

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

Daniel Halpern
University of Toronto

daniel.halpern@mail.utoronto.ca

Alexandros Psomas
Google Research
alexpsomi@cs.berkeley.edu

Ariel D. Procaccia
Harvard University

arielpro@seas.harvard.edu

Nisarg Shah
University of Toronto
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.

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. It is not surprising, therefore, that several papers in fair division pay special attention to the case of binary additive valuations [1; 6; 11; 20; 22].

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 (Theorems 3 and 4) 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 Section 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, and 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 Section 4.

Ortega [30] 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

¹There are two popular definitions of EFX (see [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.

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 Section 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 Section 4 is to prove an *ex post* EF1 guarantee, which Ortega [30] does not provide.

Two central concepts in our work are those of 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 Section 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 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 is endowed with 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 [15; 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

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.

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.

Definition 6 (Leximin allocations). We say that \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 [31] and Freeman *et al.* [22] study leximin allocations (and variants of this definition), and show that they have related fairness properties as well.

²There 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.

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.*

Proof. This is because lexicographic comparison is a total order among utility profiles, and leximin allocations, by definition, are those whose utility profile is its greatest element. \square

Lemma 2 (Lemma 21 of Freeman *et al.* [22]). *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.

³We 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.

The following result establishes the compelling properties of MNW^{tie} . The key idea behind the 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 is formally presented as Algorithm 1.

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

Before diving into the proof, we need another concept that we will use repeatedly. The *graph of an allocation \mathbf{A}* is a directed graph $G(\mathbf{A}) = (V, E)$, where V contains a vertex for each agent, and there is a directed edge $(i, j) \in E$ if and only if there is a good in agent j 's bundle that agent i likes (i.e., $A_j \cap V_i \neq \emptyset$). Given a path $P = (u_1, \dots, u_k)$ in $G(\mathbf{A})$, let $P(\mathbf{A})$ denote an allocation obtained by transferring a good $g \in A_{u_{\ell+1}} \cap V_{u_\ell}$ from agent $u_{\ell+1}$ to agent u_ℓ for each $\ell \in [k-1]$; we refer to this operation as *passing back along P* . We characterize MNW allocations in terms of non-existence of a special path in their graph.

Lemma 4. *Let \mathbf{A} be a Pareto optimal allocation, P be a path from agent i to agent j in $G(\mathbf{A})$, and $\mathbf{A}' = P(\mathbf{A})$ be obtained by passing back along P . Then $v_j(\mathbf{A}'_j) = v_j(\mathbf{A}_j) - 1$, $v_i(\mathbf{A}'_i) = v_i(\mathbf{A}_i) + 1$, and $v_k(\mathbf{A}'_k) = v_k(\mathbf{A}_k)$ for all $k \in \mathcal{N} \setminus \{i, j\}$.*

Proof. Note that if good g is being passed from agent $u_{\ell+1}$ to agent u_ℓ on path P , then by definition u_ℓ must like g . Hence, g is a valued good. Thus, by PO, $u_{\ell+1}$ must like g as well. Thus, each agent on P except i and j loses a good she likes and receives a good she likes, agent j only loses a good she likes, and agent i only receives a good she likes. \square

Lemma 5. *A Pareto optimal allocation \mathbf{A} is an MNW allocation if and only if there is no directed path from an agent i to an agent j in $G(\mathbf{A})$ such that $v_j(\mathbf{A}_j) > v_i(\mathbf{A}_i) + 1$.*

Proof. Lemma 3 of Barman *et al.* [6] establishes that \mathbf{A} is an MNW allocation if and only if there is no directed path P such that passing back along P strictly increases Nash welfare.⁴ Given that \mathbf{A} is PO, Lemma 4 implies that this is equivalent to $(v_j(\mathbf{A}_j) - 1) \cdot (v_i(\mathbf{A}_i) + 1) > v_j(\mathbf{A}_j) \cdot v_i(\mathbf{A}_i)$, which is equivalent to $v_j(\mathbf{A}_j) > v_i(\mathbf{A}_i) + 1$. \square

We are now ready to prove Theorem 1.

Proof. Caragiannis *et al.* [16] already establish that all MNW allocations are EF1 and PO, even for general additive valuations. Hence, MNW^{tie} is also trivially EF1 and PO.

Group strategyproofness. We now establish that it is GSP. Note that this holds regardless of how ties are broken among allocations satisfying the two conditions in the definition of MNW^{tie} .

First, notice that if $\mathbf{A} = \text{MNW}^{\text{tie}}(\mathbf{v})$, then \mathbf{A} is minimally complete and PO. Hence, if agent i receives good g , she must like it. In other words, $A_i \subseteq V_i$, and thus, $v_i(\mathbf{A}_i) = |A_i|$ for each agent $i \in \mathcal{N}$. Consequently, $A_U \subseteq V_U$ for every subset of agents $U \subseteq \mathcal{N}$. We will use this observation repeatedly.

Next, for an allocation \mathbf{A} and agent $i \in \mathcal{N}$, define $L_{\mathbf{A}}^i = \{j \in \mathcal{N} \mid v_j(\mathbf{A}_j) < v_i(\mathbf{A}_i)\}$ to be the set of agents who have strictly less utility than agent i , and define $S_{\mathbf{A}}^i$ to be the set of agents reachable from $L_{\mathbf{A}}^i \cup \{i\}$ in $G(\mathbf{A})$. The following lemma shows that agents in $S_{\mathbf{A}}^i$ must collectively receive all the goods that they like.

Lemma 6. *If $\mathbf{A} = \text{MNW}^{\text{tie}}(\mathbf{v})$, then for each agent $i \in \mathcal{N}$, we have $A_{S_{\mathbf{A}}^i} = V_{S_{\mathbf{A}}^i}$.*

⁴Technically, either more agents receive positive utility, or the product of positive utilities increases.

