On Fair and Efficient Allocations of Indivisible Public Goods

Jugal Garg ✉
University of Illinois, Urbana-Champaign, IL, USA

Pooja Kulkarni ✉
University of Illinois, Urbana-Champaign, IL, USA

Aniket Murhekar ✉
University of Illinois, Urbana-Champaign, IL, USA

Abstract

We study fair allocation of indivisible public goods subject to cardinality (budget) constraints. In this model, we have n agents and m available public goods, and we want to select $k \leq m$ goods in a fair and efficient manner. We first establish fundamental connections between the models of private goods, public goods, and public decision making by presenting polynomial-time reductions for the popular solution concepts of maximum Nash welfare (MNW) and leximin. These mechanisms are known to provide remarkable fairness and efficiency guarantees in private goods and public decision making settings. We show that they retain these desirable properties even in the public goods case. We prove that MNW allocations provide fairness guarantees of Proportionality up to one good (Prop1), $1/n$ approximation to Round Robin Share (RRS), and the efficiency guarantee of Pareto Optimality (PO). Further, we show that the problems of finding MNW or leximin-optimal allocations are NP-hard, even in the case of constantly many agents, or binary valuations. This is in sharp contrast to the private goods setting that admits polynomial-time algorithms under binary valuations. We also design pseudo-polynomial time algorithms for computing an exact MNW or leximin-optimal allocation for the cases of (i) constantly many agents, and (ii) constantly many goods with additive valuations. We also present an $O(n)$-factor approximation algorithm for MNW which also satisfies RRS, Prop1, and 1/2-Prop.

2012 ACM Subject Classification Theory of computation → Mathematical optimization

Keywords and phrases Public goods, Nash welfare, Leximin, Proportionality

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2021.22

Related Version *Full Version*: https://arxiv.org/abs/2107.09871 [23]

Funding Supported by NSF Grant CCF-1942321 (CAREER).

1 Introduction

The problem of fair division was formally introduced by Steinhaus [32], and has since been extensively studied in economics and computer science [10, 28]. Recent work has focused on the problem of fair and efficient allocation of indivisible *private* goods. We label this setting as the PrivateGoods model. Here, goods have to be partitioned among agents, and a good provides utility only to the agent who owns it. However, goods are not always private, and may provide utility to multiple agents simultaneously, e.g., books in a public library. The fair and efficient allocation of such *indivisible public goods* is an important problem.

In this paper we study the setting of PublicGoods, where a set of n agents have to select a set of at most k goods from a set of m given goods. This simple cardinality constraint models several real world scenarios. While previous work has largely focused on the $k < n$ case, e.g., for voting and committee selection [2, 13], there is much less work available for the case of $k \geq n$. This setting is important in its own right. We present a few compelling examples.

41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021).
Editors: Mikołaj Bojańczyk and Chandra Chekuri; Article No. 22; pp. 22:1–22:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

▶ **Example 1.** A *public library* wants to buy k books that adhere to preferences of n people who might use the library. Clearly, the number of books has to be much greater than the number of people using the library, hence $k \gg n$.

▶ **Example 2.** A family (or a group of friends) of size n wants to decide on a list of k movies to watch together for a few months. Here too, $k > n$. Another example of the same flavor is a committee tasked with inviting speakers at a year-long weekly seminar.

▶ **Example 3.** Another important example is that of *diverse search results* for a query. Given a query (say of "computer scientist images") on a database, we would like to output k search results which reflect diversity in terms of n specified features (like "gender, race and nationality"). Once again, $k \geq n$.

A related setting PublicDecisions of public decision making [15] models the scenario in which n agents are faced with m issues with multiple alternatives per issue, and they must arrive at a decision on each issue. Conitzer et al. [15] showed that this model subsumes the PrivateGoods setting.

Connections between the models. A central question motivating this work is:

▶ **Question 1.** *Can we establish fundamental connections between the three models* PrivateGoods*,* PublicGoods*, and* PublicDecisions*?*

To answer this question, we first describe two well-studied solution concepts for allocating goods in the PrivateGoods and PublicDecisions models, namely the *maximum Nash welfare* (MNW) and *leximin* mechanisms. These mechanisms have been shown to produce allocations that are fair and efficient in the models of PrivateGoods and PublicDecisions. The MNW mechanism returns an allocation that maximizes the geometric mean of agents' utilities, and the leximin mechanism returns an allocation that maximizes the minimum utility, and subject to this, maximizes the second minimum utility, and so on. We label the problems of computing the Nash welfare maximizing (resp. leximin optimal) allocation in the three models as PrivateMNW, PublicMNW, DecisionMNW (resp. PrivateLex, PublicLex, DecisionLex).

We answer Question 1 positively by presenting novel polynomial-time reductions from the model of PrivateGoods to PublicGoods, and from PublicGoods to PublicDecisions for the problem of computing a Nash welfare maximizing allocation.

$$\boxed{\text{PrivateMNW} \leq \text{PublicMNW} \leq \text{DecisionMNW}} \tag{1}$$

More notably, these reductions also work for the MNW problem when restricted to binary valuations. Apart from establishing fundamental connections between these models, our reductions also determine the complexity of the MNW problem, as we detail below. We also develop similar reductions between the models for the leximin mechanism, showing:

$$\boxed{\text{PrivateLex} \leq \text{PublicLex} \leq \text{DecisionLex}} \tag{2}$$

Fairness and efficiency considerations. We next describe the fairness and efficiency properties that the MNW and leximin mechanisms have been shown to satisfy in the PrivateGoods and PublicDecisions models.

The standard notion of economic efficiency is Pareto-optimality (PO). An allocation is said to be PO if no other allocation makes an agent better off without making anyone worse off. The classical fairness notion of *proportionality* requires that every agent gets her *proportional value*, i.e., $1/n$-fraction of the maximum value she can obtain in any allocation. However, proportional allocations are not guaranteed to exist.[1] Hence, we study the notion of Proportionality up to one good (Prop1) for PublicGoods. We say an allocation is Prop1 if for every agent i who does not get her proportional value, i gets her proportional value after swapping some unselected good with a selected one. For PrivateGoods and PublicDecisions, Prop1 is defined similarly – in the former, an agent is given an additional good [6, 27]; and in the latter, an agent is allowed to change the decision on a single issue [15]. While Prop1 is an individual fairness notion, it is still important for allocating public goods. For instance, in Example 1, we want allocations in which every agent has some books that cater to her taste, even if her taste differs from the rest of the agents. Likewise, in Example 2, a fair selection of movies must ensure that there are some movies every member can enjoy.

We also consider the fairness notion of Round-Robin Share (RRS) [15], which demands that each agent i receives at least the utility which she would get if agents were allowed to pick goods in a round-robin fashion, with i picking last.

In the PrivateGoods and PublicDecisions models, an MNW allocation satisfies Prop1 in conjunction with PO [11, 15]. Similarly in both these models, the leximin-optimal allocation satisfies RRS and PO [15]. It is therefore natural to ask:

▶ **Question 2.** *What guarantee of fairness and efficiency do the MNW and leximin mechanisms provide in the* PublicGoods *model?*

Answering this question, we show that an MNW allocation satisfies Prop1, $1/n$-approximation to RRS, and is PO. Further, for all agents, a leximin-optimal allocation satisfies RRS[2], Prop1 and PO.

Complexity of computing MNW and leximin-optimal allocations. Given the desirable fairness and efficiency properties of these mechanisms, we investigate the complexity of computing MNW and leximin-optimal allocations in the PublicGoods model. It is known that PrivateMNW is APX-hard [26, 21] (hard to approximate) and DecisionMNW [15] is NP-hard. Likewise, PrivateLex too is NP-hard [9]. Therefore, we ask:

▶ **Question 3.** *What is the complexity of* PublicMNW *and* PublicLex*?*

Since PrivateMNW and PrivateLex are known to be NP-hard, our reductions (1) and (2) immediately show that PublicMNW and PublicLex are NP-hard. However, we show stronger results that PublicMNW and PublicLex remain NP-hard even when the valuations are binary. These results are in stark contrast to the PrivateGoods case, which admits polynomial-time algorithms for binary valuations [16, 20]. Further, our reductions between PublicGoods and PublicDecisions also directly enable us to show NP-hardness of DecisionMNW and DecisionLex. Note that the hardness of these problems is known through the connection between PrivateMNW (PrivateLex) and DecisionMNW (DecisionLex) [15]. However, a feature of our reductions (Observation 8) enables us to shows that DecisionMNW is NP-hard even for binary valuations, highlighting the utility of our reductions. We also show that PublicMNW and PublicLex remain NP-hard even when there are only two agents. We note that for the case

[1] Consider for example, two agents A and B and six public goods $\{g_1, g_2, g_3, g_4, g_5, g_6\}$. Agent A has value 1 for g_1, g_2, g_3 and B has value 1 for g_4, g_5, g_6. All other valuations are 0. Suppose we want to select three of these goods. The proportional share of both agents is 1.5. However, in any allocation, the value of at least one agent is at most 1, implying that proportional allocations need not exist.

