



Fair Division with Binary Valuations: One Rule to Rule Them All

Daniel Halpern¹, Ariel D. Procaccia¹, Alexandros Psomas^{2(✉)},
and Nisarg Shah³

¹ Harvard University, Cambridge, USA

dhalpern@g.harvard.edu, arielpro@seas.harvard.edu

² Purdue University, West Lafayette, USA

apsomas@cs.purdue.edu

³ University of Toronto, Toronto, Canada

nisarg@cs.toronto.edu

Abstract. We study fair allocation of indivisible goods among agents. Prior research focuses on additive agent preferences, which leads to an impossibility when seeking truthfulness, fairness, and efficiency. We show that when agents have binary additive preferences, a compelling rule—maximum Nash welfare (MNW)—provides all three guarantees. Specifically, we show that deterministic MNW with lexicographic tie-breaking is group strategyproof in addition to being envy-free up to one good and Pareto optimal. We also prove that fractional MNW—known to be group strategyproof, envy-free, and Pareto optimal—can be implemented as a distribution over deterministic MNW allocations, which are envy-free up to one good. Our work establishes maximum Nash welfare as the ultimate allocation rule in the realm of binary additive preferences.

Keywords: Fair division · Mechanism design

1 Introduction

Fair division [13, 28] is a sprawling field that cuts across scientific disciplines. Among its many challenges, the division of indivisible goods—an ostensible oxymoron—is arguably the most popular in recent years. The goods are “indivisible” in the sense that each must be allocated in its entirety to a single agent (think of pieces of jewelry or tickets to different football games in a season). Each agent has her own *valuation function*, which represents the benefit the agent derives from bundles of goods.

A fully expressive model of valuation functions would have to account for combinatorial preferences. Classic examples include a right shoe that is worthless without its matching left shoe (complementarities), and two identical refrigerators (substitutes). However, rich preferences can be difficult to elicit. It is often assumed, therefore, that the valuation functions are *additive*, that is, that each agent’s value for a bundle of goods is the sum of her values for individual goods

in the bundle. Additive valuations strike a balance between expressiveness and ease of elicitation; in particular, each agent need only report her value for each good separately.

Another advantage of additive valuations is that they admit a practical rule that is both (economically) efficient and fair. Specifically, the Maximum Nash Welfare (MNW) solution—which maximizes the product of valuations and, therefore, is obviously Pareto optimal (PO)—is envy-free up to one good (EF1): for any two agents i and j , it is always the case that i prefers her own bundle to that of j , possibly after removing a single good from the latter bundle [16].

The MNW solution, however, is not *strategyproof*, that is, agents can benefit by misreporting their preferences. In fact, under additive valuations, the only Pareto optimal and strategyproof rule is *serial dictatorship*, which is patently unfair [24]. This profound clash between efficiency and truthfulness holds true even when agents can only have three possible values for goods!

The only hope for reconciling efficiency, fairness and truthfulness, therefore, is to assume that agents' values for goods are *binary*. This assumption is not just a theoretical curiosity: while it obviously comes at a significant cost to expressiveness, it leads to extremely simple elicitation. In this sense, it arguably represents another natural point on the conceptual expressiveness-elicitation Pareto frontier. The same bold tradeoff has long been considered sensible in the literature on voting, where binary values are implicitly represented as *approval* votes [12]; in fact, the assumption underlying some of the recent work on approval-based multi-winner elections [17, 26] is nothing but that of binary additive valuations. Thus, it is not surprising that many works in fair division pay special attention to the case of binary additive valuations [1, 6, 11, 20, 23].

With this rather detailed justification for binary additive valuations in mind, our primary research question is this: *do binary additive valuations admit rules that are efficient, fair, and truthful?*

1.1 Our Contribution

We provide a positive answer—and then some. Specifically, Theorem 1 asserts that, under binary additive valuations, a particular form of the MNW solution is Pareto optimal, EF1, *group* strategyproof (even a coalition of agents cannot misreport its members' preferences in a way that benefits them all) and polynomial-time computable.

Furthermore, we show that by randomizing over MNW allocations, we can achieve ex ante envy-freeness (each agent's expected value for their random allocation is at least as high as for any other agent's), ex ante Pareto optimality, ex ante group strategyproofness, and ex post EF1 simultaneously in polynomial time. In other words, randomization allows us to circumvent the mild unfairness that is inherent in deterministic allocations of indivisible goods without losing the other guarantees. In our view, these results are essentially the final word on how to divide indivisible goods under binary additive valuations.

1.2 Related Work

There is an extensive body of work on fair division, much too large to survey here. Instead, we focus on the most closely related work on fair division with *binary* valuations.