Proof. We have already established that $A_{S_{\mathbf{A}}^i} \subseteq V_{S_{\mathbf{A}}^i}$. Suppose for contradiction that there exists a good $g \in V_{S_{\mathbf{A}}^i} \setminus A_{S_{\mathbf{A}}^i}$. Then, by the construction of $G(\mathbf{A})$, there would have been an edge from an agent in $S_{\mathbf{A}}^i$ who likes g to an agent outside of $S_{\mathbf{A}}^i$ who is allocated g under \mathbf{A} (note that g is valued, so it must be allocated under \mathbf{A}). However, the definition of $S_{\mathbf{A}}^i$ implies that it cannot have any outgoing edges, otherwise the set of agents reachable from $L_{\mathbf{A}}^i \cup \{i\}$ could be expanded. Hence, we have $A_{S_{\mathbf{A}}^i} = V_{S_{\mathbf{A}}^i}$. \square

Next, we show that even though $S_{\mathbf{A}}^i$ contains all agents reachable from $L_{\mathbf{A}}^i \cup \{i\}$, an agent in $S_{\mathbf{A}}^i$ cannot have much higher utility than agent i does.

Lemma 7. *If $\mathbf{A} = \text{MNW}^{\text{tie}}(\mathbf{v})$, then for each agent $i \in \mathcal{N}$ and each agent $j \in S_{\mathbf{A}}^i$, we have that $v_j(A_j) \leq v_i(A_i) + 1$, and if $j \geq i$, then $v_j(A_j) \leq v_i(A_i)$.*

Proof. This is trivial for $j = i$, so assume $j \neq i$, and suppose for contradiction that the statement is false. Hence, there exist agents $i \in \mathcal{N}$ and $j \in S_{\mathbf{A}}^i$ such that either $v_j(A_j) \geq v_i(A_i) + 2$, or $v_j(A_j) = v_i(A_i) + 1$ and $j > i$.

Since $j \in S_{\mathbf{A}}^i$, there exists a path from an agent $k \in L_{\mathbf{A}}^i \cup \{i\}$ to agent j . Further, $k \in L_{\mathbf{A}}^i \cup \{i\}$ implies that $v_k(A_k) \leq v_i(A_i)$ by definition.

Now, in the former case, we would have that there exists a path from agent k to agent j and $v_j(A_j) \geq v_i(A_i) + 2 \geq v_k(A_k) + 2$. However, this contradicts Lemma 5.

In the latter case, we consider two sub-cases. If $k \neq i$, then $k \in L_{\mathbf{A}}^i$. Hence, $v_k(A_k) < v_i(A_i)$. This implies $v_j(A_j) = v_i(A_i) + 1 \geq v_k(A_k) + 2$, which leads to a contradiction as pointed out above. If $k = i$, then we have a path from agent i to agent $j > i$ with $v_j(A_j) = v_i(A_i) + 1$. Once again, passing back along this path would result in an allocation \mathbf{A}' under which $v_t(A'_t) = v_t(A_t)$ for all $t \neq i, j$, $v_i(A'_i) = v_i(A_i) + 1 = v_j(A_j)$, and $v_j(A'_j) = v_j(A_j) - 1 = v_i(A_i)$. Since \mathbf{A}' has the same utility profile as \mathbf{A} , by Lemma 2, \mathbf{A}' is also a maximum Nash welfare allocation. Further, since a lower-indexed agent receives higher utility, \mathbf{A}' is lexicographically better than \mathbf{A} , which contradicts the fact that \mathbf{A} was returned by MNW^{tie} . \square

We are now ready to show that MNW^{tie} is GSP. Suppose for contradiction that there exist valuation profiles \mathbf{v} and \mathbf{v}' , and a set of agents $C \subseteq \mathcal{N}$ such that $v_t = v'_t$ for all $t \notin C$ and $v_j(A_j^{\text{lie}}) > v_j(A_j^{\text{truth}})$ for each $j \in C$, where $\mathbf{A}^{\text{truth}} = \text{MNW}^{\text{tie}}(\mathbf{v})$ and $\mathbf{A}^{\text{lie}} = \text{MNW}^{\text{tie}}(\mathbf{v}')$.

Let $i = \min [\text{argmin}_{t \in C} v_t(A_t^{\text{truth}})]$ be the agent in C who has the lowest index among all agents in C having the minimum utility under honest reporting. For simplicity, let us denote $S = S_{\mathbf{A}^{\text{truth}}}^i$. We have that for every $j \in C$, $|V_j \cap A_j^{\text{truth}}| < |V_j \cap A_j^{\text{lie}}|$. Further, since $A_j^{\text{truth}} \subseteq V_j$, this simplifies to $|A_j^{\text{truth}}| < |V_j \cap A_j^{\text{lie}}|$. When $j \in S \cap C$, we get $|A_j^{\text{truth}}| < |V_j \cap A_j^{\text{lie}}| \leq |V_S \cap A_j^{\text{lie}}|$ because $V_j \subseteq V_S$.

Let $R \subseteq S$ be the set of agents in S from which some agent in C is reachable in $G(\mathbf{A}^{\text{lie}})$. We now establish that some non-manipulating agent in R must receive strictly fewer goods under \mathbf{A}^{lie} than under $\mathbf{A}^{\text{truth}}$.

Lemma 8. *There exists $j^* \in R \setminus C$ with $|A_j^{\text{lie}}| < |A_j^{\text{truth}}|$.*

Proof. Suppose for a contradiction that for all $j \in R \setminus C$, $|A_j^{\text{lie}}| \geq |A_j^{\text{truth}}|$. Take a $j \in R \setminus C$. Since $j \notin C$, she reports $v'_j = v_j$. Hence, we have $A_j^{\text{lie}} \subseteq V'_j = V_j$. Further, since $j \in R \subseteq S$, we have $V_j \subseteq V_S$ by definition. We conclude that for each $j \in R \setminus C$, $A_j^{\text{lie}} \subseteq V_S$, so $|A_j^{\text{lie}} \cap V_S| = |A_j^{\text{lie}}| \geq |A_j^{\text{truth}}|$.

Additionally, for each $j \in R \cap C \subseteq C$, we have that $|A_j^{\text{lie}} \cap V_S| \geq |A_j^{\text{lie}} \cap V_j| > |A_j^{\text{truth}}|$. Since bundles of an allocation are disjoint, we can add these inequalities over all $j \in (R \setminus C) \cup (R \cap C) = R$ to get $|A_R^{\text{lie}} \cap V_S| > |A_R^{\text{truth}}|$. The inequality is strict because $R \cap C \neq \emptyset$ as $i \in R \cap C$ by definition. Now, recall that by Lemma 6, $A_S^{\text{truth}} = V_S$. Hence, this becomes $|A_R^{\text{lie}} \cap A_S^{\text{truth}}| > |A_R^{\text{truth}}|$.

