envied agent ($\lceil m/n \rceil$ bads as opposed to $\lfloor m/n \rfloor$ bads). However, when m is a multiple of n , it turns out that the allocation is in fact ex-post EF1, as the following lemma shows.

Lemma 5. *For additive valuation functions over bads, recursive probabilistic serial and its polynomial-time variant from Section 3.2 produce randomized allocations that are ex-ante envy-free (EF) and ex-post envy-free up to one bad (EF1) whenever the number of bads is divisible by the number of agents.*

Proof. Proof. Envy-freeness follows directly from Proposition 6. For EF1, note that the exact same reasoning as in the proof of Proposition 6 applies, and we have $v_i(\{b_{i,t} : 1 \leq t < \lfloor m/n \rfloor\}) \geq$

$v_i(\{b_{j,t} : 2 \leq t \leq \lfloor m/n \rfloor\}) \geq v_i(A_j)$ for every pair of agents $i, j \in N$. However, since n divides m , we are guaranteed that there is only one bad missing from the LHS, which is that allocated to agent i in round $\lfloor m/n \rfloor = m/n$. Hence, we have $v_i(A_i \setminus b_{i,m/n}) \geq v_i(A_j)$. \square

What about the case where m is not a multiple of n ? It turns out that a simple modification of RPS (or its polynomial-time variant) works. One can simply add a number of dummy bads (which all agents value at 0) such that the total number of bads becomes divisible by the number of agents n , then use RPS or its polynomial-time variant to compute an ex-ante EF + ex-post EF1 allocation (Lemma 5), and finally remove the dummy bads from the computed allocation, which does not affect its fairness guarantees. This yields the following result.

Theorem 6. *For additive valuation functions over bads, a randomized allocation that is ex-ante envy-free (EF) and ex-post envy-free up to one bad (EF1) always exists, and can be computed in strongly polynomial time.*

7.1.1 The case of mixed manna.

Let us briefly examine the *mixed manna* setting in which we need to allocate *items*, and each item can be either a good or a bad for each agent. We generically denote an item by o . That is, for each agent $i \in N$ and item $o \in M$, we only require that $v_{i,o} \in \mathbb{R}$.

Envy-freeness up to one item can be generalized to the mixed manna setting. It requires that if agent i envies agent j , then the envy can be eliminated by either removing one item from agent j 's bundle (which must be a good for i), or one item from agent i 's bundle (which must be a bad for i).

Definition 15 (Envy-Freeness Up To One Item (EF1); [8]). An integral allocation A is *envy-free up to one item* if for every pair of agents $i, j \in N$, either i does not envy j (i.e., $v_i(A_i) \geq v_i(A_j)$), or there exists an item $o \in A_i \cup A_j$ such that $v_i(A_i \setminus \{o\}) \geq v_i(A_j \setminus \{o\})$.

It is easy to check that when the instance consists of only goods or only bads, this definition reduces to the EF1 definitions we previously introduced. However, our algorithms achieve a relaxation of this, which we call *weak envy-freeness up to one item* (w-EF1), that allows *both* the removal of an item from agent j 's bundle and an item from agent i 's bundle. Note that when the instance consists of only goods or only bads, w-EF1 still reduces to the standard EF1 definition.

Definition 16 (Weak Envy-Freeness Up to One Item (w-EF1)). An integral allocation A is *weak envy-free up to one item* if for every pair of agents $i, j \in N$, either i does not envy j (i.e., $v_i(A_i) \geq v_i(A_j)$), or there exists a good $o_j \in A_j$ and/or a bad $o_i \in A_i$ such that $v_i(A_i \setminus \{o_i\}) \geq v_i(A_j \setminus \{o_j\})$.

While we were unable to settle the existence of an ex-ante EF + ex-post EF1 allocation for the mixed manna setting, we are able to guarantee ex-ante EF + ex-post w-EF1. Thus, Theorem 7 generalizes both Theorem 2 and Theorem 6.

Theorem 7. *For additive valuation functions in the mixed manna setting, a randomized allocation that is ex-ante envy-free (EF) and ex-post weak envy-free up to one item (w-EF1) always exists, and can be computed in strongly polynomial time.*

Proof. Proof. Once again, we add dummy items (which all agents value at 0) to guarantee that m is an integer multiple of n , and remove them at the end. When we run RPS (or its polynomial-time variant), we can again apply the same argument as in the proof of Proposition 6 to obtain $v_i(\{o_{i,t} : 1 \leq t < m/n\}) \geq v_i(\{o_{j,t} : 2 \leq t \leq m/n\})$. That is, $v_i(A_i \setminus \{o_{i,m/n}\}) \geq v_i(A_j \setminus \{o_{j,1}\})$, which is the desired ex-post w-EF1 guarantee. Ex-ante EF follows with the same reasoning as before. \square

Since our algorithms only take the ordinal preferences of agents over items as input, this in fact guarantees ex-ante SD-EF and ex-post SD-w-EF1 (which can be defined similarly to SD-EF1 using the SD preference relation).

Finally, we note that the techniques from this section can also be applied to the decomposition of probabilistic serial (PS) achieved in the subsequent work of Aziz [7] to obtain an ex-ante SD-EF and ex-post SD-w-EF1 allocation of mixed manna that retains additional axiomatic properties of PS.

7.2 Ex-Ante GF + Ex-Post Prop1 + Ex-Post EF₁¹ for Bads

We now show that our main result from Section 5 for the case of goods carries over easily to the case of bads. This is somewhat striking because it has been observed that many results for the case of goods do not easily carry over to the case of bads [17]. For example, while it is known that a competitive equilibrium with equal incomes (CEEI) fractional allocation exists for the case of bads, and such an allocation is envy-free and Pareto optimal, it is no longer obtained by maximizing or minimizing the product of (dis)utilities, and whether such an allocation can be computed in polynomial time is a major open question [20, 34]. However, if such an allocation is given, Brânzei and Sandomirskiy [20] show that one can round it to obtain an integral Prop1+EF₁¹+fPO allocation in strongly polynomial time.

Interestingly, we do not need to re-prove the utility guarantee (Lemma 3) for the case of bads. Instead, we can simply reduce the case of bads to the case of goods by considering a modified valuation function $\hat{v}_i = -v_i$ (thus guaranteeing $\hat{v}_i \geq 0$), and since Lemma 3 bounds the difference between the utility of an agent in an integral allocation and her utility in the original fractional allocation from both below and above, this guarantee remains useful for the case of bads. We then show that this guarantee is sufficient to establish Prop1 and EF₁¹ in the case of bads as well.