[2] Note that here we assume we scale the valuations so that RRS $= 1$ for every agent.

Table 1 Complexity of computing MNW and leximin-optimal allocations.

Problem	PrivateGoods	PublicGoods	PublicDecisions
MNW $\{0,1\}$ valuations	P [8, 16]	NP-hard (Theorem 16)	NP-hard (Corollary 23)
Leximin $\{0,1\}$ valuations	P [8, 16]	NP-hard (Theorem 21)	?
MNW two agents	NP-hard	NP-hard (Theorem 20)	?
Leximin two agents	NP-hard	NP-hard (Theorem 22)	?

of two agents, the NP-hardness of PrivateMNW and PrivateLex does not imply NP-hardness of PublicMNW and PublicLex because our reductions between the models do not preserve the number of agents.

We summarize our results in Table 1.

In light of the above computational hardness, we turn to approximation algorithms and exact algorithms for special cases. We design a polynomial-time algorithm that returns an allocation which approximates the MNW to a $O(n)$-factor when $k \geq n$, and is also Prop1 and satisfies RRS.

Finally, we obtain pseudo-polynomial time algorithms for computing MNW and leximin-optimal allocations for constant n. These are essentially tight in light of the NP-hardness for constant n. In interest of space, we skip some proofs from this version. All these proofs can be found in full version of the paper [23].

1.1 Other related work

Maximum Nash welfare. The problem of approximating maximum Nash welfare for private goods is well-studied, see e.g., [14, 7, 12, 22]. [18] showed that the MNW problem is NP-hard for allocating public goods subject to matroid or packing constraints. It has also been studied in the context of voting, or multi-winner elections [1]. Fluschnik et al. [19] studied the fair multi-agent knapsack problem, wherein each good has an associated budget, and a set of goods is to be selected subject to a budget constraint. In this context, they studied the objective of maximizing the geometric mean of $(1 + u_i)$ where u_i is the utility of the i^{th} agent. They showed that maximizing this objective is NP-hard, even for binary valuations or constantly many agents with equal budgets and presented a pseudo-polynomial time algorithm for constant n.

Leximin. Leximin was developed as a fairness notion in itself [30]. Plaut and Roughgarden [29] showed that for private goods, leximin can be used to construct allocations that are envy-free up to any good. Freeman et al. [20] showed that in the PrivateGoods model the MNW and leximin-optimal allocations coincide when valuations are binary.

Core. Core is a strong property that enforces both PO and proportionality-like fairness guarantees for all subsets of agents. It is well-studied in many settings, including game theory and computer science [31, 25]. The core of indivisible public goods might be empty. Fain et al. [18] proved that under matroid constraints, a 2-additive approximation to core exists. On an individual fairness level, 1-additive core is weaker than Prop1 [18].

Participatory Budgeting. The participatory budgeting problem [3, 4] consists of a set of n agents (or voters), and a set of k projects that require funds, a total available budget, and the preferences of the voters over the projects. The problem is to allocate the budget in a fair and efficient manner. Here typically $k \ll n$. Fain et al. [17] showed that the fractional core outcome is polynomial-time computable. This could be modeled as a public goods problem with goods as the projects.

Voting and Committee Selection. These settings involve selecting a set of k members from a set of m candidates based on the preferences of n agents. Usually, here $k \ll n$ and the fairness notions studied are group fairness like Justified Representation [2], and a core-like notion called *stability* [13].

2 Notation and Preliminaries

Problem setting. For $t \in \mathbb{N}$, let $[t]$ denote $\{1, \ldots, t\}$. An instance of the $\mathsf{PublicGoods}$ allocation problem is given by a tuple $\mathcal{I} = (\mathcal{A}, \mathcal{G}, k, \{v_i\}_{i\in\mathcal{A}})$ of a set $\mathcal{A} = [n]$ of $n \in \mathbb{N}$ agents, a set $\mathcal{G} = [m]$ of $m \in \mathbb{N}$ public goods, an integer $0 \le k \le m$, and a set of valuation functions $\{v_i\}_{i\in\mathcal{A}}$, one per agent, where each $v_i : 2^{\mathcal{G}} \to \mathbb{Z}_{\ge 0}$. Unless specified, we assume that $k \ge n$. For a subset of goods $S \subseteq \mathcal{G}$, $v_i(S)$ denotes the utility agent i derives from the goods in S. Unless specified, we assume the valuations are *additive*. In this case, each v_i is specified by m non-negative integers $\{v_{ij}\}_{j\in\mathcal{G}}$, where v_{ij} denotes the value of agent i for good j. Then for $S \subseteq \mathcal{G}$, $v_i(S) = \sum_{j\in S} v_{ij}$. We assume without loss of generality that for every agent i, there is at least one good j with $v_{ij} > 0$. For brevity, we write $v_i(g_1, \ldots, g_r)$ in place of $v_i(\{g_1, \ldots, g_r\})$ for a set $\{g_1, \ldots, g_r\} \subseteq \mathcal{G}$. An *allocation* is a subset $\mathbf{x} \subseteq \mathcal{G}$ of goods which satisfies the cardinality constraint $|\mathbf{x}| \le k$.

Nash welfare. The Nash welfare (NW) of an allocation $\mathbf{x}$ is given by $\mathsf{NW}(\mathbf{x}) = (\prod_{i\in\mathcal{A}} v_i(\mathbf{x}))^{1/n}$. An allocation with the maximum NW is called an MNW allocation or a Nash optimal allocation.[3] We also refer to the product of the agents' utilities as the *Nash product*. An allocation $\mathbf{x}$ *approximates* MNW to a factor of α if $\mathsf{NW}(\mathbf{x}) \ge \alpha \cdot \mathsf{NW}(\mathbf{x}^*)$, where $\mathbf{x}^*$ is an MNW allocation.

Leximin. Given an allocation $\mathbf{x}$, let $\hat{\mathbf{x}}$ denote the vector of agent's utilities under $\mathbf{x}$, sorted in non-decreasing order. For two allocations $\mathbf{x}, \mathbf{y}$, we say $\mathbf{x}$ *leximin-dominates* $\mathbf{y}$ if there exists $i \in [n]$ such that $\hat{\mathbf{x}}_i > \hat{\mathbf{y}}_i$ and $\forall j < i, \hat{\mathbf{x}}_j = \hat{\mathbf{y}}_j$. An allocation is leximin-optimal if no other allocation leximin-dominates it.

Fairness notions. We now discuss fairness notions for the $\mathsf{PublicGoods}$ setting. The *proportional share* of an agent i, denoted by Prop_i is a $1/n$-share of the maximum value she can obtain from any allocation. Formally:

$$\mathsf{Prop}_i = \frac{1}{n} \cdot \max_{\mathbf{x}\subseteq\mathcal{G}, |\mathbf{x}|\le k} v_i(\mathbf{x}).$$

The round-robin share of agent i, denoted by RRS_i, is the minimum value an agent can be guaranteed if the agents pick k goods in a round-robin fashion, with i picking last. Therefore, this value equals the maximum value of any $\lfloor k/n \rfloor$ sized subset. Formally:

$$\mathsf{RRS}_i = \max_{\mathbf{x}\subseteq\mathcal{G}, |\mathbf{x}|\le\lfloor k/n\rfloor} v_i(\mathbf{x}).$$

[3] If the NW is 0 for all allocations, MNW allocations are defined as those which give non-zero utility to maximum number of agents, and then maximize the product of utilities for those agents. Note if $k \ge n$, every agent positively values at least one good and thus MNW > 0.

For $\alpha \in (0,1]$, an allocation $\mathbf{x}$ is said to satisfy:

1. α-Proportionality (α-Prop) if $\forall i \in \mathcal{A}$, $v_i(\mathbf{x}) \geq \alpha \mathsf{Prop}_i$;
2. α-Proportionality up to one good (α-Prop1) if $\forall i \in \mathcal{A}$, $\exists g \in \mathbf{x}, g' \in \mathcal{G}$, such that $v_i((\mathbf{x} \setminus g) \cup g') \geq \alpha \mathsf{Prop}_i$,
3. α-RRS if for all agents $i \in \mathcal{A}$, $v_i(\mathbf{x}) \geq \alpha \mathsf{RRS}_i$.

Due to the cardinality constraints in the PublicGoods model, the notion of Prop1 requires that for every agent, there is a way to swap one preferred unpicked good with one picked good, after which the agent gets her proportional share. Since Prop1 in PrivateGoods requires only giving an extra good, this makes the definition of Prop1 in PublicGoods slightly more demanding than that in PrivateGoods.