This implies that there must exist a good g that is in both A_R^{lie} and A_S^{truth} but not in A_R^{truth} . Therefore, there exist agents $t \in R$ and $k \in S \setminus R$ such that $g \in A_t^{\text{lie}}$ and $g \in A_k^{\text{truth}}$. The latter implies $v_k(g) = 1$ due to Pareto optimality of $\mathbf{A}^{\text{truth}}$.

Since $k \notin R$, by definition k does not have a path to an agent in C under $G(\mathbf{A}^{\text{lie}})$. This trivially implies $k \notin C$ since every vertex is reachable from itself. Since only members of C changed their reported valuations, $v'_k(g) = v_k(g) = 1$. It follows that there must be an edge from agent k to agent t in $G(\mathbf{A}^{\text{lie}})$. Thus, all vertices reachable from t are also reachable from k . But then, $t \in R$ implies $k \in R$, which is a contradiction. \square

Consider an agent $j^* \in R \setminus C$ as per Lemma 8. Since $j^* \in R$, there must exist a path P from j^* to some agent $k \in C$. Let \mathbf{A}' denote the allocation obtained by passing back along path P . We show that \mathbf{A}' must be preferred to \mathbf{A}^{lie} by MNW^{tie} given valuation profile \mathbf{v}' , contradicting the fact that $\text{MNW}^{\text{tie}}(\mathbf{v}') = \mathbf{A}^{\text{lie}}$.

Note that since \mathbf{A}^{lie} is PO under valuation profile \mathbf{v}' , when constructing \mathbf{A}' from \mathbf{A}^{lie} , we get $v'_t(A'_t) = v'_t(A_t^{\text{lie}})$ for all $t \neq j^*, k$, $v'_k(A'_k) = v'_k(A_k^{\text{lie}}) - 1$, and $v_{j^*}(A'_{j^*}) = v_{j^*}(A_{j^*}^{\text{lie}}) + 1$ due to Lemma 4; recall that $j^* \notin C$, so $v_{j^*} = v'_{j^*}$. Further, the set of goods allocated does not change. Hence, \mathbf{A}' remains minimally complete.

If $v_{j^*}(A_{j^*}^{\text{lie}}) + 2 \leq v'_k(A_k^{\text{lie}})$, then it can be checked that \mathbf{A}^{lie} violates the Pigou-Dalton principle due to the existence of \mathbf{A}' , which, given Lemma 3, contradicts the fact that \mathbf{A}^{lie} is an MNW allocation. Hence, we must have

$$v_{j^*}(A_{j^*}^{\text{lie}}) + 1 \geq v'_k(A_k^{\text{lie}}) = |A_k^{\text{lie}}| \geq v_k(A_k^{\text{lie}}) \geq v_k(A_k^{\text{truth}}) + 1 \geq v_i(A_i^{\text{truth}}) + 1, \quad (1)$$

where the second transition is because \mathbf{A}^{lie} is minimally complete and PO under valuation profile \mathbf{v}' , the fourth transition is because $k \in C$, and the last transition is due to the choice of i . On the other hand, we also have

$$v_{j^*}(A_{j^*}^{\text{lie}}) + 1 \leq v_{j^*}(A_{j^*}^{\text{truth}}) \leq v_i(A_i^{\text{truth}}) + 1, \quad (2)$$

where the first transition holds because, due to Lemma 8, $v_{j^*}(A_{j^*}^{\text{lie}}) \leq |A_{j^*}^{\text{lie}}| < |A_{j^*}^{\text{truth}}| = v_{j^*}(A_{j^*}^{\text{truth}})$, and the second transition holds due to Lemma 7 and the fact that $j^* \in R \subseteq S$.

Putting Equations (1) and (2) together, we have

$$v_{j^*}(A_{j^*}^{\text{lie}}) + 1 = v_{j^*}(A_{j^*}^{\text{truth}}) = v_i(A_i^{\text{truth}}) + 1 = v_k(A_k^{\text{truth}}) + 1 = v'_k(A_k^{\text{lie}}).$$

By the second equality and Lemma 7, we must have $j^* < i$. By the third equality, the fact that k and i have the same utility under $\mathbf{A}^{\text{truth}}$, and the definition of i , we have that $k \geq i$. Therefore, $k > j^*$. Then, as argued in the proof of Lemma 7, under the valuation profile \mathbf{v}' , \mathbf{A}' has the same utility profile as \mathbf{A}^{lie} , and thus, by Lemma 2, it is an MNW allocation. Further, it is lexicographically better than \mathbf{A}^{lie} under \mathbf{v}' , which contradicts the fact that $\mathbf{A}^{\text{lie}} = \text{MNW}^{\text{tie}}(\mathbf{v}')$.

Polynomial-time computability. Finally, we show that MNW^{tie} can be computed in polynomial time. Fix an instance $(\mathcal{N}, \mathcal{M}, \mathbf{v})$. Without loss of generality, suppose there are no non-valued goods. This is because if there are any non-valued goods, we can simply remove them, and run the algorithm below on the remaining instance.

Let \mathbf{u}^{tie} be the utility vector that is lexicographically greatest among the utility vectors of all MNW allocations. Our goal is to compute an allocation that achieves this utility vector. Our algorithm relies on the following important lemma.

Lemma 9. *Suppose \mathbf{A} is an MNW allocation with utility vector $\mathbf{u} = (u_1, \dots, u_n) \neq \mathbf{u}^{\text{tie}}$. Let i be smallest index such that $u_i \neq u_i^{\text{tie}}$. Then, $u_i^{\text{tie}} = u_i + 1$, and there exists $j > i$ such that there is a path from i to j in $G(\mathbf{A})$ and $u_j = u_i + 1$.*

Proof. Given two allocations \mathbf{A} and \mathbf{B} , we define the transformation graph $G(\mathbf{A}, \mathbf{B})$ similarly to Freeman *et al.* [22]. It has a vertex corresponding to each agent, and for each good g , there is a directed edge (i, j) if $g \in A_i$, $g \in B_j$, and $i \neq j$; note that this may be a multi-graph. This edge signifies that g must be passed from agent i to agent j in order to transform \mathbf{A} to \mathbf{B} . Let $\mathcal{E}_+ = \{i \in \mathcal{N} : |A_i| > |B_i|\}$ and $\mathcal{E}_- = \{i \in \mathcal{N} : |B_i| > |A_i|\}$.