Theorem 8. *For additive valuation functions over bads, there is a strongly polynomial-time algorithm that, given any fractional proportional (Prop) allocation as input, computes an implementation of it using integral allocations that are proportional up to one bad (Prop1). If, in addition, the input is a fractional CEEI allocation, then the integral allocations in the support also satisfy envy-freeness up to one bad more-and-less (EF₁¹).*

The proof of Theorem 8 is presented in Section 12.2.

Corollary 2. *For additive valuation functions over bads, there exists a randomized allocation that is ex-ante group fair (GF), ex-post proportional up to one bad (Prop1), and ex-post envy-free up to one bad more-and-less (EF₁¹). Further, it can be computed in strongly polynomial time given a fractional CEEI allocation.*

We note that in Corollary 2, unlike in the case of goods, we have to add the condition “given a fractional CEEI allocation”. This is because a fractional CEEI allocation is known to be computable in strongly polynomial time in the case of goods, but this remains an interesting open question in the case of bads [20, 34].

8 Discussion

Perhaps the most fascinating open question that stems from our work is whether ex-ante envy-freeness (or even ex-ante proportionality) is compatible with ex-post EF1 and ex-post PO. While our simulations did not find a counterexample, simulations have been known to be misleading for existential results in the past, e.g., in case of the (non)existence of MMS allocations; see [38].

Open Question: *Does there always exist a randomized allocation that is ex-ante EF, ex-post EF1, and ex-post PO? What about ex-ante Prop, ex-post EF1, and ex-post PO?*

The difficulty in approaching this question is that there are very few available methods of finding integral EF1+PO allocations [10, 24], so finding many such allocations and randomizing

over them is tricky. Also, unlike the set of integral EF1 allocations, which we somewhat understand, not much is known about the set of integral EF1+PO allocations other than the fact that it is always non-empty. For example, round-robin method allows us to control exactly which agents might envy which other agents in addition to obtaining EF1. The following questions, which thus resolve positively for EF1, are open for EF1+PO to the best of our knowledge.

Open Question: *Given agents i and j , does there always exist an integral EF1+PO allocation in which agent i does not envy agent j ? What about an integral EF1+PO allocation in which agent i does not envy any other agent? Given a priority ordering over the agents, does there always exist an EF1+PO allocation in which no agent envies an agent with lower priority?*

Various other open problems remain. For instance, other related notions such as envy-freeness up to any good (EFX) and approximate maximin-share guarantee (MMS) are natural candidates for ex-post fairness. (For additive valuations, whether an integral EFX allocation always exists is an open problem, and while integral MMS allocations do not always exist [38], 3/4-approximate MMS allocations are known to always exist [33, 35].)

Future work could also consider the price of fairness: what fraction of the optimal social welfare must be sacrificed (in the worst case) in order to guarantee a fair allocation? Bertsimas et al. [14] and Caragiannis et al. [25] study the price of EF, while Bei et al. [12] and Barman et al. [11] study the price of EF1 for indivisible goods. What is the price of achieving both ex-ante EF and ex-post EF1 together?

More broadly, the next step would be to achieve ex-ante and ex-post fairness guarantees simultaneously in a variety of other problems such as voting, matching, and public decision-making.

Acknowledgments

We are grateful to anonymous reviewers from EC'20 and to Dominik Peters for helpful comments. Rohit Vaish acknowledges support from ONR #N00014-17-1-2621. Nisarg Shah was partially supported by an NSERC Discovery Grant.

References

- [1] A. Abdulkadiroğlu and T. Sönmez. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica*, 66(3):689–701, 1998.
- [2] M. Akbarpour and A. Nikzad. Approximate Random Allocation Mechanisms. *The Review of Economic Studies*, 2020. ISSN 0034-6527. doi: 10.1093/restud/rdz066.
- [3] M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online Fair Division: Analysing a Food Bank Problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2540–2546, 2015.
- [4] K. J. Arrow and M.D. Intriligator, editors. *Handbook of Mathematical Economics*. North-Holland, 1982.
- [5] H. Aziz. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. In *The Future of Economic Design*, pages 45–50. Springer, 2019.
- [6] H. Aziz and C. Ye. Cake Cutting Algorithms for Piecewise Constant and Piecewise Uniform Valuations. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, pages 1–14, 2014.
- [7] Haris Aziz. Simultaneously Achieving Ex-Ante and Ex-Post Fairness. *arXiv preprint arXiv:2004.02554*, 2020.

- [8] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair Allocation of Indivisible Goods and Chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2019.
- [9] S. Barman and S.K. Krishnamurthy. On the Proximity of Markets with Integral Equilibria. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1748–1755, 2019.
- [10] S. Barman, S. K. Krishnamurthy, and R. Vaish. Finding Fair and Efficient Allocations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.
- [11] S. Barman, U. Bhaskar, and N. Shah. Settling the Price of Fairness for Indivisible Goods. *arXiv preprint arXiv:2007.06242*, 2020.
- [12] X. Bei, X. Lu, P. Manurangsi, and W. Suksompong. The Price of Fairness for Indivisible Goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence IJCAI*, pages 81–87, 7 2019.
- [13] M. Berliant, W. Thomson, and K. Dunz. On the Fair Division of a Heterogeneous Commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- [14] D. Bertsimas, V. F. Farias, and N. Trichakis. The Price of Fairness. *Operations Research*, 59(1):17–31, 2011.
- [15] G. Birkhoff. Three Observations on Linear Algebra. *Universidad Nacional de Tucumán, Revista A*, 5:147–151, 1946.
- [16] A. Bogomolnaia and H. Moulin. A New Solution to the Random Assignment Problem. *Journal of Economic Theory*, 100:295–328, 2001.
- [17] A. Bogomolnaia, H. Moulin, F. Sandomirskiy, and E. Yanovskaya. Competitive Division of a Mixed Manna. *Econometrica*, pages 1847–1871, 2017.
- [18] S. Bouveret, Y. Chevaleyre, N. Maudet, and H. Moulin. *Fair Allocation of Indivisible Goods*, pages 284–310. Cambridge University Press, 2016.
- [19] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [20] S. Brânzei and F. Sandomirskiy. Algorithms for Competitive Division of Chores. *arXiv preprint arXiv:1907.01766*, 2019.
- [21] J. Brustle, J. Dippel, VV Narayan, M Suzuki, and A Vetta. One Dollar Each Eliminates Envy. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 23–39, 2020.
- [22] E. Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [23] E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review*, 103(2):585–623, 2013.
- [24] I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, and J. Wang. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