Pareto-optimality. An allocation $\mathbf{y}$ is said to Pareto-dominate an allocation $\mathbf{x}$ if for all agents $i \in \mathcal{A}$, $v_i(\mathbf{y}) \geq v_i(\mathbf{x})$, with at least one of the inequalities being strict. We say $\mathbf{x}$ is Pareto-optimal (PO) if there is no allocation that Pareto-dominates $\mathbf{x}$.

Related models.

1. PrivateGoods. The classic problem of *private goods allocation* concerns partitioning a set of goods $\mathcal{G}$ among the set $\mathcal{A}$ of agents. Thus, a feasible allocation $\mathbf{x}$ is an n-partition $(\mathbf{x}_1, \ldots, \mathbf{x}_n)$ of $\mathcal{G}$, where agent i is allotted $\mathbf{x}_i \subseteq \mathcal{G}$, and derives utility $v_i(\mathbf{x}_i)$ only from $\mathbf{x}_i$.
2. PublicDecisions. In this model, a set $\mathcal{A}$ of agents are required to make *decisions* on a set $\mathcal{G}$ of issues. Each issue $j \in \mathcal{G}$ has a set $\mathcal{G}_j$ of k_j alternatives, given by $\mathcal{G}_j := \{(j,1),(j,2),\ldots,(j,k_j)\}$. A feasible allocation or outcome $\mathbf{x} = (x_1, \ldots, x_m)$ comprises of m decisions, where $x_j \in [k_j]$ is the decision made on issue j. Assuming the valuations are additive, each agent has a value $v_i(j,\ell)$ for the ℓ^{th} alternative of issue j. The valuation of the agent for the outcome $\mathbf{x}$ is then $v_i(\mathbf{x}) = \sum_{j \in \mathcal{G}} v_i(j, x_j)$.

3 Relating the models

We first show rigorous mathematical connections between the PrivateGoods, PublicGoods and PublicDecisions models w.r.t. computing optimal MNW and leximin allocations.

▶ **Theorem 4.** PublicMNW *polynomial-time reduces to* DecisionMNW.

Proof. Let $\mathcal{I} = (\mathcal{A}, \mathcal{G}, k, \{v_i\}_{i \in \mathcal{A}})$ be an instance of the PublicGoods model. For $k = m$, the MNW problem is trivial, since we can select all the m goods. For $n \leq k < m$, we can construct an instance $\mathcal{I}' = (\mathcal{A}', \mathcal{G}', \{\mathcal{G}_j\}_{j \in \mathcal{G}'} \{v'_i\}_{i \in \mathcal{A}'})$ of PublicDecisions from $\mathcal{I}$ in polynomial time, such that given an MNW allocation of $\mathcal{I}'$, we can compute an MNW allocation of $\mathcal{I}$ in polynomial time. Let $V = \max_{i,j} v_{ij}$. We create m public issues: corresponding to each good $j \in \mathcal{G}$, we create an issue j with two alternatives $(j,1)$ and $(j,2)$. That is, $\mathcal{G}' = [m]$, and $\mathcal{G}_j = \{(j,1),(j,2)\}$ for $j \in \mathcal{G}'$. We create $\mathcal{A}' = [n + mT]$, where $T = \lceil 2mn \log mV \rceil$. The first n agents here correspond to the n agents in $\mathcal{I}$. The last mT agents are of two types: kT agents $\{n+1, \ldots, n+kT\}$ of type A, and $(m-k)T$ agents $\{n+kT+1, \ldots, n+mT\}$ of type B. The valuations are as follows: each agent $i \in [n]$ values alternative "1" of the issue $j \in \mathcal{G}'$ at v_{ij}, the agents of type A value only alternative "1", agents of type B value only alternative "2". Formally, for $i \in \mathcal{A}'$, and an alternative (j,c) of the issue $j \in \mathcal{G}'$, where $c \in \{1,2\}$:

$$v_i'(j,c) = \begin{cases} v_{ij}, & \text{if } c = 1 \text{ and } i \in [n]; \\ 1, & \text{if } n < i \leq n + kT \text{ and } c = 1; \\ 1, & \text{if } n + kT < i \leq n + mT \text{ and } c = 2; \\ 0, & \text{otherwise.} \end{cases}$$

Let $\mathbf{x}'$ be an allocation for the instance $\mathcal{I}'$. For $c \in \{1,2\}$, let S_c be the set of issues j with decision c in $\mathbf{x}'$. That is, $S_c = \{j \in [m] : \mathbf{x}'_j = c\}$. Let $k' = |S_1|$. Then we have:

$$\mathsf{NW}(\mathbf{x}') = \left(\prod_{i \in [n]} v_i'(\mathbf{x}') \cdot (k')^{kT} \cdot (m-k')^{(m-k)T} \right)^{\frac{1}{n+mT}}.$$

We now relate $\mathbf{x}'$ to the PublicGoods instance $\mathcal{I}$. The decision $(j,1)$ corresponds to selecting the public good j. Let $\mathbf{x} = S_1 \subseteq \mathcal{G}$ be the corresponding set of public goods. Then for any $i \in [n]$ we have that $v_i(\mathbf{x}) = v_i'(\mathbf{x}')$, since $v_i'(j,2) = 0$ for every $j \in [m]$. Thus:

$$\mathsf{NW}(\mathbf{x}') = \left(\mathsf{NW}(\mathbf{x})^n \cdot (k')^{kT} \cdot (m-k')^{(m-k)T}\right)^{\frac{1}{n+mT}}. \tag{3}$$

We now have to prove that $\mathbf{x}$ satisfies $|\mathbf{x}| \leq k$. Let W_ℓ be the Nash product of any MNW allocation for the PublicGoods instance $\mathcal{I}_\ell = (\mathcal{A}, \mathcal{G}, \ell, \{v_i\}_{i \in \mathcal{A}})$, $0 \leq \ell \leq m$. Clearly, $0 = W_0 \leq W_1 \leq \ldots W_m \leq (mV)^n$. As $k \geq n$, $W_k \geq 1$, since we assume every agent has at least one good that she values positively. Define $g : [m] \to \mathbb{Z}$, as $g(a) = a^k(m-a)^{m-k}$. Then if $\mathbf{x}'$ is an MNW allocation for $\mathcal{I}'$, (3) becomes:

$$\mathsf{NW}(\mathbf{x}') = (W_{k'} \cdot g(k')^T)^{1/(n+mT)}. \tag{4}$$

Let G_1 and G_2 denote the largest and second-largest values that g attains over its domain. We observe that g increases in $[0,k]$, and decreases in $[k,m]$. Hence, $G_1 = g(k)$ implying:

$$G_1 = k^k(m-k)^{m-k}; G_2 = \max(g(k-1), g(k+1)).$$

We now claim the following and prove it in Appendix A:

▷ Claim 5. $G_1^T > W_m \cdot G_2^T$.

Using Claim 5, we have for all $k' \in [m] \setminus \{k\}$:

$$W_k \cdot g(k)^T \geq G_1^T > W_m \cdot G_2^T \geq W_{k'} \cdot g(k')^T,$$

Hence, the quantity $W_{k'} \cdot g(k')^T$ is maximized when $k' = k$. Recalling (4), we conclude that for the MNW allocation $\mathbf{x}'$ of $\mathcal{I}'$, the corresponding set $\mathbf{x}$ has cardinality exactly k. Further $\mathbf{x}$ also maximizes the NW among all allocations of the instance $\mathcal{I}$ satisfying this cardinality constraint. Thus, $\mathbf{x}$ in fact is an MNW allocation for $\mathcal{I}$. Finally, it is clear that this is a polynomial time reduction. ◀

We next relate the MNW problem in the PrivateGoods model with the PublicGoods model.

▶ **Theorem 6.** PrivateMNW *polynomial-time reduces to* PublicMNW.

Proof. Let $\mathcal{I} = (\mathcal{A} = [n], \mathcal{G} = [m], V)$ be a PrivateGoods instance, using which we create a PublicGoods instance I' as follows. We create $n + 2m$ agents, i.e. $\mathcal{A}' = [n + 2m]$. The first n agents correspond to the n agents in $\mathcal{I}$. The last $2m$ are dummy agents. We create $n \cdot m$ public goods: for each good $j \in [m]$, we create a set of n copies $S_j = \{j_1, j_2, \dots, j_n\}$, $\mathcal{G}' = \bigcup_{j \in \mathcal{G}} S_j$. We set $k = m$. The valuations for $i \in \mathcal{A}'$, $j_\ell \in \mathcal{G}'$ are:

$$v'_i(j_\ell) = \begin{cases} v_{ij}, \text{ if } i = \ell \text{ and } i \in [n]; \\ 1, \text{ if } i \in \{n + 2j - 1, n + 2j\}; \\ 0, \text{ otherwise,} \end{cases}$$

i.e. each agent $i \in [n]$ values exactly one copy, j_i for each $j \in \mathcal{G}$ at v_{ij}, and for each good $j \in \mathcal{G}$, there are exactly two dummy agents who value all copies of j.