Corollary 3 by Freeman *et al.* [22] establishes that edges in $G(\mathbf{A}, \mathbf{B})$ can be decomposed into a set of cycles \mathcal{C} and a set of paths \mathcal{P} such that each path begins at an agent in \mathcal{E}_+ and ends at an agent in \mathcal{E}_- . Although they do not mention this in the statement of their corollary, they in fact prove something stronger: for every agent $i \in \mathcal{E}_+$ (resp. \mathcal{E}_-), there is a path $P \in \mathcal{P}$ beginning (resp. ending) at i .

Next, we notice that the transformation graph $G(\mathbf{A}, \mathbf{B})$ is closely related to graphs $G(\mathbf{A})$ and $G(\mathbf{B})$. Specifically, if \mathbf{A} and \mathbf{B} are both PO, and there is an edge (i, j) in $G(\mathbf{A}, \mathbf{B})$, then there must be an edge (j, i) in $G(\mathbf{A})$ and an edge (i, j) in $G(\mathbf{B})$. To see this, note that if (i, j) is an edge in $G(\mathbf{A}, \mathbf{B})$, then there is a good $g \in B_j \cap A_i$. Since both allocations are PO, both i and j must like g (i.e. $g \in V_j \cap V_i$). Now, $g \in A_i \cap V_j$ implies that edge (j, i) exists in $G(\mathbf{A})$, and $g \in B_j \cap V_i$ implies that edge (i, j) exists in $G(\mathbf{B})$. Extending this argument, we get that a path from i to j in $G(\mathbf{A}, \mathbf{B})$ implies a path from j to i in $G(\mathbf{A})$ and a path from i to j in $G(\mathbf{B})$.

We let \mathbf{A}^{tie} be some MNW allocation with utility vector \mathbf{u}^{tie} , and consider $G(\mathbf{A}, \mathbf{A}^{\text{tie}})$. Since both \mathbf{A} and \mathbf{A}^{tie} are PO, and there are no non-valued goods, we have that $v_j(A_j) = |A_j|$ and $v_j(A_j^{\text{tie}}) = |A_j^{\text{tie}}|$ for all agents j .

Consider agent i defined in the lemma statement. First, note that $u_i > u_i^{\text{tie}}$ would violate lexicographic maximality of \mathbf{u}^{tie} . Hence, we must have $u_i < u_i^{\text{tie}}$, i.e., $i \in \mathcal{E}_-$. Therefore, there must exist a path P from some agent $j \in \mathcal{E}_+$ to agent i in $G(\mathbf{A}, \mathbf{A}^{\text{tie}})$. Note that this means there is a path from j to i in $G(\mathbf{A}^{\text{tie}})$, and, crucially, a path from i to j in $G(\mathbf{A})$. However, for all $j < i$, we have $u_j = u_j^{\text{tie}}$, i.e., $|A_j| = |A_j^{\text{tie}}|$, i.e., j belongs to neither \mathcal{E}_+ nor \mathcal{E}_- . Hence, we must have $j > i$.

Since both \mathbf{A} and \mathbf{A}^{tie} are MNW allocations, and there is a path from i to j in $G(\mathbf{A})$ and a path from j to i in $G(\mathbf{A}^{\text{tie}})$, by Lemma 5, we have that

$$|A_j| \leq |A_i| + 1 \text{ and } |A_i^{\text{tie}}| \leq |A_j^{\text{tie}}| + 1. \quad (3)$$

In addition, we have $i \in \mathcal{E}_-$ and $j \in \mathcal{E}_+$. Hence,

$$|A_i| < |A_i^{\text{tie}}| \text{ and } |A_j^{\text{tie}}| < |A_j|. \quad (4)$$

Combining Equations (3) and (4), we have

$$|A_i| < A_i^{\text{tie}} \leq A_j^{\text{tie}} + 1 < |A_j| + 1. \quad (5)$$

Since all values are integers, we have $|A_i| + 2 \leq |A_j| + 1$, i.e., $|A_i| + 1 \leq |A_j|$. Given Equation (3), this implies $|A_j| = |A_i| + 1$. Substituting this equality in Equation (5), we also get $|A_i^{\text{tie}}| = |A_i| + 1$, as desired. \square

Suppose \mathbf{A} is an MNW allocation with $\mathbf{u} \neq \mathbf{u}^{\text{tie}}$. Hence, Lemma 9 holds. Consider the agents i, j and path P identified in the lemma. Let $\mathbf{A}' = P(\mathbf{A})$ have utility vector \mathbf{u}' . Then, by Lemma 4, we have that $u'_j = u_j - 1 = u_i$ and $u'_i = u_i + 1 = u_j = u_i^{\text{tie}}$. Hence, it can be checked that \mathbf{A}' is an MNW allocation, and its utility vector \mathbf{u}' has a strictly longer prefix matching \mathbf{u}^{tie} than \mathbf{u} does. Consequently, if \mathbf{A} is in fact an MNW allocation with $\mathbf{u} = \mathbf{u}^{\text{tie}}$, then no such path can exist.

We are now ready to describe our algorithm. It starts by computing any MNW allocation \mathbf{A} . Barman *et al.* [6]; Darmann and Schauer [20] provide efficient algorithms for computing an MNW allocation under binary additive preferences, which can be used. Then, our algorithm iteratively finds the smallest index i

that has a path to some $j > i$ with $|A_j| = |A_i| + 1$ and passes back along such a path. By the above arguments, this must terminate in at most n iterations at an MNW allocation with lexicographically greatest utility vector, which the algorithm returns. A somewhat simpler but equivalent description of the algorithm is given as Algorithm 1.