The most closely related work is that of Babaioff et al. [5], who, independently and in parallel to our work, also discovered some of the results that we present for the deterministic MNW rule. Specifically, their *prioritized egalitarian* mechanism is identical to our deterministic MNW^{tie} mechanism presented in Sect. 3. They show that this rule is strategyproof, EFX,¹ PO, Lorenz-dominating, and polynomial-time computable. This is very similar to our Theorem 1. The difference is that we strengthen strategyproofness to group strategyproofness, but only establish EF1 (weaker than EFX) and do not establish Lorenz-dominance. We note that the EFX property is also established by Amanatidis et al. [3]. We view these results as complementary to ours. Together, they establish that MNW^{tie} is group strategyproof, EFX, PO, Lorenz-dominating, and polynomial-time computable, making it even more compelling. We note that Babaioff et al. [5] do not study randomized allocation rules, which we focus on in Sect. 4.

Ortega [31] studies a slightly more general problem where there may be multiple copies of each good, but each agent can receive at most one copy of any good. His *egalitarian solution* is identical to our fractional MNW rule in terms of the probability of each good going to each agent, but he does not discuss how to implement these fractional allocations as a distribution over integral allocations with good properties. He shows that this rule is ex ante envy-free, ex ante PO, and ex ante group strategyproof. However, he uses a weaker notion of strategyproofness, where agents are only allowed to report a good that they like as one that they do not like, but not vice-versa. As we note in Sect. 4, in our (standard) setting with a single copy of each good, these guarantees (including the stronger strategyproofness notion, or even group strategyproofness) follow directly from prior work [25]. Hence, our main focus in Sect. 4 is to prove an *ex post* EF1 guarantee, which Ortega [31] does not provide.

Two central concepts in our work are maximum Nash welfare (MNW) and leximin allocations. Aziz and Rey [4] show that under binary additive valuations, all leximin allocations are also MNW allocations. As we observe in Sect. 3, this, together with known properties of the two solutions, immediately implies that the sets of MNW and leximin allocations are identical. Benabbou et al. [7] extend this equivalence to a more general valuation class.

On the computation front, our polynomial-time computability result for the deterministic MNW^{tie} rule builds upon on efficient algorithms by Darmann and Schauer [20] and Barman et al. [6] for finding an MNW allocation under binary additive valuations; specifically, our algorithm starts from an arbitrary MNW allocation computed by either of these algorithms, and then iteratively finds a *special* MNW allocation that MNW^{tie} outputs. Benabbou et al. [7] also show

¹ There are two popular definitions of EFX [3]; this result holds for the stronger one: an allocation is EFX if the envy that one agent has toward another can be eliminated by removing *any* good from the envied agent's bundle.

that an MNW allocation can be computed efficiently under their more general valuation class.

2 Preliminaries

For $k \in \mathbb{N}$, let $[k] = \{1, \dots, k\}$. Let $\mathcal{N} = [n]$ denote a set of *agents*, and \mathcal{M} denote a set of m indivisible *goods*. Each agent i has a *valuation* function $v_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ such that $v_i(\emptyset) = 0$. It is assumed that valuations are additive: $\forall T \subseteq \mathcal{M}$, $v_i(T) = \sum_{g \in T} v_i(\{g\})$. To simplify notation, we write $v_i(g)$ instead of $v_i(\{g\})$.

We focus on a subclass of additive valuations known as binary additive valuations, under which $v_i(g) \in \{0, 1\}$ for all $i \in \mathcal{N}$ and $g \in \mathcal{M}$. We say that agent i *likes* good g if $v_i(g) = 1$. Sometimes it is easier to think of the valuation function of agent i as the set of goods that agent i likes, denoted $V_i = \{g \in \mathcal{M} : v_i(g) = 1\}$. Note that $v_i(T) = |V_i \cap T|$ for all $T \subseteq \mathcal{M}$. For a set of agents $S \subseteq \mathcal{N}$, let $V_S = \bigcup_{i \in S} V_i$ be the set of goods that at least one agent in S likes. The vector of agent valuations $\mathbf{v} = (v_1, \dots, v_n)$ is called the *valuation profile*. A problem instance is given by the tuple $(\mathcal{N}, \mathcal{M}, \mathbf{v})$.

For a set of goods $T \subseteq \mathcal{M}$ and $k \in \mathbb{N}$, let $\Pi_k(T)$ denote the set of partitions of T into k bundles. We say that $\mathbf{A} = (A_1, \dots, A_n)$ is an allocation if $\mathbf{A} \in \Pi_n(T)$ for some $T \subseteq \mathcal{M}$. Here, A_i is the bundle of goods allocated to agent i , and $v_i(A_i)$ is the *utility* to agent i . Let us denote $A_S = \bigcup_{i \in S} A_i$ for $S \subseteq \mathcal{N}$. Let $\mathbb{A} = \bigcup_{T \subseteq \mathcal{M}} \Pi_n(T)$ denote the set of all allocations.