- [25] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.
- [26] Y. Chen and T. Sönmez. Improving Efficiency of On-Campus Housing: An Experimental Study. *American Economic Review*, 92(5):1669–1686, 2002.
- [27] V. Conitzer, R. Freeman, and N. Shah. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.
- [28] V. Conitzer, R. Freeman, N. Shah, and J.W. Vaughan. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1853–1860, 2019.
- [29] G. Debreu and H. Scarf. A Limit Theorem on the Core of an Economy. *International Economic Review*, 4(3):235–246, 1963.
- [30] Edmund Eisenberg and David Gale. Consensus of Subjective Probabilities: The Parimutuel Method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [31] D. Foley. Resource Allocation and the Public Sector. *Yale Economics Essays*, 7:45–98, 1967.
- [32] T. Gajdos and J. Tallon. Fairness under Uncertainty. *Economics Bulletin*, 4(18):1–7, 2002.
- [33] J. Garg and S. Taki. An Improved Approximation Algorithm for Maximin Shares. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, 2020. Forthcoming.
- [34] Jugal Garg and Peter McGlaughlin. Computing Competitive Equilibria with Mixed Manna. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 420–428, 2020.
- [35] M. Ghodsi, M. HajiAghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair Allocation of Indivisible Goods: Improvements and Generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 539–556, 2018.
- [36] M. Grötschel, L. Lovász, and A. Schrijver. The Ellipsoid Method and its Consequences in Combinatorial Optimization. *Combinatorica*, 1(2):169–197, 1981.
- [37] A. Hylland and R. Zeckhauser. The Efficient Allocation of Individuals to Positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
- [38] D. Kurokawa, A.D. Procaccia, and J. Wang. Fair Enough: Guaranteeing Approximate Maximin Shares. *Journal of the ACM (JACM)*, 65(2):1–27, 2018.
- [39] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.
- [40] A.S. Nesterov. Fairness and Efficiency in Strategy-Proof Object Allocation Mechanisms. *Journal of Economic Theory*, 170:145–168, 2017.
- [41] J.B. Orlin. Improved Algorithms for Computing Fisher’s Market Clearing Prices. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 291–300, 2010.
- [42] H. Scarf. An Analysis of Markets with a Large Number of Participants. In *Recent Advances in Game Theory*, pages 127–155, 1962.

- [43] E. Segal-Halevi and B.R. Sziklai. Monotonicity and Competitive Equilibrium in Cake-Cutting. *Economic Theory*, 68(2):363–401, 2019.
- [44] H. Steinhaus. The Problem of Fair Division. *Econometrica*, 16:101–104, 1948.
- [45] W. Thomson. A Study of Choice Correspondences in Economies with a Variable Number of Agents. *Journal of Economic Theory*, 46(2):237–254, 1988.
- [46] William Thomson. Fair Allocation Rules. In *Handbook of Social Choice and Welfare*, volume 2, pages 393–506. Elsevier, 2011.
- [47] H. Varian. Equity, Envy and Efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- [48] L.A. Végh. A Strongly Polynomial Algorithm for a Class of Minimum-Cost Flow Problems with Separable Convex Objectives. *SIAM Journal on Computing*, 45(5):1729–1761, 2016.
- [49] J. von Neumann. A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem. In W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 2, pages 5–12. Princeton University Press, 1953.

9 Omitted Material from Section 3

9.1 Recursive Probabilistic Serial

Algorithm 2 EATING

- 1: **Input:** A (possibly fractional) set of goods M' .
 - 2: **Output:** A fractional allocation X , integral allocations $\{A^k\}_{k \in [\ell]}$, and coefficients $\{w^k\}_{k \in [\ell]}$
 - 3: $X \leftarrow$ fractional (partial) allocation of M' by running Probabilistic Serial for one unit of time (one unit is the time it takes any agent to consume a single unit of any good).
 - 4: **return** X
-

Algorithm 3 A Possibly Exponential-Time Implementation of Recursive Probabilistic Serial (RPS)

- 1: **Input:** A (possibly fractional) set of goods M^t .
 - 2: **Output:** A fractional allocation Y^t .
 - 3: **if** $M^t = \emptyset$ **then**
 - 4: $Y^t \leftarrow \emptyset$
 - 5: **else**
 - 6: $X^t \leftarrow \text{EATING}(M^t, 1)$
 - 7: BvN: $X^t = \sum_{k=1}^{\ell} w^{t,k} A^{t,k}$
 - 8: $Y^t \leftarrow X^t + \sum_{k=1}^{\ell} w^{t,k} \cdot \text{RPS}(M^t \setminus A^{t,k})$
 - 9: **end if**
 - 10: **return** Y^t
-

Algorithm 4 A Polynomial-Time Implementation of Recursive Probabilistic Serial (RPS)

```

1: Input: A (possibly fractional) set of goods  $M^t$ .
2: Output: A fractional allocation  $Y^t$ .
3: if  $M^t = \emptyset$  then
4:    $Y^t \leftarrow \emptyset$ 
5: else
6:    $X^t \leftarrow \text{EATING}(M^t, 1)$ 
7:   BvN:  $X^t = \sum_k \bar{w}^{t,k} \bar{A}^{t,k}$ 
8:   Proposition 3:  $X^t = \sum_{k=1}^{s_t} w^{t,k} A^{t,k}$  where  $s_t \leq mn + 1$  is the support-size after compression.
9:    $q \leftarrow 1$ 
10:   $M^{t+q} \leftarrow M^{t+q-1} \setminus \cup_{i \in N} X_i^{t+q-1}$   $\triangleright$  Residual set of goods after the  $q^{\text{th}}$  round of eating
11:  while  $M^{t+q} \neq \emptyset$  do
12:     $\setminus$ * The fractional allocation at the end of round  $t + q - 1$  is  $X^{t+q-1} = \sum_{k=1}^{s_{t+q-1}} w^{t+q-1,k} A^{t+q-1,k}$  and has support-size  $s_{t+q-1} \leq mn + 1$ . * $\setminus$ 
13:    for  $k = 1, 2, \dots, s_{t+q-1}$  do
14:       $Z^{t+q,k} \leftarrow \text{EATING}(M^{t+q} \setminus A^{t+q-1,k}, 1)$ 
15:      BvN:  $Z^{t+q,k}$  has polynomial-sized support.
16:    end for
17:     $X^{t+q} \leftarrow \sum_{k=1}^{s_{t+q-1}} w^{t+q-1,k} \cdot Z^{t+q,k}$ 
18:    Proposition 3: Compress the support-size of  $X^{t+q}$  to at most  $mn + 1$ .
19:     $q \leftarrow q + 1$ 
20:  end while
21:  return  $Y^t \leftarrow X^{t+q-1}$ 
22: end if