Algorithm 1 A polynomial-time algorithm to compute MNW^{tie} for binary additive valuations

```

1: Compute an MNW allocation  $\mathbf{A}^0$ 
2: for  $i = 1, \dots, n$  do
3:   if there is a path  $P$  in  $G(\mathbf{A}^{i-1})$  from agent  $i$  to some agent  $j > i$  with  $|A_j^{i-1}| = |A_i^{i-1}| + 1$  then
4:      $\mathbf{A}^i \leftarrow P(\mathbf{A}^{i-1})$ 
5:   else
6:      $\mathbf{A}^i \leftarrow \mathbf{A}^{i-1}$ 
7:   end if
8: end for
9: return  $\mathbf{A}^n$ 

```

To formally argue correctness, we use induction on i and show that \mathbf{A}^i is an MNW allocation with $v_k(A_k^i) = u_k^{\text{tie}}$ for $k \in [i]$ (that is, its utility vector matches \mathbf{u}^{tie} in the first i components). The base case with $i = 0$ is trivial as \mathbf{A}^0 is an MNW allocation. Suppose the induction hypothesis holds for $i - 1$.

Now, if there is a path P from i to some $j > i$ in $G(\mathbf{A}^{i-1})$ as identified in Lemma 9, then we know that $v_i(A_i^{i-1}) + 1 = u_i^{\text{tie}}$. However, as argued above, passing back along this path results in an MNW allocation \mathbf{A}^i under which $v_i(A_i^i) = v_i(A_i^{i-1}) + 1 = u_i^{\text{tie}}$. Further, it does not change the utilities to any agent $k < i$. Hence, \mathbf{A}^i is an MNW allocation whose utility vector matches \mathbf{u}^{tie} in the first i components, as desired. On the other hand, if there is no such path, then by Lemma 9 and the induction hypothesis, it must be the case that $v_i(A_i^{i-1}) = u_i^{\text{tie}}$, so setting $\mathbf{A}^i \leftarrow \mathbf{A}^{i-1}$ achieves the desired goal.

To see that the running time is polynomial, first note that we can compute an arbitrary MNW allocation in polynomial time for binary additive utilities. In each iteration of the for loop, constructing $G(\mathbf{A}^{i-1})$ and searching for a desired path in this graph can also be done in polynomial time. Finally, computing an allocation by passing back along a path can be done in polynomial time. Since the for loop runs for n iterations, the overall running time is polynomial. \square

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.* [23] 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. Our positive result circumvents this impossibility for binary additive valuations. Additionally, it satisfies GSP, which Freeman *et al.* do not consider. 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* $\bar{\mathbf{A}}$ is a probability distribution over deterministic allocations.

There is a natural fractional allocation \mathbf{A} associated with each randomized allocation $\bar{\mathbf{A}}$, where $A_i(g)$ is the probability of good g being allocated to agent i under $\bar{\mathbf{A}}$. In this case, we say that randomized allocation $\bar{\mathbf{A}}$ *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 $\bar{\mathbf{A}}$ as simply the utility of agent i under $\bar{\mathbf{A}}$. 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.⁵ We say that a randomized allocation $\bar{\mathbf{A}}$ 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 $\bar{\mathbf{A}}$. 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 $\bar{\mathbf{A}}$ 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 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; 29; 32]. Our first result shows that under binary additive valuations, these two fundamental 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.*

Proof. We begin by showing that there exists a utility vector \mathbf{u}^{lex} (resp. \mathbf{u}^{mnw}) such that the set of fractional leximin allocations (resp. fractional MNW allocations) is exactly the set of all fractional allocations with

⁵In case of Pareto optimality of a fractional allocation, we require that no other *fractional* allocation Pareto-dominate it.

utility vector \mathbf{u}^{lex} (resp. \mathbf{u}^{mnw}). In fact, this step holds even under general additive valuations. Then, for binary additive valuations, we will show that $\mathbf{u}^{\text{lex}} = \mathbf{u}^{\text{mnw}}$, implying the desired result.

Let \mathbf{A} be a fractional leximin allocation with utility vector \mathbf{u} . Trivially, every fractional allocation with utility vector \mathbf{u} is also a fractional leximin allocation. We want to show that there is no fractional leximin allocation \mathbf{A}' with utility vector $\mathbf{u}' \neq \mathbf{u}$. Suppose for contradiction that there is one. Consider the fractional allocation $\mathbf{A}'' = 1/2 \cdot \mathbf{A} + 1/2 \cdot \mathbf{A}'$, i.e., $A''_i(g) = 1/2 \cdot A_i(g) + 1/2 \cdot A'_i(g)$ for each agent i and good g . Because valuations are additive, its utility vector is $\mathbf{u}'' = 1/2 \cdot \mathbf{u} + 1/2 \cdot \mathbf{u}'$.

The key step to observe is that if \mathbf{u} and \mathbf{u}' are utility vectors of two fractional leximin allocations, and $\mathbf{u} \neq \mathbf{u}'$, then \mathbf{u}'' is strictly better than both \mathbf{u} and \mathbf{u}' according to leximin comparison, which yields the desired contradiction. To see why this is true, note that because both \mathbf{A} and \mathbf{A}' are fractional leximin allocations, their utility profiles must be identical; call it \mathbf{s}^* . Then, for any $k \in [n]$, the sum of the k lowest utilities under \mathbf{A}'' is the average of the sum of utilities of the corresponding k agents under \mathbf{A} and \mathbf{A}' . Since each sum is at least the sum of the first k components of \mathbf{s}^* , it follows that the sum of the k lowest utilities under \mathbf{u}'' is at least as much as the sum of the k lowest utilities under \mathbf{u} or \mathbf{u}' , i.e., \mathbf{u}'' is at least as good as \mathbf{u} and \mathbf{u}' under leximin comparison. To see why it is strictly better, recall that $\mathbf{u} \neq \mathbf{u}'$. Let agent i be such that $u_i \neq u'_i$, and among such agents, one with the lowest u''_i . Without loss of generality, assume $u_i < u'_i$. Then, $u_i < u''_i < u'_i$. Let $N = \{j \in \mathcal{N} : u''_j < u''_i\}$ and $k = |N|$. By the definition of N , for each $j \in N$, we have $u_j = u'_j = u''_j$. Hence, the k smallest values in \mathbf{u}'' also appear in \mathbf{u} and the $k + 1^{\text{st}}$ smallest value in \mathbf{u}'' (which is u''_i) is strictly higher than the $k + 1^{\text{st}}$ smallest value in \mathbf{u} (which is less than u''_i), which shows that \mathbf{u}'' is strictly better than \mathbf{u} under leximin comparison. The comparison to \mathbf{u}' follows since \mathbf{u} and \mathbf{u}' have identical sorted order, and thus are equivalent under leximin comparison.