We say that good g is *non-valued* if $v_i(g) = 0$ for all agents i ; all the remaining goods are called *valued*. Let \mathcal{Z} denote the set of non-valued goods. We say that an allocation \mathbf{A} is *complete* if it allocates every valued good, i.e., if $A_{\mathcal{N}} \supseteq \mathcal{M} \setminus \mathcal{Z}$; we say that it is *minimally complete* if it is complete and does not allocate any non-valued goods, i.e., if $A_{\mathcal{N}} = \mathcal{M} \setminus \mathcal{Z}$.

We are interested in *fair* allocations. One of the most prominent notions of fairness is envy-freeness [21].

Definition 1 (Envy-freeness). *An allocation \mathbf{A} is called envy-free (EF) if, for all agents $i, j \in \mathcal{N}$, $v_i(A_i) \geq v_i(A_j)$.*

Envy-freeness requires that no agent prefer another agent's bundle over her own. This cannot be guaranteed (imagine two agents liking a single good). Prior literature focuses on its relaxations, such as envy-freeness up to one good [14, 27], which can be guaranteed.

Definition 2 (Envy-freeness up to one good). *An allocation \mathbf{A} is called envy-free up to one good (EF1) if, for all agents $i, j \in \mathcal{N}$ such that $A_j \neq \emptyset$, there exists $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.*

EF1 requires that it should be possible to remove envy between any two agents by removing at most one good from the envied agent's bundle. We remark that there is a stronger fairness notion called envy-freeness up to the least positively valued good (EFX) [16], which coincides with EF1 under binary additive

valuations.² Finally, another classic desideratum in resource allocation is Pareto optimality, which is a notion of economic efficiency.

Definition 3 (Pareto optimality). *An allocation \mathbf{A} is called Pareto optimal (PO) if there does not exist an allocation \mathbf{A}' such that for all agents $i \in \mathcal{N}$, $v_i(A'_i) \geq v_i(A_i)$, and at least one inequality is strict.*

It is easy to see that with binary additive valuations, Pareto optimality is equivalent to ensuring that each valued good is allocated to one of the agents who likes it, i.e., that the utilitarian social welfare (sum of utilities) is maximized and is equal to the number of valued goods.

3 Deterministic Setting

In this section, our main goal is to establish the existence of a deterministic allocation rule that is fair, efficient, and truthful under binary additive valuations. Our rule builds upon the concept of maximum Nash welfare allocations [16], which we define below. All missing proofs can be found in the full version of this paper.

Definition 4 (Maximum Nash welfare allocation). *We say that \mathbf{A} is a maximum Nash welfare (MNW) allocation if, among the set of allocations \mathbb{A} , it maximizes the number of agents receiving positive utility and, subject to that, maximizes the product of positive utilities. Formally, let $W(\mathbf{A}) = \{i \in \mathcal{N} : v_i(A_i) > 0\}$ and $\mathbb{A}_M = \operatorname{argmax}_{\mathbf{A} \in \mathbb{A}} |W(\mathbf{A})|$. Then, $\operatorname{argmax}_{\mathbf{A} \in \mathbb{A}_M} \prod_{i \in W(\mathbf{A})} v_i(A_i)$ is the set of MNW allocations.*

Even under general additive valuations, all maximum Nash welfare allocations satisfy EF1 and PO [16]. Our work uses a connection between MNW allocations and the classic concept of leximin allocations, that holds under binary additive valuations.

Definition 5 (Leximin comparison). *For an allocation \mathbf{A} , let its utility vector be $(v_1(A_1), \dots, v_n(A_n))$, and its utility profile be the utility vector sorted in a non-descending order. Given two utility profiles $\mathbf{s} = (s_1, \dots, s_n)$ and $\mathbf{s}' = (s'_1, \dots, s'_n)$, we say that \mathbf{s} leximin-dominates \mathbf{s}' , denoted $\mathbf{s} \succ_{\text{lex}} \mathbf{s}'$, if there exists $k \in [n]$ such that $u_k > u'_k$ and $u_r = u'_r$ for all $r < k$. We say that \mathbf{s} weakly leximin-dominates \mathbf{s}' , denoted $\mathbf{s} \succeq_{\text{lex}} \mathbf{s}'$, if $\mathbf{s} \succ_{\text{lex}} \mathbf{s}'$ or $\mathbf{s} = \mathbf{s}'$. Note that this is a total order among utility profiles. We extend these comparisons to utility vectors by applying them to the utility profiles they induce, and call two utility vectors leximin-equivalent if they induce the same utility profile.*

² There are two popular definitions of EFX [3]. The original definition by Caragiannis et al. [16] asks that agent i not envy agent j after removal of any good from agent j 's bundle that has positive value for agent i , whereas a latter definition omits the requirement of "positive value". Under binary additive valuations, the former definition is equivalent to EF1 whereas the latter definition is stronger than EF1.