```

10 Omitted Material from Section 5

10.1 MNW satisfies Group Fairness

For completeness, we provide a proof of the result mentioned by Conitzer et al. [28], that MNW satisfies group fairness. Indeed, we prove a stronger result. We show that this holds not only for the case of goods division, but in the more general cake-cutting setting.

In the cake-cutting setting, we have a single heterogeneous good modeled by $M = [0, 1]$, and each agent has an integrable valuation function $v_i : [0, 1] \rightarrow \mathbb{R}_+$. Each agent's allocation, A_i , is a finite union of intervals, and $v_i(A_i)$ is given by the area under the curve v_i over the intervals in A_i .

A price function p is a measure on M . Say that a subset $Z \subseteq M$ is a *positive slice* if there exists an agent i with $v_i(Z) > 0$. A pair (X, p) of allocation X and price p is a *strong competitive equilibrium from equal incomes* (s-CEEI) [43] if the following conditions hold:

1. For every positive slice $Z \subseteq M$, $p(Z) > 0$.
2. For every pair of positive slices $Z \subseteq M$ and $Z_i \subseteq X_i$, $\frac{v_i(Z_i)}{p(Z_i)} \geq \frac{v_i(Z)}{p(Z)}$.
3. For every $i \in N$, $p(X_i) = 1$

Segal-Halevi and Sziklai [43] show that the MNW allocation is an s-CEEI allocation when p is the standard price measure defined by, for all $i \in N$ and all $Z_i \subseteq X_i$, $p(Z_i) = \frac{v_i(Z_i)}{v_i(X_i)}$.

Proposition 7. *MNW satisfies group fairness in the cake-cutting setting.*

Proof. Proof. Consider an instance of cake cutting and suppose that the MNW/s-CEEI allocation X violates group fairness. Then there exist subsets of agents $S, T \subseteq N$ with $|T| \neq 0$ and an allocation Y of $\cup_{j \in T} X_j$ among agents in S such that $\frac{|S|}{|T|} \cdot v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in S$ and the inequality is strict for at least one $i \in S$.

By the s-CEEI condition (2), we know that $v_i(X_i) = \frac{v_i(X_i)}{p(X_i)} \geq \frac{v_i(Y_i)}{p(Y_i)}$ for all $i \in S$. Combining this with the group fairness condition yields

$$\frac{|S|}{|T|} \cdot v_i(Y_i) \geq \frac{v_i(Y_i)}{p(Y_i)} \implies p(Y_i) \geq \frac{|T|}{|S|}$$

for all $i \in S$, with at least one inequality strict. Taking the sum over all $i \in S$ gives us $\sum_{i \in S} p(Y_i) > |T|$. Finally, because $\sum_{i \in S} p(Y_i) = p(\cup_{i \in S} Y_i) = p(\cup_{j \in T} X_j) = \sum_{j \in T} p(X_j)$, it must be the case that $p(X_j) > 1$ for some $j \in T$, violating the assumption that (X, p) is s-CEEI. \square

10.2 Decomposition result of Budish et al. [23]

Let X be a fractional allocation. Recall that X satisfies *column-wise* feasibility constraints, namely $0 \leq \sum_{i \in N} X_{i,r} \leq 1$ for all $r \in M$. More generally, we can impose capacity constraints of the form $\underline{q}_S \leq \sum_{(i,r) \in S} X_{i,r} \leq \bar{q}_S$, where S is a *constraint set* comprising of a collection of agent-object pairs, and \underline{q}_S and \bar{q}_S are the *lower* and *upper quotas* for S , respectively. The set of all capacity constraints imposed by a given problem is called the *constraint structure* \mathcal{H} of the problem, and is specified as a collection of all constraint sets and the corresponding quotas $(\underline{q}_S, \bar{q}_S)_{S \in \mathcal{H}}$. Given a constraint structure \mathcal{H} , we say that the fractional allocation X admits a *feasible* implementation $\mathbf{X} := \{(p^k, A^k)\}_{k \in [\ell]}$ if every integral allocation in its support also satisfies the constraints in \mathcal{H} . That is, for every $k \in [\ell]$, we have

$$\underline{q}_S \leq \sum_{(i,r) \in S} A_{i,r}^k \leq \bar{q}_S \text{ for every } S \in \mathcal{H}.$$

Definition 17 (Hierarchy and bihierarchy). A constraint structure \mathcal{H} is said to be a *hierarchy* (or a laminar family) if for every $S, S' \in \mathcal{H}$, we have that either $S \subset S'$, or $S' \subset S$, or $S \cap S' = \emptyset$. We say that \mathcal{H} is a *bihierarchy* if it can be partitioned into two hierarchies, i.e., if there exist hierarchies \mathcal{H}_1 and \mathcal{H}_2 such that $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ and $\mathcal{H}_1 \cap \mathcal{H}_2 = \emptyset$.

As an example, consider the fractional allocation X in Figure 2. The row constraints (shown as red or blue solid rectangles) as well as all singleton constraints of the form $0 \leq X_{i,r} \leq 1$ (not shown in the figure) together constitute a hierarchy, say \mathcal{H}_1 , since for any pair of constraint sets, either they are disjoint or one is completely contained inside the other. Similarly, the column constraints (shown as gray dotted rectangles) form another hierarchy \mathcal{H}_2 . Furthermore, $\mathcal{H} := \mathcal{H}_1 \cup \mathcal{H}_2$ is a bihierarchy since any constraint set (rectangle or singleton) belongs to exactly one of \mathcal{H}_1 or \mathcal{H}_2 .