Thus, we have established that there exists a utility vector, say \mathbf{u}^{lex} , such that the set of fractional leximin allocations is the set of allocations with utility vector \mathbf{u}^{lex} . It is easy to see that the above argument holds for fractional MNW allocations as well. Crucially, the key step in the paragraph above holds because the MNW objective function (product of utilities of agents who like at least one good) is a strictly concave function. Hence, if $\mathbf{u} \neq \mathbf{u}'$ have equal objective value, then \mathbf{u}'' has a strictly better objective value than both of them. Let \mathbf{u}^{mnw} denote the utility vector for fractional MNW allocations.

Finally, we need to show that $\mathbf{u}^{\text{lex}} = \mathbf{u}^{\text{mnw}}$. Fix arbitrary fractional leximin and fractional MNW allocations \mathbf{A}^{lex} and \mathbf{A}^{mnw} . Suppose this is not true. Let \mathbf{w}^{lex} and \mathbf{w}^{mnw} be the utility profiles corresponding to \mathbf{u}^{lex} and \mathbf{u}^{mnw} , respectively. Let k be the smallest index such that $w_k^{\text{lex}} \neq w_k^{\text{mnw}}$. Because \mathbf{w}^{lex} is the leximin-optimal utility profile, we must have $w_k^{\text{lex}} > w_k^{\text{mnw}}$. Because fractional leximin and fractional MNW allocations are PO, they have identical sum of utilities. Hence, there exists t such that $w_t^{\text{lex}} < w_t^{\text{mnw}}$. Choose the smallest such t . Then, we have that for each $k < t$, $w_k^{\text{lex}} \geq w_k^{\text{mnw}}$, and for at least one $k < t$, $w_k^{\text{lex}} > w_k^{\text{mnw}}$. Thus, $\sum_{k=1}^{t-1} w_k^{\text{lex}} > \sum_{k=1}^{t-1} w_k^{\text{mnw}}$. It is also worth noting that $w_{t-1}^{\text{mnw}} \leq w_{t-1}^{\text{lex}} \leq w_t^{\text{lex}} < w_t^{\text{mnw}}$.

Let N denote the set of agents with the $t - 1$ lowest utilities under \mathbf{u}^{mnw} . Since the collective utility these agents receive is at the minimum the sum of the first $t - 1$ values of \mathbf{w}^{lex} , they receive strictly higher total utility under \mathbf{u}^{lex} than under \mathbf{u}^{mnw} . Thus, there must exist an agent $i \in N$ and a good $g \in V_i$ such that a positive fraction of g is allocated to an agent $j \notin N$ under \mathbf{A}^{mnw} . Since $u_j^{\text{mnw}} \geq w_t^{\text{mnw}} > w_{t-1}^{\text{mnw}} \geq u_i^{\text{mnw}}$, it follows that transferring a small enough fraction of good g from agent j to agent i in \mathbf{A}^{mnw} strictly improves the Nash welfare, which is a contradiction. Hence, $\mathbf{u}^{\text{lex}} = \mathbf{u}^{\text{mnw}}$. Consequently, the set of fractional leximin and fractional MNW allocations coincide, and these allocations have identical utility vectors. \square

Note that the identical utility vector guarantee in Theorem 2 is much stronger than the identical utility profile guarantee in the deterministic case (Lemmas 1 and 2).

Even under general additive valuations, it is known that every fractional MNW allocation is ex ante EF and ex ante PO [32], and one such allocation can be computed in strongly polynomial time [29; 33]. 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 below, 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}). The proof of the next result is straightforward.

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.*

Proof. We show that binary additive valuations satisfy the four requirements laid out by Kurokawa *et al.* [25] in Section 3.2 of their paper. Using their notation, \mathcal{A} denotes the set of feasible allocations, \mathcal{P} denotes the set of possible preferences (i.e. weak order over \mathcal{A}) that agents may have, and \mathcal{U} denotes the set of valuation functions that the agents may have. In our setting, \mathcal{A} is the set of all fractional allocations, and each preference in \mathcal{P} has a unique valuation function in \mathcal{U} consistent with it, just the natural one which assigns value 1 to each good g for which the agent strictly prefers the allocation $\{g\}$ to \emptyset , and 0 to every other good. We now specify the four requirements laid out by Kurokawa *et al.* [25] and argue why they are satisfied in our domain.

1. *Convexity.* Given two feasible allocations $\mathbf{A}, \mathbf{A}' \in \mathcal{A}$ and $\lambda \in [0, 1]$, we need to show there exists an allocation $\mathbf{A}'' \in \mathcal{A}$ such that $u_i(A_i'') = \lambda u_i(A_i) + (1 - \lambda)u_i(A_i')$ for each agent i . We can simply let \mathbf{A}'' to be the fractional allocation induced by the randomized allocation that selects \mathbf{A} with probability λ and \mathbf{A}' with probability $1 - \lambda$.
2. *Equality.* Kurokawa *et al.* [25] only use this property to achieve a guarantee known as proportionality. While in our domain envy-freeness implies proportionality under a complete allocation, their result requires the equality requirement to guarantee proportionality even in domains where it is not implied by envy-freeness. Hence, we do not need to show this requirement in our domain.
3. *Shifting Allocations.* Given an allocation $\mathbf{A} \in \mathcal{A}$ and agents $i, j \in \mathcal{N}$, we need to show there exists an allocation $\mathbf{A}' \in \mathcal{A}$ such that $v_k(A'_k) = v_k(A_k)$ for all agents $k \in \mathcal{N} \setminus \{i, j\}$ and $v_i(A'_i) \geq v_i(A_j)$. For this, we can choose \mathbf{A}' such that $A'_k = A_k$ for all $k \in \mathcal{N} \setminus \{i, j\}$, $A'_i = A_i \cup A_j$, and $A'_j = \emptyset$.
4. *Optimal Utilization.* This requires that if $\mathbf{A} \in \mathcal{A}$ is returned by the fractional MNW rule in a given instance, then for any valuation function $v \in \mathcal{U}$, $v_i(A_i) \geq v(A_i)$. Crucially, because our rule outputs a minimally complete fractional MNW allocation, agent i must like each good that she is assigned a positive fraction of. Hence, $v_i(A_i) = \sum_{g \in \mathcal{M}} A_i(g) \cdot 1 \geq v(A_i)$ for any binary additive valuation function v .