Definition 6 (Leximin allocations). *We say that an allocation \mathbf{A} is a leximin allocation if, among all allocations, it lexicographically maximizes the utility profile, i.e., maximizes the minimum utility, subject to that maximizes the second minimum, and so on. Thus, leximin allocations are those whose utility profile is the greatest element of the total order \succ_{lex} . We also extend the notions of leximin-dominance and weak leximin-dominance to allocations by comparing their utility vectors.*

Leximin is a refinement of the traditional Rawlsian fairness, which requires maximization of the minimum utility. Plaut and Roughgarden [32] and Freeman et al. [23] study leximin allocations (and variants of this definition), and show that they have related fairness properties as well.

Important to our work is the observation that for binary additive valuations, the sets of leximin and MNW allocations coincide. This is established under a more general valuation class by the contemporary work of Benabbou et al. [7], but for binary additive valuations, this can also be inferred easily from the following observations, which we will use in our work.

Lemma 1. *All leximin allocations have the same utility profile. Further, any allocation with this utility profile is a leximin allocation.*

Lemma 2 (Lemma 21 of Freeman et al. [23]). *Under binary additive valuations, all maximum Nash welfare allocations have the same utility profile. Further, any allocation with this utility profile is a maximum Nash welfare allocation.*

Under binary additive valuations, given the observations above, the sets of MNW and leximin allocations can be either identical or disjoint. Aziz and Rey [4] shows that all leximin allocations are also MNW allocations, which implies that the two sets are identical.

Lemma 3. *Under binary additive valuations, the set of maximum Nash welfare allocations coincides with the set of leximin allocations.*

Henceforth, we will use the terms “MNW allocation” and “leximin allocation” interchangeably. Before we define our deterministic rule, let us define this concept formally. Fix the set of agents \mathcal{N} and the set of goods \mathcal{M} . A *deterministic rule* f takes a valuation profile \mathbf{v} as input and returns an allocation \mathbf{A} . Note that f is not allowed to return ties. We say that f is EF1 (resp. PO) if it always outputs an allocation that is EF1 (resp. PO). The game-theoretic literature offers the following strong desideratum to prevent strategic manipulations by agents.

Definition 7 (Group strategyproofness). *A deterministic rule f is called group strategyproof (GSP) if there do not exist valuation profiles \mathbf{v} and \mathbf{v}' , and a group of agents $C \subseteq \mathcal{N}$, such that $v'_k = v_k$ for all $k \in \mathcal{N} \setminus C$ and $v_j(A'_j) > v_j(A_j)$ for all $j \in C$, where $\mathbf{A} = f(\mathbf{v})$ and $\mathbf{A}' = f(\mathbf{v}')$.*

A weaker requirement, which only imposes the above property for group C of size 1 (i.e. prevents manipulations by a single agent) is commonly known as strategyproofness (SP). We are now ready to define our rule, which chooses a special MNW allocation.

Definition 8 (MNW^{tie}). *The deterministic rule MNW^{tie} returns an allocation \mathbf{A} such that:*

1. \mathbf{A} is an MNW allocation with lexicographically greatest utility vector among all MNW allocations (i.e., among all MNW allocations, it maximizes $v_1(A_1)$, subject to that maximizes $v_2(A_2)$, and so on);³ and
2. \mathbf{A} is minimally complete (i.e. $A_{\mathcal{N}} = \mathcal{M} \setminus \mathcal{Z}$).

If there are several allocations satisfying both conditions, MNW^{tie} arbitrarily picks one.

First, observe that MNW^{tie} is well-defined, i.e., that the set of allocations satisfying both conditions is non-empty. Indeed, the set of allocations satisfying the first condition is trivially non-empty. And for any allocation in this set, there is a corresponding minimally complete allocation—obtained by throwing away all non-valued goods—which has the same utility vector, and therefore still satisfies the first condition.

The following result establishes the compelling properties of MNW^{tie} . The key idea for polynomial-time computability is as follows. Darmann and Schauer [20] and Barman et al. [6] show that under binary additive valuations, an MNW allocation can be computed efficiently. Starting from this MNW allocation, we keep moving to lexicographically better MNW allocations, as in the definition of MNW^{tie} . The algorithm and proof are formally presented in the full version of this paper.

Theorem 1. *Under binary additive valuations, MNW^{tie} is envy-free up to one good, Pareto optimal, group strategyproof, and polynomial-time computable.*

4 Randomized Setting

In the previous section, we established the existence of a deterministic rule which is EF1, PO, and GSP. For deterministic rules, it is necessary to relax EF to EF1. For example, in case of a single good that is liked by two agents, giving it to either agent would be EF1 but not EF. However, if one is willing to randomize, the natural solution of assigning the good to an agent chosen at random would be “ex ante EF” in addition to being “ex post EF1”. This is because each deterministic allocation in the support is EF1, but in expectation, no agent envies the other. This leads to a natural question. *Can randomness help achieve ex ante EF and ex post EF1, in addition to PO and GSP?*

In this section, we answer this question affirmatively for binary additive valuations. In parallel to our work, Freeman et al. [22] show that ex ante EF and ex post EF1 can be achieved simultaneously even under general additive valuations, but they show an impossibility when ex ante PO is added to the combination.