Figure 2: Decomposition of a fractional proportional allocation X into integral Prop1 allocations A^1 , A^2 , and A^3 . The underlying fair division instance comprises of four goods and two agents with valuations $v_{1,1} = 10$, $v_{1,2} = 6$, $v_{1,3} = 4$, $v_{1,4} = 2$ and $v_{2,1} = 2$, $v_{2,2} = 10$, $v_{2,3} = 6$, $v_{2,4} = 4$.

Budish et al. [23] showed that bihierarchy constraint structure is a sufficient condition for a fractional allocation to admit a feasible implementation. We recall this result in Proposition 8.

Proposition 8 ([23]). *Given a fractional allocation X satisfying a bihierarchy constraint structure \mathcal{H} , one can compute, in strongly polynomial time, a set of coefficients $p_1, \dots, p_\ell \in [0, 1]$ and integral allocations A^1, \dots, A^ℓ such that (a) $\sum_{k=1}^\ell p^k = 1$, (b) each A^k satisfies the constraints in \mathcal{H} , and (c) $X = \sum_{k=1}^\ell p^k A^k$.*

Observe that the well-known Birkhoff-von Neumann theorem is a special case of Proposition 8 when \mathcal{H}_1 consists of all singleton as well as row constraints, \mathcal{H}_2 consists of all column constraints, and the lower and upper quotas for each constraint set are 0 and 1, respectively. It is worth pointing out that while Budish et al. [23] only note a polynomial running time, it is easy to check that their (combinatorial) algorithm, in fact, runs in *strongly* polynomial time.

Budish et al. [23] use Proposition 8 to establish the utility guarantee (Proposition 4).

10.3 Proof of Lemma 3

Recall the statement of Lemma 3.

Lemma 3 (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, \dots, A^ℓ such that for every $k \in [\ell]$ 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 $v_i(A_i^k) - v_i(g_i^+) < v_i(X_i)$.

Proof. Proof. In their proof of Proposition 4, Budish et al. [23] propose the following method for computing an implementation of a given fractional allocation X . Consider a fixed agent $i \in N$. Suppose the goods in M are indexed as $g_{i,1}, \dots, g_{i,m}$ so that $v_i(g_{i,k}) \geq v_i(g_{i,k+1})$ for each $k \in [m-1]$. For simplicity, we will write $v_{i,k} := v_i(g_{i,k})$ for all $k \in [m]$ and $v_{i,m+1} := 0$.

For any $i \in N$ and any $k \in [m]$, define $Q_{i,k} := \sum_{t=1}^k X_{i,g_{i,t}}$ as the total fractional amount of the k most preferred goods assigned to agent i under X . Consider the following set of constraints

on a generic fractional allocation Y :

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

Observe that X trivially satisfies these constraints. Budish et al. show that these constraints have the so-called ‘‘bihierarchy’’ structure (refer to Section 10.2 for a formal definition), which allows computing, in strongly polynomial time, an implementation of X whose support consists of integral allocations A^1, \dots, A^ℓ that also satisfy these constraints.

Lastly, Budish et al. show that any integral allocation satisfying the constraints in Equation (4) must satisfy the guarantee in Proposition 4. We show that it in fact satisfies the slightly stronger guarantee that we seek. For simplicity, let us write \widehat{A} to denote a generic integral allocation satisfying the constraints in Equation (4), and $\widehat{Q}_{i,k} := \sum_{t=1}^k \widehat{A}_{i,g_{i,t}}$ for all $i \in N$ and $k \in [m]$.

Let us first analyze the case where $v_i(\widehat{A}_i) < v_i(X_i)$. Then, there must exist some good $g \in M$ such that $\widehat{A}_{i,g} < X_{i,g}$. Since $\widehat{A}_{i,g} \in \{0, 1\}$ and $X_{i,g} \in [0, 1]$, this is equivalent to $\widehat{A}_{i,g} = 0 < X_{i,g}$. Let k^- be the smallest index such that $\widehat{A}_{i,g_{i,k^-}} < X_{i,g_{i,k^-}}$, i.e., g_{i,k^-} is agent i 's most preferred good satisfying this condition. Hence, $g_{i,k^-} \notin \widehat{A}_i$ and $X_{i,g_{i,k^-}} > 0$. Further, for all $k < k^-$, we have $\widehat{A}_{i,g_{i,k}} \geq X_{i,g_{i,k}}$, and, as a result, $\widehat{Q}_{i,k} \geq Q_{i,k}$. Thus,

$$\begin{aligned}v_i(X_i) - v_i(\widehat{A}_i) &= \sum_{k=1}^m v_{i,k} \cdot (X_{i,g_{i,k}} - \widehat{A}_{i,g_{i,k}}) = \sum_{k=1}^m (v_{i,k} - v_{i,k+1}) \cdot (Q_{i,k} - \widehat{Q}_{i,k}) \\ &\leq \sum_{k=k^-}^m (v_{i,k} - v_{i,k+1}) \cdot (Q_{i,k} - \widehat{Q}_{i,k}) < \sum_{k=k^-}^m (v_{i,k} - v_{i,k+1}) \cdot 1 = v_{i,k^-},\end{aligned}$$

where the second transition is a simple algebraic exercise, the third transition holds because we noted that $\widehat{Q}_{i,k} \geq Q_{i,k}$ for all $k < k^-$, and the fourth transition holds because \widehat{A} satisfies \mathcal{H}_1 in Equation (4), and therefore, we have that $Q_{i,k} - \widehat{Q}_{i,k} \leq Q_{i,k} - [Q_{i,k}] < 1$. Taking $g_i^- := g_{i,k^-}$, we notice that this is the guarantee we desire when $v_i(\widehat{A}_i) < v_i(X_i)$.

Next, consider the other case where $v_i(\widehat{A}_i) > v_i(X_i)$. Then, there must exist some good $g \in M$ such that $\widehat{A}_{i,g} > X_{i,g}$. Note that this is equivalent to $\widehat{A}_{i,g} = 1 > X_{i,g}$. Let k^+ be the smallest index such that g_{i,k^+} satisfies this condition. Then, we have that $g_{i,k^+} \in \widehat{A}_i$, $X_{i,g_{i,k^+}} < 1$, and by an argument similar to the one above, $v_i(\widehat{A}_i) - v_i(X_i) < v_{i,k^+}$. Hence, in this case, we can take $g_i^+ := g_{i,k^+}$, as desired. \square

11 Omitted Material from Section 6

11.1 Proof of Proposition 5

Recall the statement of Proposition 5.