We use the following claim in our proof. We prove it in Appendix A.

▷ Claim 7. Any MNW allocation $\mathbf{x}'$ of $\mathcal{I}'$ does not select two goods from same $S_j, j \in [m]$.

Consider any MNW allocation $\mathbf{x}'$ of $\mathcal{I}'$. We construct a partition, $\mathbf{x}$ of goods for $\mathcal{I}$ from this in the following way. For $i \in [n]$, $j \in [m]$, define $x_{ij} = 1$ if $j_i \in \mathbf{x}'$, and 0 otherwise. Let $\mathbf{x}_i = \{j \in \mathcal{G} : x_{ij} = 1\}$. Thus, the value that agent i gets in $\mathbf{x}$ is

$$\begin{aligned} v_i(\mathbf{x}_i) = \sum_{j \in \mathcal{G}} v_{ij} x_{ij} &= \sum_{j \in \mathcal{G}} v_{ij} \mathbf{1}(j_i \in \mathbf{x}'), \\ &= \sum_{j \in \mathcal{G}} v'_i(j_i) \mathbf{1}(j_i \in \mathbf{x}'), \\ &= v'_i(\mathbf{x}'). \end{aligned}$$

Thus, if $m \geq n$, $\mathsf{NW}(\mathbf{x}) = \mathsf{NW}(\mathbf{x}')^{(n+2m)/n}$ and the partition corresponding to $\mathbf{x}'$ as defined above gives an MNW solution for $\mathcal{I}$. On the other hand, if $m < n$, then $\mathbf{x}'$ already gives non-zero value to all dummy agents by Claim 7. Thus, to maximize the total number of agents who get non-zero value, it maximizes the number of agents in $[n]$ who get non-zero value. Call this set S^*. Thus partition $\mathbf{x}$ has maximum number of agents getting a non-zero value. Finally, it maximizes the Nash product over $S^* \cup \{n + 1, \dots, n + 2m\}$. Claim 7 also implies that all dummy agents get value 1. Thus, $\prod_{i \in S^*} v_i(\mathbf{x}_i) = \prod_{i \in S^*} v_i(\mathbf{x}')$. Thus even in this case the allocation $\mathbf{x}$ corresponds to an MNW allocation in $\mathcal{I}$. ◀

▶ **Observation 8.** *A desirable feature of the above reductions for the MNW problem from instance $\mathcal{I} = (\mathcal{A}, \mathcal{G}, V)$ to $\mathcal{I}' = (\mathcal{A}', \mathcal{G}', V')$ is that $V' = V \cup \{0, 1\}$, i.e., the reduction only creates instances $\mathcal{I}'$ which have 0 and 1 as the only potentially additional values as compared to $\mathcal{I}$. We use this feature in establishing the computational complexity of computing an MNW allocation in the* PublicDecisions *model with binary values, see Corollary 23.*

Similar polynomial-time reductions hold between the three models for the problem of computing a leximin-optimal allocation. We give the theorem statements here and the proofs can be found in full version of the paper.

▶ **Theorem 9.** PublicLex *polynomial-time reduces to* DecisionLex.

▶ **Theorem 10.** PrivateLex *polynomial-time reduces to* PublicLex.

4 Properties of MNW and Leximin

We prove that MNW and leximin-optimal allocations satisfy desirable fairness and efficiency properties in the PublicGoods model as well. First, we show some interesting relations between our three fairness notions – Prop, Prop1, and RRS in the PublicGoods model where $k \geq n$.[4] Our results are presented in Table 2.

Table 2 Relations between the fairness notions for $k \geq n$. Each cell (R, C) contains a factor α s.t. any allocation satisfying the row property R implies an α-approximation to the column property C. Cells with $\alpha = 1$ are marked with ✓, and with $\alpha = 0$ are marked with ✗.

	RRS	Prop	Prop1
RRS	✓	$\frac{n}{2n-1}$ (Lem. 12)	✓(Lem. 11)
Prop	$1/n$ (Lem. 13)	✓	✓
Prop1	✗(Ex. 14)	✗(Ex. 14)	✓

▶ **Lemma 11.** *Any allocation that satisfies* RRS *also satisfies* Prop1*.*

Proof. Fix any agent i. Let $\mathbf{x} = \{h_1, h_2, \dots, h_k\}$ be any allocation that satisfies RRS. Let $\mathbf{x}_k^* = \{g_1, g_2, \dots, g_k\}$ denote the top k goods for agent i. We assume that the goods both in $\mathbf{x}$ and $\mathbf{x}_k^*$ are ordered in decreasing order of valuations according to agent i. Now, suppose that top ℓ goods of $\mathbf{x}$ match with top ℓ goods of $\mathbf{x}_k^*$, i.e. $v_i(h_j) = v_i(g_j), \forall j \leq \ell$ and $v_i(h_{\ell+1} < v_i(g_{\ell+1}))$. Note that since $\mathbf{x}_k^*$ is the top k goods of agent i, we cannot have that $v_i(h_j) > v_(g_j)$ for any $j \leq \ell$. We want to prove that RRS implies Prop1. If $\mathbf{x}$ was already satisfying proportionality, it is obvious that $\mathbf{x}$ is Prop1. If $\ell \geq d$, it is again easy to see that $\mathbf{x}$ is Prop1. This is because, if $k = d$ then we already have top k goods, giving a proportional allocation. If $k > d$, then we can remove any good from $h_{d+1}, \dots, h_k$ and exchange it with g_{d+1} to ensure proportionality, making the original allocation Prop1. Finally, if n divides k then we have proportionality implied by RRS from Lemma 12.

Thus, we now assume that $\ell < d$, $k = nd + r$ with $r \leq n - 1$ and that $\mathbf{x}$ is not already a proportional allocation. We know that $v(h_1, \dots, h_\ell) = v(g_1, \dots, g_\ell)$ and $v(h_1, \dots, h_k) < \frac{1}{n} v(g_1, g_2, \dots, g_k)$. Thus,

$$v(h_{\ell+1}, \dots, h_k) < \frac{1}{n} v(g_{\ell+1}, \dots, g_k) \tag{5}$$

Now, $v(h_k) \leq \frac{1}{k-\ell} v(h_{\ell+1}, \dots, h_k)$. Thus,

$$v(h_k) \leq \frac{1}{n \cdot (k-\ell)} v(g_{\ell+1}, \dots, g_k) \tag{6}$$

Now, consider the good $g_{\ell+1}$. It is the good with highest value that is not in $\mathbf{x}$. We prove that removing h_k and adding $g_{\ell+1}$ gives us an allocation that is proportional. Since $\ell < d$, $v_i(g_{\ell+1}) \geq v_i(g_{nd+j}), \forall j \leq r$. Combining with the fact that $r < n$,

$$(n-1) \cdot v_i(g_{\ell+1}) \geq v_i(g_{nd+1}, \dots, g_{nd+r}). \tag{7}$$

[4] Note that when $k < n$, RRS is 0. Any agent who gets 0 value satisfies Prop1 when $k < n$ trivially. Thus, RRS and Prop1 coincide when $k < n$. On the other hand, the proportional value will be non-zero even when $k = 1$ if the agent likes at least one good. Thus, there can be no multiplicative relation between RRS and Prop when $k < n$.

Again since the goods are arranged in decreasing order of valuations, $v_i(g_1, \dots, g_d) \geq v_i(g_{jd+1}, \dots, g_{(j+1)d}), \forall 1 \leq j \leq (n-1)$. Thus,

$$(n-1) \cdot v_i(g_1, \dots, g_d) \geq v_i(g_{d+1}, \dots, g_{nd}). \tag{8}$$

Define, $LHS = (n-1)v_i(g_{\ell+1}) + (n-1)v_i(g_1, \dots, g_d)$. Combining (7) and (8),

$$\begin{aligned} LHS &\geq v_i(g_{nd+1}, \dots, g_{nd+r}) + v_i(g_{d+1}, \dots, g_{nd}) \\ &= v_i(g_{d+1}, \dots, g_k) \\ &= v_i(g_{\ell+1}, \dots, g_k) - v_i(g_{\ell+1}, \dots, g_d) \end{aligned}$$

Thus we get,

$$(n-1)v_i(g_{\ell+1}) + (n-1)v_i(g_1, \dots, g_\ell) \geq v_i(g_{\ell+1}, \dots, g_k) - nv_i(g_{\ell+1}, \dots, g_d)$$

Now, $v_i(g_{\ell+1}) \geq \frac{1}{k-\ell} v_i(g_{\ell+1}, \dots, g_k)$. Hence,

$$\begin{aligned} nv_i(g_{\ell+1}) + (n-1)v_i(g_1, \dots, g_\ell) &\geq v_i(g_{\ell+1}, \dots, g_k) - nv_i(g_{\ell+1}, \dots, g_d) + \frac{1}{k-\ell} v_i(g_{\ell+1}, \dots, g_k) \\ &\geq v_i(g_{\ell+1}, \dots, g_k) - nv_i(h_{\ell+1}, \dots, h_k) + nv_i(h_k), \end{aligned}$$

where the second inequality follows because $\mathbf{x}$ is RRS and from (6). Rearranging the above terms and using the fact that $v_i(g_1, \dots, g_\ell) = v_i(h_1, \dots, h_\ell)$, we get

$$nv_i(g_{\ell+1}) + nv_i(h_1, \dots, h_k) - nv_i(h_k) \geq v_i(g_1, \dots, g_k)$$

which implies that $\mathbf{x}$ is Prop1. ◀

▶ **Lemma 12.** *Any allocation that is α-RRS is also $\alpha \cdot \frac{n}{2n-1}$-Prop. Further, when n divides k, α-RRS implies α-Prop.*

Proof. We will prove a stronger result assuming the valuations are monotone and subadditive.