³ We note that tie-breaking by agent index is without loss of generality. One can break ties according to any given ordering of the agents, and the corresponding rule will still satisfy all the desiderata.

Our positive result circumvents this impossibility for binary additive valuations. Additionally, it satisfies GSP, which Freeman et al. [22] do not consider. Missing proofs can be found in the full version of this paper.

Let us first formally extend our framework to include randomness.

Definition 9 (Fractional and randomized allocations). A fractional allocation $\mathbf{A} = (A_1, \dots, A_n)$ is such that $A_i(g) \in [0, 1]$ denotes the fraction of good g allocated to agent i and $\sum_{i \in \mathcal{N}} A_i(g) \leq 1$ for each good g . A randomized allocation $\overline{\mathbf{A}}$ is a probability distribution over deterministic allocations.

There is a natural fractional allocation \mathbf{A} associated with each randomized allocation $\overline{\mathbf{A}}$, where $A_i(g)$ is the probability of good g being allocated to agent i under $\overline{\mathbf{A}}$. In this case, we say that randomized allocation $\overline{\mathbf{A}}$ implements fractional allocation \mathbf{A} . There may be several randomized allocations implementing a given fractional allocation.

We refer to the expected utility of agent i under a randomized allocation $\overline{\mathbf{A}}$ as simply the utility of agent i under $\overline{\mathbf{A}}$. Note that this is equal to the utility of agent i from the corresponding fractional allocation \mathbf{A} , defined as $v_i(A_i) = \sum_{g \in \mathcal{M}} A_i(g) \cdot v_i(g)$. With this notation, the definitions of envy-freeness and Pareto optimality extend naturally to fractional allocations.⁴ We say that a randomized allocation $\overline{\mathbf{A}}$ is ex ante envy-free (resp. ex ante Pareto optimal) if the corresponding fractional allocation \mathbf{A} is envy-free (resp. Pareto optimal).

With a fixed set of agents \mathcal{N} and a fixed set of goods \mathcal{M} , a randomized rule f takes a valuation profile \mathbf{v} as input and returns a randomized allocation $\overline{\mathbf{A}}$. We say that f is ex ante envy-free (resp. ex ante Pareto optimal) if it always returns a randomized allocation that is ex ante envy free (resp. ex ante Pareto optimal). We say that f is ex ante group strategyproof if no group of agents can misreport their preferences so that each agent in the group receives strictly greater expected utility. Note that these ex ante guarantees depend only on the fractional allocation corresponding to the randomized allocation returned by f . Hence, when talking about ex ante guarantees, we will think of the randomized rule f as directly returning a fractional allocation. However, when talking about ex post guarantees, we would need to specify which randomized allocation f returns.

Definition 10 (Ex post EF1). We say that a randomized allocation $\overline{\mathbf{A}}$ is ex post envy-free up to one good if each deterministic allocation in its support is EF1. A randomized rule is ex post EF1 if it always returns a randomized allocation that is ex post EF1.

Fractional leximin allocations, like their deterministic counterpart, lexicographically maximize the utility profile among all fractional allocations. The same can be said about fractional MNW allocations; however, we can skip the first step of maximizing the number of agents who receive positive utility because

⁴ In case of Pareto optimality of a fractional allocation, we require that no other fractional allocation Pareto-dominate it.

in the fractional case we can simultaneously give positive utility to every agent who likes at least one good (and thus can possibly get positive utility).

Definition 11 (Fractional MNW allocations). *We say that a fractional allocation is a fractional maximum Nash welfare allocation if it maximizes the product of utilities of agents who do not have zero value for every good.*

Bogomolnaia and Moulin [9], Bogomolnaia et al. [10], and Kurokawa et al. [25] study fractional leximin allocations under an assignment setting, and establish several desirable properties. In addition, fractional MNW allocations, also known as competitive equilibria with equal incomes (CEEI), are widely studied in fair division with additive valuations [18, 19, 30, 33]. Our first result shows that under binary additive valuations, these two concepts coincide.

Theorem 2. *Under binary additive valuations, the set of fractional leximin allocations coincides with the set of fractional maximum Nash welfare allocations. All such allocations have identical utility vectors.*

Note that the identical utility vector guarantee in Theorem 2 is much stronger than the identical utility profile guarantee in the deterministic case.

Even under general additive valuations, it is known that every fractional MNW allocation is ex ante EF and ex ante PO [33], and one such allocation can be computed in strongly polynomial time [30, 34]. Hence, these properties carry over to our binary additive valuations domain, and due to Theorem 2, also apply to fractional leximin allocations.