Proposition 5. *The fractional MNW rule satisfies replication invariance.*

Proof. Proof. Suppose that it violates replication invariance. That is, for some instance $I = (N, M, (v_i)_{i \in N})$ and $k \in \mathbb{N}$, there exists a fractional allocation $Y \in \text{MNW}(I) = \arg \max_{X \in \mathcal{X}} \prod_{i \in N} v_i(X_i)$ such that $kY \notin \text{MNW}(kI)$. For an agent $i \in N$ and $\ell \in [k]$, let i_ℓ denote the ℓ^{th} copy of agent i in the replicated instance kI . Then, there exists an allocation Z on kI such that

$$\prod_{i \in N} \prod_{\ell=1}^k v_i(Z_{i_\ell}) > \prod_{i \in N} \prod_{\ell=1}^k v_i(Y_i).$$

For every $i \in N$ and $\ell \in [k]$, let $Z'_{i_\ell} := Z'_i := \frac{\sum_{\ell'=1}^k Z_{i_{\ell'}}$. That is, Z' takes the total allocation to all copies of i and distributes it evenly across the copies so that each copy of agent i receives Z'_i . Because valuations are additive and each copy of i has the same valuation function, it can be seen that the Nash Welfare of Z' is at least as large as that of Z . That is,

$$\prod_{i \in N} \prod_{\ell=1}^k v_i(Z'_i) \geq \prod_{i \in N} \prod_{\ell=1}^k v_i(Z_{i_\ell}) > \prod_{i \in N} \prod_{\ell=1}^k v_i(Y_i) \Rightarrow \prod_{i \in N} v_i(Z'_i) > \prod_{i \in N} v_i(Y_i),$$

which contradicts the fact that $Y \in \text{MNW}(I)$. \square

11.2 Proof of Lemma 4

Recall the statement of Lemma 4.

Lemma 4. *For every instance I , and every Pareto optimal allocation $A \notin \text{MNW}(I)$, there exists a pair of agents i, j and $X \subseteq A_i$ with $v_i(X) \in (0, v_i(A_i))$ and $v_j(X) > 0$, such that*

$$v_j(X) > v_j(A_j) \cdot \frac{v_i(X)}{v_i(A_i \setminus X)}. \quad (2)$$

Proof. Proof. Let $Z \in \text{MNW}(I)$ and let $U = \{i \in N : v_i(A_i) > v_i(Z_i)\}$ be the set of agents with more utility in allocation A than in Z . First, we claim that there exist agents $i \in U$ and $j \notin U$ with $A_i \cap Z_j \neq \emptyset$. If this were not the case, then $\cup_{i \in U} A_i = \cup_{i \in U} Z_i$. Further, since $A \notin \text{MNW}(I)$, there exists an agent h with $v_h(A_h) \neq v_h(Z_h)$. And since A is Pareto optimal, we must have that $U \neq \emptyset$. Hence, the allocation B with $B_i = A_i$ for every $i \in U$ and $B_i = Z_i$ for every $i \notin U$ has strictly higher Nash welfare than Z , a contradiction.

Fix agents $i \in U$ and $j \notin U$ such that $A_i \cap Z_j \neq \emptyset$. Let $X = \epsilon \cdot (A_i \cap Z_j)$, where $\epsilon > 0$ is small enough that

$$\frac{v_i(A_i \setminus X)}{v_i(Z_i)} > \frac{v_j(A_j)}{v_j(Z_j \setminus X)}, v_i(X) \in (0, v_i(A_i)), \text{ and } v_j(X) \in (0, v_j(Z_j)). \quad (5)$$

It is possible to choose such an ϵ because the first inequality in Equation (5) holds for $\epsilon = 0$ and both sides are continuous in ϵ . Further, we know that $v_i(X) > 0$ by Pareto optimality of A , and that $v_j(X) > 0$ by Pareto optimality of Z .

We also know, by the definition of a fractional MNW allocation, that

$$\frac{v_j(X)}{v_j(Z_j \setminus X)} \geq \frac{v_i(X)}{v_i(Z_i)}. \quad (6)$$

Multiplying the left and right hand sides of the inequalities in Equations (5) and (6) gives the desired result. \square

11.3 Maximum Nash Welfare Variant

Relaxing group fairness. We now present a replication-invariant allocation rule that is not a refinement of MNW. As a consequence of Theorem 5, it therefore fails group fairness. However, we show that it satisfies a relaxation of group fairness that still generalizes many of the fairness properties common in the literature.

Say that good g is a *weak* good if there exists only a single agent with $v_i(g) > 0$. Otherwise, we say that g is a *strong* good. Let $W_i = \{g \in M : v_j(g) = 0 \forall j \neq i\}$ denote the set of all weak goods valued positively by agent i .

Mechanism *Maximum Nash Welfare Variant* (MNW-V) works as follows. Given an instance I , let I_{-W} denote the instance with weak goods removed. Then $\text{MNW-V}(I)_i = \text{MNW}(I_{-W})_i \cup W_i$. That is, we begin by allocating only the strong goods according to MNW. We then allocate to each agent the weak goods that she values positively.

Example 3. Suppose there are three agents and two goods, g and g^* , with $v_1(g) = v_2(g) = v_3(g) = 1$ and $v_1(g^*) = 1$ and $v_2(g^*) = v_3(g^*) = 0$. Note that g is a strong good and g^* is a weak good, valued only by agent 1. MNW allocates agents 1 and 2 half of g each, and all of g^* to agent 3. MNW-V, on the other hand, allocates each agent a $1/3$ share of good g , and all of g^* to agent 3. Note that MNWV does not satisfy group fairness. If we let $S = N$ and T consist only of agent 3, then we can partition agent 3's bundle so that $B_1 = B_2 = \frac{1}{6}g$, and $B_3 = g^*$, for utilities $v_1(B_1) = v_2(B_2) = 1/6$ and $v_3(B_3) = 1$. When these utilities are scaled by a factor of $|S|/|T| = 3$, each agent is strictly happier than under the original MNW-V allocation.

The above example shows that MNW-V violates group fairness. But we can show that the allocation is fair to any group S provided that S is small relative to group T . We formalize this using the following property.

Definition 18. An allocation X is *group fair for less* if for all non-empty subsets of agents $S, T \subseteq N$ with $|S| \leq |T|$, there is no allocation Y of the set of goods $\cup_{i \in T} X_i$ among agents in S such that $\frac{|S|}{|T|} \cdot v_i(Y_i) \geq v_i(X_i)$ for all agents $i \in S$ and at least one inequality is strict.