Hence, it follows from the result of Kurokawa *et al.* [25] that fractional MNW rule is ex ante envy-free, ex ante Pareto optimal, and ex ante group strategyproof. \square

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.* [14], which is a generalization of the classic Birkhoff-von Neumann theorem [8; 34]. 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 .

$$\begin{aligned} \mathcal{H}_1 : \sum_{i \in \mathcal{N}} A_i(g) &= \sum_{i \in \mathcal{N}} A_i^*(g), \forall g \in \mathcal{M}, \\ \mathcal{H}_2 : \lfloor u_i^* \rfloor &\leq \sum_{g \in \mathcal{M}} A_i(g) \cdot v_i(g) \leq \lceil u_i^* \rceil, \forall i \in \mathcal{N}. \end{aligned} \quad (6)$$

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.* [14]. 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 Equation (6). 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 $\bar{\mathbf{A}}$ be the randomized allocation implementing \mathbf{A}^* that is returned by the polynomial-time algorithm of Budish *et al.* [14] with the bihierarchy constraints in Equation (6). Let \mathcal{A} denote the set of deterministic allocations in the support of $\bar{\mathbf{A}}$. 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} ,

$$\sum_{r \in [k]} \sum_{i \in S_r} u_i \leq \sum_{r \in [k]} \sum_{i \in S_r} u_i^*, \forall k \in [t]. \quad (7)$$

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} , Equation (7) 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 (Equation (6)).

$$\begin{aligned}\mathcal{H}_2 &: \lfloor u_i^* \rfloor \leq u_i \leq \lceil u_i^* \rceil, \forall i \in \mathcal{N}, \\ \mathcal{H}_3 &: \sum_{i \in S_k} u_i = \sum_{i \in S_k} u_i^*, \forall k \in [t].\end{aligned}\tag{8}$$

We say that a utility vector is a *rounded* if it satisfies the constraints in Equation (8), 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.

Lemma 10. *The set of rounded allocations coincides with the set of maximum Nash welfare allocations.*

Proof. Because leximin and MNW are equivalent concepts for deterministic allocations (Lemma 3), we will refer to MNW allocations as leximin allocations in this proof. To establish the desired result, it is sufficient to show that given a rounded allocation $\mathbf{A}^{\text{round}}$ and an arbitrary allocation \mathbf{A} , $\mathbf{A}^{\text{round}}$ weakly leximin-dominates \mathbf{A} , and if \mathbf{A} is not rounded, then $\mathbf{A}^{\text{round}}$ strictly leximin-dominates \mathbf{A} . Let $\mathbf{u}^{\text{round}}$ and \mathbf{u} be the utility vectors of $\mathbf{A}^{\text{round}}$ and \mathbf{A} , respectively.

First, assume that \mathbf{A} is rounded. Then, both $\mathbf{u}^{\text{round}}$ and \mathbf{u} satisfy Equation (8). However, it is easy to see that any two utility vectors satisfying Equation (8) induce the same utility profile, and thus $\mathbf{u}^{\text{round}}$ trivially weakly leximin-dominates \mathbf{u} . To see this, note that for each $k \in [t]$, the sum of utilities of agents in S_k is fixed due to \mathcal{H}_3 . And further, for each agent $i \in S_k$, \mathcal{H}_2 implies that either $u_i = u_i^* = L_k$ if $u_i^* = L_k$, or $u_i \in \{L_k, L_k + 1\}$ if $u_i^* \in (L_k, L_k + 1)$. Thus, together, \mathcal{H}_2 and \mathcal{H}_3 fix the number of agents in S_k that have utility L_k and those that have utility $L_k + 1$. Thus, any two utility vectors satisfying Equation (8) induce identical utility vectors.

Next, assume that \mathbf{A} is not rounded, i.e., \mathbf{u} violates either \mathcal{H}_2 or \mathcal{H}_3 . Let $k \in [t]$ be the smallest index such that either \mathcal{H}_2 is violated for some agent $i \in S_k$, or \mathcal{H}_3 is violated for S_k . Then, by the above argument, the partial utility vectors $(u_i^{\text{round}})_{i \in \cup_{r < k} S_r}$ and $(u_i)_{i \in \cup_{r < k} S_r}$ induce identical utility profiles, and therefore, are leximin-equivalent.

Suppose \mathcal{H}_3 is violated for S_k . Then, because \mathbf{u} satisfies Equation (7), and \mathbf{u} and $\mathbf{u}^{\text{round}}$ match on the total utility of agents in S_r for $r < k$, we have $\sum_{i \in S_k} u_i < \sum_{i \in S_k} u_i^* = \sum_{i \in S_k} u_i^{\text{round}}$. Because $(u_i^{\text{round}})_{i \in S_k}$ has a higher sum than $(u_i)_{i \in S_k}$, and because it distributes that higher sum as equally as possible, $(u_i^{\text{round}})_{i \in S_k}$ strictly leximin-dominates $(u_i)_{i \in S_k}$. Hence, $(u_i^{\text{round}})_{i \in \cup_{r \leq k} S_r}$ also strictly leximin-dominates $(u_i)_{i \in \cup_{r \leq k} S_r}$. To argue that $\mathbf{u}^{\text{round}}$ strictly leximin-dominates \mathbf{u} , we need to argue that adding the remaining utilities does not change the comparison. To that end, note that for any $r' > k$, $r \leq k$, $i' \in S_{r'}$ and $i \in S_r$, we have $u_{i'}^{\text{round}} \geq L_{r'} \geq U_r \geq u_i^{\text{round}}$. Thus, because we are only adding utilities to $\mathbf{u}^{\text{round}}$ that are at least as high as the ones already added, the strict leximin-dominance continues to hold.