For ex ante GSP, we build on the literature on fractional leximin allocations. Kurokawa et al. [25] show that returning a fractional leximin allocation satisfies ex ante EF, ex ante PO, and ex ante GSP whenever four key requirements are satisfied. We describe them in the full version, and show that they are easily satisfied under binary additive valuations, if we return a minimally complete leximin allocation. Hence, we define our fractional leximin/MNW rule to always return a minimally complete fractional leximin/MNW allocation (like our deterministic rule MNW^{tie}).

Definition 12 (Fractional maximum Nash welfare rule). *The fractional maximum Nash welfare rule returns a minimally complete fractional maximum Nash welfare allocation.*

Theorem 3. *Under binary additive valuations, every fractional maximum Nash welfare (equivalently, leximin) allocation is ex ante envy-free and ex ante Pareto optimal. Further, the fractional maximum Nash welfare rule is ex ante group strategyproof.*

The only missing property at this point is ex post EF1. Therefore, the main question we seek to answer in this section is the following: *Can every fractional MNW allocation be implemented as a distribution over deterministic EF1 allocations?* We go one step further and show that it can in fact be implemented as a distribution over deterministic MNW allocations, which are in turn EF1.

Our main tool is the bihierarchy framework introduced by Budish et al. [15], which is a generalization of the classic Birkhoff-von Neumann theorem [8, 29]. At a high level, the framework allows implementing any fractional allocation \mathbf{A} using deterministic allocations which satisfy a set of constraints, as long as the set of constraints forms a bihierarchy structure and the fractional allocation itself satisfies those constraints.

In our case, we start with a minimally complete fractional MNW allocation \mathbf{A}^* . Let u_i^* denote the utility to agent i under this allocation. We want to implement this as a randomized allocation. We impose the following constraints on a deterministic allocation \mathbf{A} in the support, where \mathbf{A} is represented as a matrix in which $A_i(g) \in \{0, 1\}$ indicates whether good g is allocated to agent i .

$$\begin{aligned} \mathcal{H}_1 : \sum_{i \in \mathcal{N}} A_i(g) &= \sum_{i \in \mathcal{N}} A_i^*(g), \forall g \in \mathcal{M}, \\ \mathcal{H}_2 : \lfloor u_i^* \rfloor &\leq \sum_{g \in \mathcal{M}} A_i(g) \cdot v_i(g) \leq \lceil u_i^* \rceil, \forall i \in \mathcal{N}. \end{aligned} \quad (1)$$

The first family of constraints ensures that under each deterministic allocation \mathbf{A} , the set of goods allocated matches that under \mathbf{A}^* . Since \mathbf{A}^* is minimally complete, this implies that \mathbf{A} must be minimally complete as well. Crucially, the second family of constraints ensures that each agent has utility that is either the floor or the ceiling of her utility under \mathbf{A}^* . That is, \mathbf{A} is not allowed to stray far from \mathbf{A}^* .

It can be checked that these constraints form a bihierarchy (each of \mathcal{H}_1 and \mathcal{H}_2 is a hierarchy); for a formal definition of a hierarchy, we refer the reader to the work of Budish et al. [15]. Importantly, they also provide a polynomial-time algorithm that computes a random allocation such that (a) it implements the fractional allocation \mathbf{A}^* , and (b) each deterministic allocation \mathbf{A} in its support satisfies the constraints in 1. We show that in this case, every deterministic allocation in the support must be a deterministic MNW allocation, yielding the desired result.

Theorem 4. *Under binary additive valuations, given any fractional maximum Nash welfare allocation, one can compute, in polynomial time, a randomized allocation which implements it and has only deterministic maximum Nash welfare allocations in its support.*

Proof. Let \mathbf{A}^* be a given fractional MNW allocation with utility vector \mathbf{u}^* . Let \mathbf{A} be the randomized allocation implementing \mathbf{A}^* that is returned by the polynomial-time algorithm of Budish et al. [15] with the bihierarchy constraints in Eq. 1. Let \mathcal{A} denote the set of deterministic allocations in the support of \mathbf{A} . Our goal is to show that every allocation in \mathcal{A} is an MNW allocation.

First, let us partition the set of agents \mathcal{N} into sets S_1, \dots, S_t such that any two agents i and j are in the same set if and only if $\lfloor u_i^* \rfloor = \lfloor u_j^* \rfloor$. For $k \in [t]$, let L_k denote the common floor of utilities of agents in S_k under \mathbf{A}^* , and $U_k = L_k + 1$. Hence, for $k \in [t]$ and each agent $i \in S_k$, $u_i^* \in [L_k, U_k)$. Further, order the sets so that $U_k \leq L_{k+1}$ for each $k \in [t - 1]$. This ensures that if $i \in S_r$, $j \in S_{r'}$, and $r' > r$, then $u_j^* > u_i^*$.