Let $\mathbf{x}$ denote any subset of k items that satisfies $\alpha \cdot$ RRS. Fix any agent i. We have,

$$v_i(\mathbf{x}) \geq \alpha \cdot \max_{|\mathbf{y}| \leq \lfloor k/n \rfloor} v_i(\mathbf{y}).$$

Let $\mathbf{x}^*$ denote the set of top k goods of agent i. Let $k = n * d + r$ where $r < n$. We can partition $\mathbf{x}^*$ by dividing it into n bundles, each of size $\lfloor k/n \rfloor$ and r more bundles, each of size 1. Note that when $k \geq n$, $\lfloor k/n \rfloor \geq 1$ and $r < n$. Thus, we get at most $2n-1$ bundles each of size at most $\lfloor k/n \rfloor$. We denote these bundles by $S_1, S_2, \dots, S_l$, with $l \leq 2n-1$. Thus, we have,

$$\begin{aligned} v_i(\mathbf{x}^*) &= v_i(\cup_{i \in [l]} S_i), \\ &\leq \sum_{i \in [l]} v_i(S_i), \\ &\leq \sum_{i \in [l]} \frac{1}{\alpha} \cdot v_i(\mathbf{x}), \\ &\leq v_i(\mathbf{x}) \cdot \frac{2n-1}{\alpha}. \end{aligned} \tag{9}$$

Here the second inequality follows from subadditivity and third follows because $\mathbf{x}$ is RRS. Thus, we have

$$v_i(\mathbf{x}) \geq \frac{\alpha}{2n-1} v_i(\mathbf{x}^*) = \alpha \cdot \frac{n}{2n-1} \mathsf{Prop}_i.$$

Further, when n divides k, $r = 0$ and we get $l = n$ bundles each of size k/n. Thus, we have from (9)

$$v_i(\mathbf{x}^*) \leq \frac{n}{\alpha} \cdot v_i(\mathbf{x}).$$

Thus,

$$v_i(\mathbf{x}) \geq \frac{\alpha}{n} v_i(\mathbf{x}^*) = \alpha \mathsf{Prop}_i.$$ ◀

▶ **Lemma 13.** *Any allocation that satisfies α-Prop gives an α/n multiplicative approximation to RRS, and this is tight.*

Proof. Suppose a given allocation, $\mathbf{x}$ satisfies α-Prop. Fix any agent i.

$$\begin{aligned} v_i(\mathbf{x}) &\geq \alpha \cdot \frac{1}{n} \cdot \max_{|\mathbf{y}| \leq k} v_i(\mathbf{y}), \\ &\geq \alpha \cdot \frac{1}{n} \cdot \max_{|\mathbf{y}| \leq \lfloor k/n \rfloor} v_i(\mathbf{y}), \\ &= \frac{\alpha}{n} \cdot \mathsf{RRS}. \end{aligned}$$

For the tightness of lemma, consider the following example: We have $n = 2$ agents and $m = 5$ goods. Agent 1 values goods 1 and 2 at 1 each, does not value goods $3, 4, 5$. Agent 2 values all goods at 1. If $k = 4$, the RRS value of agent 1 is 2. Her proportional value is 1. Thus, picking goods $1, 3, 4, 5$ gives agent 1 her Prop share but only ensures $1/n$ of her RRS share. ◀

Finally, we note in the following example that Prop1 will not give any approximation to either Prop or RRS.

▶ **Example 14** (Prop1 does not approximate Prop or RRS)**.** Finally, we note that a Prop1 allocation might not give an α approximation to RRS for any $\alpha > 0$. Consider an instance of public goods allocation with $n = 2$. We have 3 goods. Agent 1 values goods 1, 2 at value of 1 and values good 3 at 0. Agent 2 values goods 1, 2 at 0 and values good 3 at 1. If we want to select $k = 2$ goods, then, selecting goods 1 and 2 gives agent 2 value 0. This allocation is Prop1, but provides no multiplicative approximation to either RRS or Prop for agent 2.

Next, we show that MNW allocations are fair:

▶ **Lemma 15.** *All MNW allocations satisfy Prop1.*

Proof. Suppose there exists an MNW allocation $\mathbf{x}$ that is not Prop1. This implies for some agent $i \in \mathcal{A}$, for all pairs of goods $j \in \mathbf{x}$ and $j' \notin \mathbf{x}$, $v_i((\mathbf{x} \setminus j) \cup j') < \mathsf{Prop}_i$. If $k < n$, $\mathsf{Prop}_i \leq \max_{j \in \mathcal{G}} v_{ij}$, and swapping any good in $\mathbf{x}$ with this good will give her her proportional share.

Consider now $k \geq n$. Since we assume each agent positively values at least one good, the MNW value is non-zero. Since MNW is scale-invariant, we scale the valuations of agents so that $v_h(\mathbf{x}) = 1 \ \forall h \neq i$. Let g' be the highest-valued good of i not in $\mathbf{x}$, i.e., $g' = \mathsf{argmax}_{j \in \mathcal{G} \setminus \mathbf{x}} v_{ij}$. Let $\mathbf{x}_0 = \{j \in \mathbf{x} : v_{ij} < v_{ig'}\}$ be the set of goods in $\mathbf{x}$ that give i strictly lesser value than g'. Since i does not satisfy Prop1, $\mathbf{x}_0 \neq \emptyset$. Suppose we order the goods in $\mathcal{G}$ according to the valuation of i as $\{g_1, \ldots, g_m\}$, where $v_i(g_r) \geq v_i(g_s)$ for $1 \leq r \leq s \leq m$. Then $n \cdot \mathsf{Prop}_i = v_i(g_1, \ldots, g_k)$ by definition. Since g' is the highest-valued good for i not in $\mathbf{x}$, and further since every good in $\mathbf{x}_0$ is valued at less than $v_{ig'}$ by i, we can bound the total value to i of the top k goods $g_1, \dots, g_k$ as follows: $v_i(g_1, \dots, g_k) \leq v_i(\mathbf{x} \setminus \mathbf{x}_0) + |\mathbf{x}_0| v_{ig'}$ which, using additivity of v_i, can alternatively be written as:

$$v_i(\mathbf{x}) + \sum_{j \in \mathbf{x}_0} (v_{ig'} - v_{ij}) \geq n\mathsf{Prop}_i. \tag{10}$$

Consider a good g given by[5]:

$$g \in \mathsf{argmin}_{j \in \mathbf{x}_0} \frac{\sum_{h \in \mathcal{A} \setminus \{i\}} v_{hj}}{v_{ig'} - v_{ij}}.$$

Then by definition of g, we have:

$$\begin{aligned} \frac{\sum_{h \in \mathcal{A} \setminus \{i\}} v_{hg}}{v_{ig'} - v_{ig}} \leq \frac{\sum_{j \in \mathbf{x}_0} \sum_{h \in \mathcal{A} \setminus \{i\}} v_{hj}}{\sum_{j \in \mathbf{x}_0} v_{ig'} - v_{ij}} &\leq \frac{\sum_{h \in \mathcal{A} \setminus \{i\}} \sum_{j \in \mathbf{x}_0} v_{hj}}{n\mathsf{Prop}_i - v_i(\mathbf{x})} \\ &\leq \frac{n-1}{n\mathsf{Prop}_i - v_i(\mathbf{x})}, \end{aligned} \tag{11}$$

where the first transition follows by rearranging terms in the numerator, and using (10) in the denominator, and the final transition follows by recalling that $v_h(\mathbf{x}) = 1$ for all $h \neq i$.