Next, \mathcal{H}_2 is violated for some agent $i \in S_k$. Then, $u_i \notin \{\lfloor u_i^* \rfloor, \lceil u_i^* \rceil\}$. By the above argument, $(u_i^{\text{round}})_{i \in \cup_{r < k} S_r}$ and $(u_i)_{i \in \cup_{r < k} S_r}$ induce the same utility profile. If $u_i < \lfloor u_i^* \rfloor$, then \mathbf{u} has strictly more agents with utility less than L_k , which implies that it is strictly leximin-dominated by $\mathbf{u}^{\text{round}}$. If $u_i > \lceil u_i^* \rceil$, then \mathbf{u} has either an agent $j \in S_k$ with $u_j < \lfloor u_j^* \rfloor = L_k$ or more agents with utility L_k than $\mathbf{u}^{\text{round}}$ does. In either case, it is again easy to see that $\mathbf{u}^{\text{round}}$ strictly leximin-dominates \mathbf{u} . \square

This completes the proof of Theorem 4. \square

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 appendix, 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] 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.
- [2] G. Amanatidis, G. Birmpas, G. Christodoulou, and E. Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 545–562, 2017.
- [3] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, A. Hollender, and A. A. Voudouris. Maximum Nash welfare and other stories about EFX. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020. Forthcoming.
- [4] H. Aziz and S. Rey. Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [5] M. Babaioff, T. Ezra, and U. Feige. Fair and truthful mechanisms for dichotomous valuations. arXiv preprint arXiv:2002.10704, 2020.
- [6] S. Barman, S. K. Krishnamurthy, and R. Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 7–13, 2018.
- [7] N. Benabbou, M. Chakraborty, A. Igarashi, and Y. Zick. Finding fair and efficient allocations when valuations don’t add up. arXiv preprint arXiv:2003.07060, 2020.
- [8] G. Birkhoff. Three observations on linear algebra. *Universidad Nacional de Tucumán, Revista A*, 5:147–151, 1946.

- [9] A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72:257–279, 2004.
- [10] A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- [11] S. Bouveret and M. Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2016.
- [12] S. J. Brams and P. C. Fishburn. *Approval Voting*. Springer, 2nd edition, 2007.
- [13] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [14] 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.
- [15] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [16] 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*, 7(3): article 12, 2019.
- [17] Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. Group fairness in committee selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 263–279, 2019.
- [18] R. Cole and V. Gkatzelis. Approximating the Nash social welfare with indivisible items. *SIAM Journal on Computing*, 47(3):1211–1236, 2018.
- [19] R. Cole, V. Gkatzelis, and G. Goel. Mechanism design for fair division: Allocating divisible items without payments. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 251–268, 2013.
- [20] A. Darmann and J. Schauer. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2):548–559, 2015.
- [21] D. K. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- [22] R. Freeman, S. Sikdar, R. Vaish, and L. Xia. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2019.
- [23] R. Freeman, N. Shah, and R. Vaish. 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.
- [24] B. Klaus and E. Miyagawa. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, 30:421–435, 2001.
- [25] D. Kurokawa, A. D. Procaccia, and N. Shah. Leximin allocations in the real world. *ACM Transactions on Economics and Computation*, 6(3–4): article 11, 2018.
- [26] M. Lackner and P. Skowron. A quantitative analysis of multi-winner rules. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 407–413, 2019.

- [27] 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 Economics and Computation (EC)*, pages 125–131, 2004.
- [28] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- [29] J. B. Orlin. 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)*, pages 291–300, 2010.
- [30] J. Ortega. Multi-unit assignment under dichotomous preferences. *Mathematical Social Sciences*, 103:15–24, 2020.
- [31] B. Plaut and T. Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2584–2603, 2018.
- [32] H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- [33] L. A. Végh. Concave generalized flows with applications to market equilibria. *Mathematics of Operations Research*, 39(2):573–596, 2013.
- [34] 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.

Appendix

A Example Illustrating MNW^{tie}

The following examples illustrate how our deterministic rule MNW^{tie} works.

Example 1. Let us denote a valuation profile by a matrix, where n rows represent agents $1, \dots, n$, m columns represent goods g_1, \dots, g_m , and the entry in row i and column j is $v_i(g_j)$. Consider the valuation profile

$$\mathbf{v} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}.$$

In this case, the unique allocation \mathbf{A} that MNW^{tie} can return is given by $A_1 = \{g_1\}$ and $A_2 = \emptyset$. This is because g_2 cannot be allocated as it is non-valued, and MNW^{tie} must prefer the allocation which gives g_1 to agent 1 over the one which gives it to agent 2.

Generally, however, there could be multiple allocations that MNW^{tie} can arbitrarily choose from. For example, consider the valuation profile

$$\mathbf{v} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}.$$

MNW^{tie} may return any allocation which gives two goods to agent 1 and one good to agent 2.

B Allocating Non-valued Goods

Theorem 5. *No deterministic rule can always output an MNW allocation, always allocate all goods, and be SP for two agents and six goods.*

Proof. Suppose there are two agents $\mathcal{N} = \{1, 2\}$ and six goods $\mathcal{M} = \{g_1, g_2, g_3, g_4, g_5, g_6\}$. Suppose for a contradiction there existed such a rule F . Consider the valuation profile \mathbf{v}^1 :

$$\begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The only possible MNW allocations are those where one agent gets good 1 and the other gets good 2. Once this condition is met, any allocation of the non-valued goods is an MNW allocation. Since there are four such goods, there must be an agent that receives at least two of them. Without loss of generality, agent 1 receives goods 3 and 4 along with one of items 1 and 2, say item 1. That is, $F(\mathbf{v}^1) = \mathbf{A}^1$ such that $\{g_1, g_3, g_4\} \subseteq A_1^1$ (it is possible they received more of the nonvalued ones).

Now consider another valuation profile \mathbf{v}^2 :

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The only possible MNW allocations are those where agent 1 receives goods 3 and 4 and agent 2 receives goods 1 and 2. Therefore, regardless of the allocation chosen, the utility to agent 1 is exactly 2, that is if $F(\mathbf{v}^2) = \mathbf{A}^2$, then $v_1^2(\mathbf{A}_1^2) = 2$. However, if agent 1 misreports to match \mathbf{v}_1^1 , then their utility $v_1^2(\mathbf{A}_1^1) \geq 3$, as agent 1 likes all three of g_1, g_3 , and g_4 . Therefore, F is not SP, a contradiction. \square

We wrote an integer linear program (ILP) to check whether there exists a full allocation for every possible binary additive valuation profile with two agents and six goods such that the resulting rule is EF1, PO, and SP. We solved the ILP using CPLEX, and determined that there indeed exists such a rule. Our program does not terminate in reasonable time when run on two agents and seven goods, so whether there exists a deterministic rule that allocates all goods and is EF1, PO, and SP is an open question for $n = 2$ and $m \geq 7$.