We argue that for each $k \in [t]$, the agents in $\cup_{r \in [k]} S_r$ must be fully allocated all of the goods that they like (i.e. all the goods in $V_{\cup_{r \in [k]} S_r}$) under \mathbf{A}^* , resulting in $\sum_{r \in [k]} \sum_{i \in S_r} u_i^* = |V_{\cup_{r \in [k]} S_r}|$. If this is not true, then a positive fraction of some good $g \in V_{\cup_{r \in [k]} S_r}$ must be allocated to an agent $j \in S_{r'}$ for $r' > k$. Let $i \in \cup_{r \in [k]} S_r$ be an agent such that $g \in V_i$. Let $r \in [k]$ be such that $i \in S_r$. Then, by the above argument, we know that $u_j^* > u_i^*$. However, then, transferring a sufficiently small fraction of g from agent j to agent i in \mathbf{A}^* will improve the Nash welfare, which contradicts the fact that \mathbf{A}^* is a fractional MNW allocation.

Note that in any deterministic allocation \mathbf{A} , $|V_{\cup_{r \in [k]} S_r}|$ is the highest utility that agents in $\cup_{r \in [k]} S_r$ can collectively have; hence, in any feasible utility vector \mathbf{u} ,

$$\sum_{r \in [k]} \sum_{i \in S_r} u_i \leq \sum_{r \in [k]} \sum_{i \in S_r} u_i^*, \forall k \in [t]. \tag{2}$$

Because a convex combination of allocations in \mathcal{A} yields the allocation \mathbf{A}^* , and utilities are additive, a convex combination of their utility vectors yields the utility vector \mathbf{u}^* . Hence, for the utility vector \mathbf{u} of any allocation in \mathcal{A} , Eq. 2 must hold with equality. Further, by subtracting each equation from the next, we get that it must further satisfy the following. Here, \mathcal{H}_2 is from the bihierarchy constraints (Eq. 1).

$$\begin{aligned} \mathcal{H}_2 : [u_i^*] \leq u_i \leq \lceil u_i^* \rceil, \forall i \in \mathcal{N}, \\ \mathcal{H}_3 : \sum_{i \in S_k} u_i = \sum_{i \in S_k} u_i^*, \forall k \in [t]. \end{aligned} \tag{3}$$

We say that a utility vector is a *rounded* if it satisfies the constraints in Eq. (3), and say that a deterministic allocation is *rounded* if it has a rounded utility vector. We have already established that every allocation in \mathcal{A} is a rounded allocation. The following lemma completes the proof of Theorem 4.

Lemma 4. *The set of rounded allocations coincides with the set of maximum Nash welfare allocations.*

Let us amend the definition of the fractional MNW rule so that it uses Theorem 4 to implement a minimally complete fractional MNW allocation. Then, we have the following.

Corollary 1. *Under binary additive valuations, the fractional maximum Nash welfare rule is ex ante envy-free, ex ante Pareto optimal, ex ante group strategyproof, ex post envy-free up to one good, and polynomial-time computable.*

5 Discussion

To recap, we showed that under binary additive valuations a deterministic variant of the maximum Nash welfare rule is envy-free up to one good (EF1), Pareto optimal (PO), and group strategyproof (GSP). We also demonstrated that its randomized variant is ex ante EF, ex ante PO, ex ante GSP, and ex post EF1. All our rules are polynomial-time computable.

Amanatidis et al. [2] show that under general additive valuations, there is no deterministic rule that is envy-free up to one good (EF1) and strategyproof, even with two agents and $m \geq 5$ goods. At first glance, Theorem 1, which establishes MNW^{tie} as both GSP and EF1, seems to show that this impossibility result does not hold for the special case of binary additive valuations. However, the impossibility result of Amanatidis et al. [2] only applies to rules that allocate all the goods; by contrast, MNW^{tie} does not allocate non-valued goods. This begs the following question: *Under binary additive valuations, is there a deterministic rule that allocates all the goods and achieves EF1, PO, and GSP?* In the full version, we show that this cannot be achieved by any variant of MNW.

Another open question is whether the ex ante GSP guarantee of Corollary 1 can be strengthened to ex post GSP, which would require the randomized rule to be implementable as a probability distribution over deterministic GSP rules.

Modulo these minor caveats, though, our results are the strongest one could possibly hope for in the domain of binary additive valuations.