Let $\delta = v_{ig'} - v_{ig}$. We know $v_i(\mathbf{x}) + \delta < \mathsf{Prop}_i$. Substituting this in (11), and noting $\delta > 0$ gives:

$$\frac{\sum_{h \in \mathcal{A} \setminus \{i\}} v_{hg}}{\delta} < \frac{1}{v_i(\mathbf{x}) + \delta}. \tag{12}$$

Let us now consider the allocation $\mathbf{x}' = (\mathbf{x} \setminus g) \cup g'$. We show $\mathsf{NW}(\mathbf{x}') > \mathsf{NW}(\mathbf{x})$, thus contradicting the Nash optimality of $\mathbf{x}$. Since for any $h \neq i$, $v_h(\mathbf{x}') \geq v_h(\mathbf{x}) - v_{hg} = 1 - v_{hg}$, we have:

$$\begin{aligned} \prod_{h \in \mathcal{A}} v_h(\mathbf{x}') \geq v_i(\mathbf{x}') \prod_{h \in \mathcal{A} \setminus \{i\}} (1 - v_{hg}) &\geq (v_i(\mathbf{x}) + \delta) \left(1 - \sum_{h \in \mathcal{A} \setminus \{i\}} v_{hg}\right) \\ &> (v_i(\mathbf{x}) + \delta) \left(1 - \frac{\delta}{v_i(\mathbf{x}) + \delta}\right) = v_i(\mathbf{x}), \end{aligned}$$

where the first transition uses Weierstrass' inequality [24], and the second transition uses (12). This leads to $\mathsf{NW}(\mathbf{x}') > \mathsf{NW}(\mathbf{x})$, giving the desired contradiction. Hence any MNW allocation satisfies Prop1. ◀

Besides Prop1, the MNW allocation satisfies several other desirable properties, as our next result shows.

▶ **Theorem 16.** *All MNW allocations satisfy* PO*,* Prop1*, and* $1/n$-RRS*. Further when* $k \geq n$*, MNW allocation implies* $\frac{1}{2n-1}$*-Prop.*

Proof. If any MNW allocation did not satisfy Pareto optimality, then at least one of the agents gets a strictly higher value with values of all other agents not decreasing. Thus, if the MNW value is non-zero, we get an allocation with strictly higher Nash Product, contradicting

5 [15] considered an *issue* similarly.

the optimality of value of MNW. On the other hand, if MNW value is zero and the strict increase of value holds for one of the agents with non-zero value, then the Nash Product over these agents increases contradicting maximality of Nash Product of these agents. On the other hand, if the strict inequality holds for an agent who receives zero value, the number of agents with non-zero value increases, contradicting the maximality of number of agents who get non-zero value. In both cases, the optimality of MNW is contradicted. Thus any MNW allocation satisfies Pareto Optimality.

Next we prove that all MNW allocations satisfy $1/n$-RRS. Suppose there exists an MNW allocation $\mathbf{x}$ that is not $1/n$-RRS. This implies that for some agent $i \in \mathcal{A}$, $v_i(\mathbf{x}) < \frac{1}{n}\mathsf{RRS}_i$. Let us order the goods according to i's valuation: let $\mathcal{G} = \{g_1, g_2, \ldots, g_m\}$, such that $v_i(g_r) \geq v_i(g_s)$, for all $1 \leq r \leq s \leq m$. Let $p = \lfloor \frac{k}{n} \rfloor$. When $k < n$, $p = 0$, in that case $\mathsf{RRS}_i = 0$. Therefore, $k \geq n$. Observe that the round-robin share of i is given by $\mathsf{RRS}_i = v_i(\{g_1, \ldots, g_p\})$. We scale the valuations of the agents so that for every agent i, $v_i(\mathbf{x}) = 1$. In particular, this implies $\mathsf{RRS}_i > n$.

Let us order the goods in $\mathbf{x}$ according to i's valuation: let $\mathbf{x} = \{j_1, j_2, \ldots, j_k\}$, such that $v_i(j_r) \geq v_i(j_s)$, for all $1 \leq r \leq s \leq k$. Define for $r \in [p]$, $S_r = \{j_{rn-n+1}, \ldots, j_{rn}\}$, and $g'_r = \mathsf{argmin}_{j \in S_r} \sum_{h \in \mathcal{A}\setminus\{i\}} v_{hj}$.

We now construct another allocation $\mathbf{x}'$ as follows. We first check if $g_1 \in \mathbf{x}$. If not, we begin constructing $\mathbf{x}'$ by removing g'_1 from $\mathbf{x}$ and adding g_1. If $g_1 \in \mathbf{x}$, then we proceed to check whether $g_2 \in \mathbf{x}$ or not. For every $r \in [p]$, we remove g'_r and add g_r if g_r is not in $\mathbf{x}$. If g_r is already in $\mathbf{x}$ then for such an r no operation is done. Since we are removing g'_r and $v_i(g'_r) < v_i(g_r) \leq v_i(g_s)$ for all $s < r$, this ensures that $\{g_1, \ldots, g_p\} \subseteq \mathbf{x}'$, which shows $v_i(\mathbf{x}') \geq \mathsf{RRS}_i > n$. Observe that:

$$
\begin{aligned}
\sum_{r=1}^{p} \sum_{h \in \mathcal{A}\setminus\{i\}} v_h(g'_r) &\leq \sum_{r=1}^{p} \frac{1}{n} \sum_{h \in \mathcal{A}\setminus\{i\}} \sum_{j \in S_r} v_{hj} && \text{(def. of } g_r\text{)} \\
&\leq \frac{1}{n} \sum_{r=1}^{p} \sum_{j \in S_r} \sum_{h \in \mathcal{A}\setminus\{i\}} v_{hj} && \text{(rearranging)} \\
&\leq \frac{1}{n} \sum_{j \in \mathbf{x}} \sum_{h \in \mathcal{A}\setminus\{i\}} v_{hj} && \text{(def. of } S_r\text{)} \\
&\leq \frac{1}{n} \sum_{h \in \mathcal{A}\setminus\{i\}} v_h(\mathbf{x}^*) && \text{(rearranging)} \\
&= \frac{n-1}{n}.
\end{aligned}
$$

Then we have:

$$
\begin{aligned}
\mathsf{NW}(\mathbf{x}')^n = \prod_{h \in \mathcal{A}} v_h(\mathbf{x}') &\geq v_i(\mathbf{x}') \prod_{h \in \mathcal{A}\setminus\{i\}} v_h(\mathbf{x}'), \\
&\geq v_i(\mathbf{x}') \prod_{h \in \mathcal{A}\setminus\{i\}} \left(1 - \sum_{r=1}^{p} v_h(g'_r)\right), \\
&\geq v_i(\mathbf{x}') \left(1 - \sum_{r=1}^{p} \sum_{h \in \mathcal{A}\setminus\{i\}} v_h(g_r)\right) > n\left(1 - \frac{n-1}{n}\right) = \mathsf{NW}(\mathbf{x})^n,
\end{aligned}
$$

which contradicts the fact that $\mathbf{x}$ is Nash optimal.

Combining this with Lemma (12) and Lemma (15), we get the proof of the theorem. ◀

Similar fairness and efficiency properties for the leximin-optimal allocation. In particular, one can prove the following theorem (proof is in full version of the paper).

▶ **Theorem 17.** *All leximin-optimal allocations are PO, satisfy* RRS *and* Prop1*. Further, when $k \geq n$, a leximin-optimal allocation is also* $(n/(2n-1))$-Prop.

5 Complexity of MNW and Leximin

In this section, we show that PublicMNW and PublicLex are NP-hard. Our hardness results also hold for instances with binary values, which is in stark contrast to the private goods setting, where MNW and leximin-optimal allocations can be computed in polynomial-time. All proofs for this section can be found in the full version of paper. Since the cases of $k \geq n$ and $k < n$ are interesting in their own right, we consider them separately.