In particular, group fairness for less generalizes proportionality, envy-freeness, Pareto optimality, group envy-freeness [13, 47], and the core [42].

We show that MNW-V satisfies group fairness for less. As a consequence, a characterization of MNW in the additive utility setting therefore seems to require the full power of group fairness.

Theorem 9. *MNW-V satisfies group fairness for less.*

Proof. Proof. Denote the MNW-V allocation on some instance I by A , and fix sets S and T with $|S| \leq |T|$. Since we know that MNW is group fair, for every partition $(B_i)_{i \in S}$ of $\cup_{i \in T} (A_i \setminus W_i)$, we have either

$$\frac{|S|}{|T|} \cdot v_i(B_i) < v_i(A_i \setminus W_i) \quad (7)$$

for some $i \in S$, or

$$\frac{|S|}{|T|} \cdot v_i(B_i) = v_i(A_i \setminus W_i) \quad (8)$$

for all $i \in S$.

Fix a partition $(C_i)_{i \in S}$ of $\cup_{i \in T} A_i$, and for every $i \in S$ denote by B_i the intersection of C_i with the set of strong goods. $(B_i)_{i \in S}$ is a partition of $\cup_{i \in T} (A_i \setminus W_i)$, so we must be in one of the two cases above. Suppose we are in the first, with witness agent i . If $i \in T$, then $C_i \subseteq B_i \cup W_i$ and we have

$$\frac{|S|}{|T|} \cdot v_i(C_i) \leq \frac{|S|}{|T|} \cdot v_i(B_i) + \frac{|S|}{|T|} \cdot v_i(W_i) < v_i(A_i \setminus W_i) + v_i(W_i) = v_i(A_i),$$

where the strict inequality follows from Equation 7. If $i \notin T$, then $v_i(C_i) = v_i(B_i)$ (since no goods from W_i are contained in $\cup_{i \in T} A_i$), and we have

$$\frac{|S|}{|T|} \cdot v_i(C_i) = \frac{|S|}{|T|} \cdot v_i(B_i) < v_i(A_i \setminus W_i) \leq v_i(A_i).$$

In both cases, group fairness holds with respect to sets S and T .

A similar argument can be used if we are in the case corresponding to Equation 8, with the relevant strict inequalities replaced by equalities. \square

11.4 Fractional Group Fairness Characterizes MNW

We now show that the dependence on replication invariance can be removed from Theorem 5 if group fairness is strengthened to *fractional group fairness*. Intuitively, fractional group fairness extends group fairness to hold even for groups of agents that may consist of (hypothetical) fractional agents. Denote a fractional subset of agents $S = (x_i^S)_{i \in N}$, where $x_i^S \in [0, 1]$ denotes the fraction of agent i that is included in S . Define $|S| = \sum_{i \in N} x_i^S$.

Definition 19 (Fractional Group Fairness (FGF)). An allocation X is *fractionally group fair* if for all fractional subsets of agents $S, T \subseteq N$ with $|T| \neq 0$, there is no allocation Y of the set of goods $\cup_j x_j^T X_j$ among all agents such that $\frac{|S|}{|T|} \cdot v_i(Y_i) \geq x_i^S \cdot v_i(X_i)$ for all agents $i \in N$ and the inequality is strict for at least one $i \in N$ with $x_i^S > 0$.

Note that we recover group fairness (Definition 6) as a special case of FGF by restricting to the subsets S and T with $x_i^S \in \{0, 1\}$ and $x_i^T \in \{0, 1\}$ for all $i \in N$.

When we strengthen group fairness to fractional group fairness, we obtain a characterization of MNW that does not require replication invariance.

Theorem 10. *Any allocation rule that satisfies fractional group fairness must be a refinement of MNW.*

Proof. Proof. The proof proceeds in a similar manner to the proof of Theorem 5. Let f be an allocation rule, and suppose that there exists an instance I for which there exists an allocation $A \in f(I)$ but $A \notin \text{MNW}(I)$. If A is not Pareto optimal then it violates group fairness. So suppose A is Pareto optimal.

By Lemma 4, there exist agents i, j and $X \subseteq A_i$ with $v_i(X) \in (0, v_i(A_i))$ and $v_j(X) > 0$ such that

$$v_j(X) > v_j(A_j) \cdot \frac{v_i(X)}{v_i(A_i - X)}. \quad (9)$$

Let $T = \{i\}$ (that is, $x_i^T = 1$ and $x_k^T = 0$ for all $k \neq i$) and $S = \{i\} \cup \{\frac{v_i(X)}{v_i(A_i - X)}j\}$ (that is, $x_i^S = 1$, $x_j^S = \frac{v_i(X)}{v_i(A_i - X)}$, and $x_k^S = 0$ otherwise). Let $B_i = A_i - X$ and $B_j = X$. Then

$$\frac{|S|}{|T|} \cdot v_i(B_i) = \left(1 + \frac{v_i(X)}{v_i(A_i - X)}\right) \cdot v_i(A_i - X) = v_i(A_i - X) + v_i(X) = x_i^T \cdot v_i(A_i)$$

and

$$\begin{aligned} \frac{|S|}{|T|} \cdot v_j(B_j) \cdot \frac{1}{x_j^S} &= \left(1 + \frac{v_i(X)}{v_i(A_i - X)}\right) \cdot \frac{v_i(A_i - X)}{v_i(X)} \cdot v_j(X) \\ &= \frac{v_i(A_i)}{v_i(X)} \cdot v_j(X) > \frac{v_i(A_i - X)}{v_i(X)} \cdot v_j(X) > v_j(A_j), \end{aligned}$$

where the final inequality follows from Inequality 9. Fractional group fairness is violated. \square

As a consequence of Theorem 10 and Proposition 7, we know that the set of Maximum Nash Welfare allocations coincides exactly with the set of fractionally group fair allocations. We note that such a statement does not hold if we use only the more permissive axiom of group fairness.

Example 4. Consider an allocation problem with two goods and two agents. We have $v_1(g_1) = v_2(g_2) = 1$, and $v_1(g_2) = v_2(g_1) = \epsilon$ for some small $\epsilon > 0$. The only MNW allocation is to assign g_1 to agent 1 and g_2 to agent 2. But consider allocation A that sets $A_1 = g_1 + \frac{1}{3}g_2$ and $A_2 = \frac{2}{3}g_2$. It is easy to check that A is group fair despite not being an MNW allocation.