References

1. Aleksandrov, M., Aziz, H., Gaspers, S., Walsh, T.: Online fair division: analysing a food bank problem. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), pp. 2540–2546 (2015)
2. Amanatidis, G., Birmpas, G., Christodoulou, G., Markakis, E.: Truthful allocation mechanisms without payments: characterization and implications on fairness. In: Proceedings of the 18th ACM Conference on Economics and Computation (EC), pp. 545–562 (2017)
3. Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., Hollender, A., Voudouris, A.A.: Maximum Nash welfare and other stories about EFX. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI) (2020, forthcoming)
4. Aziz, H., Rey, S.: Almost group envy-free allocation of indivisible goods and chores. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI) (2020)
5. Babaioff, M., Ezra, T., Feige, U.: Fair and truthful mechanisms for dichotomous valuations. arXiv preprint [arXiv:2002.10704](https://arxiv.org/abs/2002.10704) (2020)
6. Barman, S., Krishnamurthy, S.K., Vaish, R.: Greedy algorithms for maximizing Nash social welfare. In: Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 7–13 (2018)
7. Benabbou, N., Chakraborty, M., Igarashi, A., Zick, Y.: Finding fair and efficient allocations when valuations don't add up. arXiv preprint [arXiv:2003.07060](https://arxiv.org/abs/2003.07060) (2020)
8. Birkhoff, G.: Three observations on linear algebra. *Universidad Nacional de Tucumán, Revista A* **5**, 147–151 (1946)
9. Bogomolnaia, A., Moulin, H.: Random matching under dichotomous preferences. *Econometrica* **72**, 257–279 (2004)
10. Bogomolnaia, A., Moulin, H., Stong, R.: Collective choice under dichotomous preferences. *J. Econ. Theory* **122**(2), 165–184 (2005)
11. Bouveret, S., Lemaître, M.: Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Auton. Agents Multi-Agent Syst.* **30**(2), 259–290 (2016)

12. Brams, S.J., Fishburn, P.C.: Approval Voting, 2nd edn. Springer, New York (2007). <https://doi.org/10.1007/978-0-387-49896-6>
13. Brams, S.J., Taylor, A.D.: Fair Division: From Cake-Cutting to Dispute Resolution. Cambridge University Press, Cambridge (1996)
14. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *J. Polit. Econ.* **119**(6), 1061–1103 (2011)
15. Budish, E., Che, Y.K., Kojima, F., Milgrom, P.: Designing random allocation mechanisms: theory and applications. *Am. Econ. Rev.* **103**(2), 585–623 (2013)
16. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum Nash welfare. *ACM Trans. Econ. Comput.* **7**(3), article 12 (2019)
17. Cheng, Y., Jiang, Z., Munagala, K., Wang, K.: Group fairness in committee selection. In: Proceedings of the 20th ACM Conference on Economics and Computation (EC), pp. 263–279 (2019)
18. Cole, R., Gkatzelis, V.: Approximating the Nash social welfare with indivisible items. *SIAM J. Comput.* **47**(3), 1211–1236 (2018)
19. Cole, R., Gkatzelis, V., Goel, G.: Mechanism design for fair division: allocating divisible items without payments. In: Proceedings of the 14th ACM Conference on Economics and Computation (EC), pp. 251–268 (2013)
20. Darmann, A., Schauer, J.: Maximizing Nash product social welfare in allocating indivisible goods. *Eur. J. Oper. Res.* **247**(2), 548–559 (2015)
21. Foley, D.K.: Resource allocation and the public sector. *Yale Economics Essays*, vol. 7, pp. 45–98 (1967)
22. Freeman, R., Shah, N., Vaish, R.: Best of both worlds: ex-ante and ex-post fairness in resource allocation. In: Proceedings of the 21st ACM Conference on Economics and Computation (EC) (2020, forthcoming)
23. Freeman, R., Sikdar, S., Vaish, R., Xia, L.: Equitable allocations of indivisible goods. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 280–286 (2019)
24. Klaus, B., Miyagawa, E.: Strategy-proofness, solidarity, and consistency for multiple assignment problems. *Int. J. Game Theory* **30**, 421–435 (2001)
25. Kurokawa, D., Procaccia, A.D., Shah, N.: Leximin allocations in the real world. *ACM Trans. Econ. Comput.* **6**(3–4), article 11 (2018)
26. Lackner, M., Skowron, P.: A quantitative analysis of multi-winner rules. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), pp. 407–413 (2019)
27. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 6th ACM Conference on Economics and Computation (EC), pp. 125–131 (2004)
28. Moulin, H.: Fair Division and Collective Welfare. MIT Press, Cambridge (2003)
29. von Neumann, J.: A certain zero-sum two-person game equivalent to the optimal assignment problem. In: Kuhn, W., Tucker, A.W. (eds.) *Contributions to the Theory of Games*, vol. 2, pp. 5–12. Princeton University Press (1953)
30. Orlin, J.B.: Improved algorithms for computing fisher’s market clearing prices: computing fisher’s market clearing prices. In: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), pp. 291–300 (2010)
31. Ortega, J.: Multi-unit assignment under dichotomous preferences. *Math. Soc. Sci.* **103**, 15–24 (2020)

32. Plaut, B., Roughgarden, T.: Almost envy-freeness with general valuations. In: Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 2584–2603 (2018)
33. Varian, H.: Equity, envy and efficiency. *J. Econ. Theory* **9**, 63–91 (1974)
34. Vég h, L.A.: Concave generalized flows with applications to market equilibria. *Math. Oper. Res.* **39**(2), 573–596 (2013)