▶ **Theorem 18.** *Given a* PublicGoods *allocation instance where $k < n$, computing an α-approximation to MNW is* NP*-hard for any $\alpha > 0$, even when all valuations are binary.*

▶ **Theorem 19.** PublicMNW *is* NP*-hard, even when all valuations are binary.*

Proof. (Sketch) We reduce from the exact regular set packing (ERSP) problem. In the input to ERSP, there are n elements $X = \{x_1, \ldots, x_n\}$, a family of subsets $\mathcal{F} = \{F_1, \ldots, F_m\}$ where each $F_j \subseteq X$ and $|F_j| = d$. The problem is to compute a subfamily $\mathcal{F}' \subseteq \mathcal{F}$, $|\mathcal{F}'| = r$, s.t. for all $F_i \neq F_j \in \mathcal{F}', F_i \cap F_j = \emptyset$. Let $\mathcal{I} = (X, \mathcal{F}, d, r)$ be an instance of ERSP. We construct a PublicGoods instance $\mathcal{I}' = \{\mathcal{A}, \mathcal{G}, k, \{v_i\}_{i\in\mathcal{A}}, T\}$ as follows. We create a set $\mathcal{A} = [n]$ of n agents, a set $\mathcal{G} = \{g_1, \ldots, g_m\} \cup \{d_1, \ldots, d_n\}$ of $m+n$ public goods. For any agent $i \in \mathcal{A}$ and good $g_j \in \mathcal{G}$, $v_i(g_j) = 1$ if $x_i \in F_j$ else 0. For any agent $i \in \mathcal{A}$ and good $d_j \in \mathcal{G}$, $v_i(d_j) = 1$. We set $k = r + n$ and $T = ((n+1)^{dr} n^{n-dr})^{1/n}$. We show that $\mathcal{I}$ is a yes-instance for ERSP iff the MNW for $\mathcal{I}'$ is at least T. ◀

▶ **Theorem 20.** PublicMNW *is* NP*-hard, even for two agents.*

We next show a similar hardness results for computing leximin-optimal allocations, which as we show, apply even for instances with binary values.

▶ **Theorem 21.** PublicLex *is* NP*-hard, even when the valuations are binary.*

▶ **Theorem 22.** PublicLex *is* NP*-hard, even for two agents.*

Proof. (Sketch) We prove this by reducing from the NP-complete problem Monotone c-SAT. In an instance of Monotone c-SAT we have $X = \{x_1, \ldots, x_n\}$ variables, formula, $F = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ in CNF form with additional constraint that all literals in it are positive. We want to determine if we can satisfy F by setting at most c variables to true. To create $\mathcal{I} = (\mathcal{A}, \mathcal{G}, k, \{v_i\}_{i\in\mathcal{A}})$, corresponding to each clause C_i, we create an agent, i and corresponding to each variable, x_j we create a good, j. Each agent likes the goods corresponding to the variables that show up in her corresponding clause. To ensure that $k \geq |\mathcal{A}|$ in the public goods instance, we create one dummy agent and $m-c+1$ dummy goods. Finally, set $k = m+1$. We show that F has a satisfying assignment with c true variables iff in the Leximin-optimal the minimum utility is $m-c+1$, and the second minimum utility is at least $m-c+2$. ◀

Using the reductions of Theorems 4 and 9 and the NP-hardness results of this section, we obtain NP-hardness results for computing MNW and leximin allocations in the public decision making model. In fact, Observation 8 implies that this NP-hardness remains for the MNW problem even with the valuations are binary.

▶ **Corollary 23.** DecisionMNW *is* NP*-hard, even when all values are binary.*

Using our reductions (Theorems 4 and 9) together with the NP-hardness of PublicMNW and PublicLex (Theorems 19 and 21) implies that:

▶ **Corollary 24.** *The problems* DecisionMNW *and* DecisionLex *are* NP*-hard.*

6 Algorithms for MNW and Leximin

In light of the above computational hardness, we turn to approximation algorithms and exact algorithms for special cases. The proofs of results in this section and the algorithms for special cases can be found in the full version of paper. We first present an algorithm that provides an $O(n)$ factor approximation to MNW and satisfies fairness properties of RRS, Prop1 when valuations $\{v_i\}_{i\in\mathcal{A}}$ are *monotone* ($v_i(S) \le v_i(S\cup g)$ for all $S\subseteq\mathcal{G}$ and $g\in\mathcal{G}\setminus S$) and *subadditive* (for all $S_1\subseteq\mathcal{G}, S_2\subseteq\mathcal{G}$, $v_i(S_1)+v_i(S_2)\ge v_i(S_1\cup S_2)$). The class of subadditive valuations captures complement-free goods, and subsumes additive valuations. Our algorithm assumes access to demand oracles[6] for the subadditive valuations. We use the following subroutine, Maximize, from [5] which takes:

- Input: Set of goods, $\mathcal{G}$, the valuation function v_i of the agent i, and an integer r; and returns:
- Output: $\mathbf{x}\subseteq\mathcal{G}$, s.t. $v_i(\mathbf{x})\ge\frac{1}{2}\max_{S\subseteq\mathcal{G},|S|\le r} v_i(S)$

Our algorithm, AlgGreedy, has two steps:

- For all $i\in\mathcal{A}$, $\mathbf{x}_i \leftarrow \mathsf{Maximize}(\mathcal{G}, v_i, \lfloor\frac{k}{n}\rfloor)$
- Return $\mathbf{x}\leftarrow\cup_{i\in\mathcal{A}}\mathbf{x}_i$

For additive valuations, we assume that Maximize returns a set of $\lfloor k/n\rfloor$ most-preferred goods for each agent. This algorithm enables us to show that:

▶ **Theorem 25.** *There exists a polynomial-time algorithm for the problem of* PublicGoods *allocation (where $k\ge n$ and agents have monotone, subadditive valuations) that returns an allocation which satisfies* RRS*,* $\frac{1}{2}$-Prop*, and approximates the MNW to a factor of $O(n)$. Further, when the valuations are additive, the allocation satisfies* Prop1*.*

We now present pseudo-polynomial time algorithms for two special cases, namely constantly many types of agents, and constantly many types of goods. Our results apply to the more general model of *budget constraints*. We denote an instance of this model by $\mathcal{I}=(\mathcal{A},\mathcal{G},B,\{c_j\}_{j\in\mathcal{G}},\{v_i\}_{i\in\mathcal{A}})$. Each good $j\in\mathcal{G}$ has an associated integral cost c_j, and in a feasible allocation the sum of costs of the picked goods must not exceed the budget B. The MNW and leximin-objectives are defined as before, but over feasible allocations that satisfy the budget constraints. Since cardinality constraints are a special case of budget constraints with uniform cost, our hardness results apply for the budget model also.

[6] Subadditive valuations are set functions and cannot in general represented efficiently. We thus assume access to the functions through some oracles. Given a set of prices p_j for each good $j\in\mathcal{G}$, a demand oracle returns any set S that maximizes $v_i(S)-\sum_{j\in S}p_j$.

Constantly many types of agents

We consider instances where the number of *agent types* is constant. We say agents i and h have the same *type* if $\forall j \in \mathcal{G}$, $v_{ij} = v_{hj}$.

▶ **Theorem 26.** *For a* PublicGoods *allocation instance,* $\mathcal{I} = (\mathcal{A}, \mathcal{G}, B, \{c_j\}_{j\in\mathcal{G}}, \{v_i\}_{i\in\mathcal{A}})$ *with* t *distinct types of agents, (i) an MNW allocation can be computed in time* $O(m \cdot (mV)^t)$*, (ii) a leximin-optimal allocation can be computed in time* $O(m \cdot n \log n \cdot (mV)^t)$*, where* $V = \max_{i\in\mathcal{A}, j\in\mathcal{G}} v_{ij}$*.*

We prove this by presenting a dynamic-programming based algorithm which computes such allocations. We also get:

▶ **Corollary 27.** *For binary valuations, with constantly many types of agents* PublicMNW *and* PublicLex *are polynomial-time solvable.*

Constantly many types of goods

We now consider instances where the number of *types of goods* is constant. We say two goods $j_1, j_2 \in \mathcal{G}$ have same type if for all agents $i \in \mathcal{A}$, $v_{ij_1} = v_{ij_2}$ and $c_{j_1} = c_{j_2}$. In this case, we can enumerate all feasible allocations efficiently, implying that an MNW or leximin-optimal allocation can be computed in polynomial-time.

▶ **Theorem 28.** *For a* PublicGoods *allocation instance* $\mathcal{I} = (\mathcal{A}, \mathcal{G}, B, \{c_j\}_{j\in\mathcal{G}}, \{v_i\}_{i\in[n]})$ *with* t *different types of goods, (i) an MNW, can be computed in time* $O(m^t)$ *(ii) a leximin-optimal allocation can be computed in time* $O(n \log n \cdot m^t)$*.*

7 Discussion

In this paper, we considered the problem of allocating indivisible public goods to agents subject to a cardinality constraint. We showed fundamental connections between the models of private goods, public goods, and public decision making, by presenting polynomial-time reductions for the popular solution concepts of maximum Nash welfare (MNW) and leximin. We also showed that MNW and leximin-optimal allocations satisfy desirable fairness properties like Prop1 and RRS, and the efficiency property of PO. Further we showed that these objectives are computationally NP-hard, including for several special cases like constantly many agents and binary valuations. Lastly, we designed an approximation algorithm for MNW and pseudo-polynomial time algorithms for the case of constantly many agents.

Our work opens up several interesting research directions. Firstly, extending our reductions to the budget model presents a challenging problem. A second question is devising an algorithm to compute a Prop1+PO or RRS+PO allocations in polynomial time, bypassing the hardness of computing MNW or leximin-optimal allocations. Appropriately defining properties like Prop1 in the budget model and investigating whether MNW and leximin satisfy them would be a third interesting research direction. Finally, designing constant-factor approximation algorithms, even for restricted cases like binary valuations, which captures a large class of voting-like scenarios, is another important open problem.