12 Omitted Material from Section 7

12.1 Proof of Proposition 6

Recall the statement of Proposition 6.

Proposition 6. *For additive valuation functions over bads, recursive probabilistic serial (RPS) and its polynomial-time variant from Section 3.2 produce randomized allocations that are ex-ante envy-free (EF) and ex-post envy-free up to two bads (EF2).*

Proof. Proof. The reason why the randomized allocation produced by both algorithms is ex-ante EF is the same as in the case of goods. In each step t , the fractional allocation X^t produced by the EATING procedure is envy-free because agents are forced to eat bads at the same rate and an agent is always eating a bad that she prefers the most among all bads not fully consumed. Since each X^t is envy-free, a probability distribution over all X^t -s in different branches is also envy-free. Hence, the expected fractional allocation induced in step t is envy-free for each t . Hence, the overall fractional allocation, which is the sum of fractional allocations induced in all rounds, is also envy-free.

For ex-post EF2, let A be an integral allocation produced by RPS (or its polynomial-time variant). We notice that Lemma 1 still holds for the case of bads. That is, for $t < \lceil m/n \rceil$, if $b_{i,t}$ denotes the bad allocated to agent i in iteration t , then agent i (weakly) prefers $b_{i,t}$ to all bads in M^{t+1} (bads unallocated after t iterations of the algorithm). Now, when $t < \lfloor m/n \rfloor$, we know that another agent j must receive a bad in round $t+1 \leq \lfloor m/n \rfloor$. Hence, for all agents $i, j \in N$ and all $t < \lfloor m/n \rfloor$, we have $v_i(b_{i,t}) \geq v_i(b_{j,t+1})$. Summing over $t < \lfloor m/n \rfloor$, we get $v_i(\{b_{i,t} : 1 \leq t < \lfloor m/n \rfloor\}) \geq v_i(\{b_{j,t} : 2 \leq t \leq \lfloor m/n \rfloor\}) \geq v_i(A_j)$. Note that in the LHS, there are at most two bads missing from A_i : the bad allocated to agent i in round $\lfloor m/n \rfloor$, and a bad which may be allocated to agent i in the final round $\lceil m/n \rceil$. Hence, we have that $v_i(A_i \setminus S_i) \geq v_i(A_j)$ for some $S_i \subseteq A_i$ with $|S_i| \leq 2$. \square

12.2 Proof of Theorem 8

Recall the statement of Theorem 8.

Theorem 8. *For additive valuation functions over bads, there is a strongly polynomial-time algorithm that, given any fractional proportional (Prop) allocation as input, computes an implementation of it using integral allocations that are proportional up to one bad (Prop1). If, in addition, the input is a fractional CEEI allocation, then the integral allocations in the support also satisfy envy-freeness up to one bad more-and-less (EF₁¹).*

Proof. Proof. Given an instance $I = (N, M, (v_i)_{i \in N})$ of bad division, let us define the corresponding instance of good division as $\widehat{I} = (N, M, (\widehat{v}_i)_{i \in N})$ of good division, where $\widehat{v}_i = -v_i$. Given a fractional allocation X of I , we treat it as an allocation of instance \widehat{I} , and apply Lemma 3 to find an implementation of X . Let A^1, \dots, A^ℓ be integral allocations in its support. Note that the ‘reduction’ to goods is used only for the decomposition step, and not for constructing the relevant fractional allocation.

First, suppose X is Prop under I . We want to show that for each $k \in [\ell]$, A^k is Prop1 under I . Since X is proportional, for every $i \in N$, $\widehat{v}_i(X_i) \leq \widehat{v}_i(\overline{I}^m)/n$, where $\widehat{v}_i(\overline{I}^m)$ is agent i ’s valuation for receiving all ‘‘goods’’ fully. Fix $k \in [\ell]$. By Lemma 3, we have that for every agent $i \in N$, either $\widehat{v}_i(A_i^k) \leq \widehat{v}_i(X_i) \leq \widehat{v}_i(\overline{I}^m)/n$, or there exists a bad $r \in A_i^k$ such that $\widehat{v}_i(A_i^k) - \widehat{v}_{i,r} < \widehat{v}_i(X_i) \leq \widehat{v}_i(\overline{I}^m)/n$. Therefore, A^k is Prop1 under I . Note that here, we are using the second part of Lemma 3, whereas in the goods case, we used the first part to establish Prop1.

Next, suppose that X is a fractional CEEI allocation under I . Since this allocation is envy-free [17], and therefore proportional, the aforementioned argument still applies for ex-post Prop1 under I . We show that in this case, A^k is also EF₁¹ under I for each $k \in [\ell]$. As mentioned before, since X is a fractional CEEI allocation under I , Equation (3) holds. Fix a pair of distinct

agents $i, j \in N$. By Lemma 3, either $\widehat{v}_i(A_i^k) \geq \widehat{v}_i(X_i)$, or there exists $b_i^- \notin A_i^k$ with $X_{i, b_i^-} > 0$ such that $\widehat{v}_i(A_i^k \cup \{b_i^-\}) > \widehat{v}_i(X_i)$. Similarly, either $\widehat{v}_j(A_j^k) \leq \widehat{v}_j(X_j)$, or there exists $b_j^+ \in A_j^k$ with $X_{j, b_j^+} < 1$ such that $\widehat{v}_j(A_j^k \setminus \{b_j^+\}) < \widehat{v}_j(X_j)$. To simplify the analysis like in the case of goods, let us assume that the second condition holds in both cases.

By summing the right-hand side inequality in Equation (3) over all $b \in A_i^k \cup \{b_i^-\}$, we get

$$\frac{\widehat{v}_j(A_i^k \cup \{b_i^-\})}{\widehat{v}_j(X_j)} \geq \frac{\widehat{v}_i(A_i^k \cup \{b_i^-\})}{\widehat{v}_i(X_i)} > 1.$$

Thus, $\widehat{v}_j(A_i^k \cup \{b_i^-\}) > \widehat{v}_j(X_j) > \widehat{v}_j(A_j^k \setminus \{b_j^+\})$, implying that $v_j(A_i^k \cup \{b_i^-\}) < v_j(X_j) < v_j(A_j^k \setminus \{b_j^+\})$, and A^k satisfies EF_1^1 , as desired. \square