References

1. S Airiau, H Aziz, I Caragiannis, J Kruger, and J Lang. Positional social decision schemes: Fair and efficient portioning. In *Proceedings of the 7th International Workshop on Computational Social Choice (COMSOC)*, 2018.
2. Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 784–790, 2015.

3 Haris Aziz, Barton E. Lee, and Nimrod Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 23–31, 2018.

4 Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. *arXiv preprint arXiv:2003.00606*, 2020.

5 Ashwinkumar Badanidiyuru, Shahar Dobzinski, and Sigal Oren. Optimization with demand oracles. *Algorithmica*, 81(6):2244–2269, 2019.

6 Siddharth Barman and Sanath Krishnamurthy. On the proximity of markets with integral equilibria. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pages 1748–1755, 2019.

7 Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 557–574, 2018.

8 Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Greedy algorithms for maximizing Nash social welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 7–13, 2018.

9 Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *SIGecom Exch.*, 5(3):11–18, April 2005.

10 S.J. Brams and A.D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

11 Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

12 Bhaskar Ray Chaudhury, Yun Kuen Cheung, Jugal Garg, Naveen Garg, Martin Hoefer, and Kurt Mehlhorn. On fair division for indivisible items. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 25:1–25:17, 2018.

13 Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. Group fairness in committee selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 263–279, 2019.

14 Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing (STOC)*, pages 371–380, 2015.

15 Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

16 Andreas Darmann and Joachim Schauer. Maximizing Nash product social welfare in allocating indivisible goods. *SSRN Electronic Journal*, 247, January 2014.

17 Brandon Fain, Ashish Goel, and Kamesh Munagala. The core of the participatory budgeting problem. In *Web and Internet Economics (WINE)*, pages 384–399, 2016.

18 Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.

19 Till Fluschnik, Piotr Skowron, Mervin Triphaus, and Kai Wilker. Fair knapsack. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1941–1948, 2019.

20 Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Equitable allocations of indivisible goods. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 280–286, 2019.

21 Jugal Garg, Martin Hoefer, and Kurt Mehlhorn. Satiation in Fisher markets and approximation of Nash social welfare. *CoRR*, abs/1707.04428, 2017.

22 Jugal Garg, Edin Husic, and László A. Végh. Approximating Nash social welfare under Rado valuations. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 1412–1425, 2021.

23 Jugal Garg, Pooja Kulkarni, and Aniket Murhekar. On fair and efficient allocations of indivisible public goods. *CoRR*, abs/2107.09871, 2021. arXiv:2107.09871.

24 M. S. Klamkin and D. J. Newman. Extensions of the Weierstrass product inequalities. *Mathematics Magazine*, 43(3):137–141, 1970. URL: http://www.jstor.org/stable/2688388.

25 Mayuresh Kunjir, Brandon Fain, Kamesh Munagala, and Shivnath Babu. ROBUS: Fair cache allocation for data-parallel workloads. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 219–234, 2017.

26 Euiwoong Lee. APX-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.

27 Peter McGlaughlin and Jugal Garg. Improving Nash social welfare approximations. *J. Artif. Intell. Res.*, 68:225–245, 2020.

28 H. Moulin. *Fair Division and Collective Welfare*. Mit Press. MIT Press, 2004.

29 Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020.

30 John Rawls. *A theory of justice*. Harvard university press, 2009.

31 Herbert E Scarf. The core of an N person game. *Econometrica: Journal of the Econometric Society*, pages 50–69, 1967.

32 Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

A Missing Proofs from Section 3

▷ Claim 5. $G_1^T > W_m \cdot G_2^T$.

Proof. Recall that W_ℓ denotes the Nash product of any MNW allocation for the PublicGoods instance $\mathcal{I}_\ell = (\mathcal{A}, \mathcal{G}, \ell, \{v_i\}_{i\in\mathcal{A}})$, for $0 \le \ell \le m$. We have $0 = W_0 \le W_1 \le \dots W_m \le (mV)^n$, and we assume $W_k \ge 1$. Recall that function $g : [m] \to \mathbb{Z}$, was defined as $g(a) = a^k(m-a)^{m-k}$. Let G_1 and G_2 denote the largest and second-largest values that g attains over its domain. We observe that g increases in $[0,k]$, and decreases in $[k,m]$. Hence:

$$G_1 = g(k) = k^k(m-k)^{m-k}.$$
$$G_2 = \max(g(k-1), g(k+1)).$$

Now observe that for $k \in [m] \setminus \{0, 1, m\}$:

$$\begin{aligned}\log g(k) - \log g(k-1) &= k(\log k - \log(k-1)) + (m-k)(\log(m-k) - \log(m-k+1)),\\ &> k \cdot \frac{1}{k-\frac{1}{2}} + (m-k)\cdot\frac{-1}{m-k} \ge \frac{1}{2k-1} \ge \frac{1}{2m},\end{aligned}$$

and for $k \in [m] \setminus \{0, m-1, m\}$:

$$\begin{aligned}\log g(k) - \log g(k+1) &= k(\log k - \log(k+1)) + (m-k)(\log(m-k) - \log(m-k-1)),\\ &> k\cdot\frac{-1}{k} + (m-k)\cdot\frac{1}{m-k-\frac{1}{2}},\\ &\ge \frac{1}{2(m-k)-1} \ge \frac{1}{2m},\end{aligned}$$

using standard properties of logarithms. Thus:

$$\log G_1 - \log G_2 > \frac{1}{2m}.$$

Then we have by recalling that $T = 2mn \log mV$,

$$T(\log G_1 - \log G_2) > 2mn \log mV \cdot \frac{1}{2m} \geq \log W_m,$$

which gives:

$$G_1^T > W_m \cdot G_2^T,$$

as required. Lastly, we consider the cases of $k = 1$ and $k = m - 1$. In both cases, $T(\log G_1 - \log G_2) = T[(m-1)\log(m-1) - \log 2 - (m-1)\log(m-2)] > 2mn \log mV \frac{1}{2m} \geq \log W_m$, which gives $G_1^T > W_m G_2^T$, as claimed. ◁

Proof of Claim 7. Consider first $m \geq n$. Suppose $\exists j \in [m]$ for which two goods $j_i, j_{i'} \in \mathbf{x}', i \neq i'$. Since exactly m goods are picked in $\mathbf{x}'$, there is some $j' \in [m]$, for which no good j'_i is picked in $\mathbf{x}'$ for any $i \in [n]$. This implies that the agents $2j' + n - 1, 2j' + n$ get zero value in $\mathbf{x}'$, making $\mathsf{NW}(\mathbf{x}') = 0$. However, choosing a good from each $j \in [m]$ gives non-zero value to all dummy agents. At the same time, since $m \geq n$, these goods can be chosen so that they give non-zero value to distinct agents in $[n]$. This makes $\mathsf{NW}(\mathbf{x}') \neq 0$ contradicting Nash optimality of $\mathbf{x}'$.

Now, if $m < n$ Nash welfare of all allocations in $\mathcal{I}$ is 0. Thus, the MNW allocation is the one that maximizes the number of agents who get non zero value and then maximizes the product of values for these agents. Consider any allocation $\bar{\mathbf{x}}$, suppose $\exists j \in [m]$ for which two goods $j_i, j_{i'} \in \bar{\mathbf{x}}, i \neq i'$. then again for some j', agents $n + 2j' - 1$ and $n + 2j'$ get value 0 making $\mathsf{NW}(\bar{\mathbf{x}}) = 0$. At the same time, even if $\bar{\mathbf{x}}$ has goods from all different S_j, since $m < n$, and each one item from S_j gives value only to one agent $i \in [n]$, the $\mathsf{NW}(\bar{\mathbf{x}}) = 0$ even in this case. Thus, if $m < n$, all allocations have Nash welfare 0 in $\mathcal{I}'$ also. Suppose the MNW allocation, $\mathbf{x}'$ had two goods from same S_j for some $j \in [m]$. Then, there exists a $j' \in [m]$ such that no good is selected from $S_{j'}$. The two goods from S_j give value to exactly four agents - the two dummy agents $2j + n - 1, 2j + n$ and two agents who receive their copy of good j. Instead, if we exchange one of these goods to a good from $S_{j'}$, we give non-zero value to at least five agents - dummy agents $2j + n - 1, 2j + n, 2j' + n - 1, 2j' + n$ and at least one of the agents in $[n]$. We did not change the value of any other agents in this process. Thus, we increase the number of agents who get non-zero value, contradicting the maximality of $\mathbf{x}'$. Thus, in both cases, all m goods are picked from different $S_j, j \in [m]$. ◁