

Fair Public Decision Making

VINCENT CONITZER, Duke University

RUPERT FREEMAN, Duke University

NISARG SHAH, Harvard University

We generalize the classic problem of fairly allocating indivisible goods to the problem of *fair public decision making*, in which a decision must be made on several social issues simultaneously, and, unlike the classic setting, a decision can provide positive utility to multiple players. We extend the popular fairness notion of proportionality (which is not guaranteeable) to our more general setting, and introduce three novel relaxations – *proportionality up to one issue*, *round robin share*, and *pessimistic proportional share* – that are also interesting in the classic goods allocation setting. We show that the Maximum Nash Welfare solution, which is known to satisfy appealing fairness properties in the classic setting, satisfies or approximates all three relaxations in our framework. We also provide polynomial time algorithms and hardness results for finding allocations satisfying these axioms, with or without insisting on Pareto optimality.

1 INTRODUCTION

The literature on mathematically rigorous fair division dates back to the work of Steinhaus [33]. In the field’s long history, most work focuses on the fair division of *private goods*, in which a set of m items must be divided among a set of n players. Agents express their preferences by specifying their value for each good, and our goal is to find a division of the goods that is fair to all players.

One particularly appealing notion of fairness is *envy-freeness* [18], which says that no player should want to switch her set of items with that of another player. This is a natural and strong notion of fairness that has long been the subject of fair division research [8, 9, 13, 21, 22, 32, 34]. It actually implies many other fairness notions such as *proportionality* [33] – each player should get at least a $1/n$ fraction of her value for the entire set of goods – and *envy-freeness up to one good (EF1)* [26] – no player should envy another player after removing at most one good from the latter player’s bundle. Unfortunately, envy-freeness cannot always be guaranteed, and therefore its relaxations have been focused on [11, 14, 26, 30].

Division of private goods, however, is not the only application in which we may desire a fair outcome. Often, we may need to make decisions where every alternative gives positive utility to many players, rather than to just one player as in the case of private goods. For instance, consider a couple, Alice and Bob, deciding where to go to dinner. Alice likes Italian food the most, but does not like Indian, whereas Bob prefers Indian food but does not like Italian. When there is only a single decision to make, we are simply in a classic bargaining game where players must attempt to arrive at a mutually agreeable solution. Nash [28] proposed maximizing the product of players’ utilities (the *Nash welfare*) as an elegant solution that uniquely satisfies several appealing properties. But no matter how we arrive at a decision – and there is a myriad of work in computational social

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. EC’17, June 26–30, 2017,

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Cambridge, Massachusetts, USA.

DOI: <http://dx.doi.org/10.1145/3033274.3085125>

ACM ISBN 978-1-4503-4527-9/17/06...\$15.00.

choice [10] discussing how exactly we should do so – some tradeoff must necessarily be made, and we may not be able to make everyone happy.

However, if we have several public decisions to make, maybe we can reach a compromise by making sure that all players are happy with at least some of the decisions. For example, if Alice and Bob are to follow their dinner with a movie, then maybe Bob will be willing to eat Italian food for dinner if he gets to pick his favorite movie, and maybe Alice will agree to this compromise.

Note that this setting generalizes the classic private goods setting, because in this special case we can view each public decision as the allocation of a single good. While envy is a compelling notion in the private goods setting, it makes less sense for public decisions. In our example, irrespective of where Alice and Bob go for dinner, because they are eating the same food, it is not clear what it would mean for Alice to envy Bob. If she could somehow trade places with Bob, she would still be sitting at the other end of the dinner table, eating the same food, and not be any better off. Thankfully, proportionality still has a sensible interpretation: Each player should get at least a $1/n$ fraction of the utility she would get if her most desired alternative was chosen for each decision. Unfortunately, as with envy-freeness, proportionality cannot always be guaranteed. Therefore in this work we consider relaxations of proportionality in order to arrive at fairness notions that can be guaranteed.

1.1 Our Results

Formally, a public decision making problem consists of m issues, where each issue has several associated alternatives. Each of n players has a utility for each alternative of each issue. Making a decision on an issue amounts to choosing one of the alternatives associated with the issue, and choosing an overall outcome requires making a decision on each issue simultaneously. The utility to a player for an outcome is the sum of her utilities for the alternatives chosen for different issues. This is a very simple setting, but one in which the problem of fairness is already non-trivial.

We propose relaxations of proportionality in two directions. The first, *proportionality up to one issue (Prop1)*, is similar in spirit to EF1, stating that a player should be able to get her proportional share if she gets to change the outcome of a single issue in her favor. The second direction is based on the guarantees provided by the round robin mechanism. This mechanism first orders the set of players, and then repeatedly goes through the ordering, allowing each player to make her favorite decision on any single issue, until decisions are made on all the issues. Our first relaxation in this direction, the *round robin share (RRS)*, guarantees each player the utility that she would have received under the round robin mechanism if she were the last player in the ordering. Note that the round robin mechanism lets each player make decisions on roughly the same number of issues. A further relaxation in this direction, the *pessimistic proportional share (PPS)*, guarantees each player the utility that she would get if her favorite alternatives were chosen for (approximately) a $1/n$ fraction of the issues, where these issues are chosen adversarially.

We examine the possibility and computational complexity of satisfying combinations of these fairness desiderata. We first observe that the round robin mechanism satisfies both Prop1 and RRS (and thus PPS). However, it fails to satisfy even the most basic efficiency property, Pareto optimality (PO), which requires that no other outcome should be able to make a player strictly better off without making at least one player strictly worse off.

When insisting on Pareto optimality, we observe that the leximin mechanism – informally, it chooses the outcome that maximizes the minimum utility to any player – satisfies RRS (therefore PPS) and PO via a simple argument. However, this argument does not extend to establishing Prop1, although we show that RRS implies a $1/2$ approximation to Prop1. To that end, we prove that the maximum Nash welfare (MNW) solution – informally, it chooses the outcome that maximizes

	PO	PPS	RRS	Prop1
MNW, Private goods	✓	✓ (Theorem 3.10)	$\frac{1}{2}$ (Theorem 3.10)	✓ (Theorem 3.4)
MNW, Public decisions	✓ (Theorem 3.4)	$\frac{1}{n}$ (Theorem 3.8)	$\frac{1}{n}$ (Theorem 3.8)	✓ (Theorem 3.4)
Leximin Mechanism	✓	✓ (Theorem 3.2)	✓ (Theorem 3.2)	$\frac{1}{2}$ (Theorem 3.2)
Round Robin Method	×	✓ (Theorem 3.1)	✓ (Theorem 3.1)	✓ (Theorem 3.1)

Table 1. Axioms satisfied or approximated by the mechanisms we consider. The MNW solution is split into private goods and general decisions because we obtain significantly stronger results for private goods. Results for the leximin mechanism and the round robin method apply equally to private goods and public decisions. The approximation results are lower bounds; we omit the upper bounds from the table for simplicity.

the *product* of utilities to players — that is known for its many desirable fairness properties in dividing private goods [14] satisfies Prop1 and PO in our public decision making framework, and simultaneously provides a $1/n$ approximation to both RRS and PPS. We also show that this approximation is tight up to a factor of $O(\log n)$. For division of private goods, these approximations improve significantly: the MNW solution completely satisfies PPS, and provides an $n/(2n-1) > 1/2$ approximation (but not better than $2/3$ approximation in the worst case) to RRS. Table 1 provides a summary of these results.

However, both the MNW outcome and the leximin outcome are \mathcal{NP} -hard to compute. It is therefore natural to consider whether our fairness properties can be achieved in conjunction with PO in polynomial time. For public decision making, the answer turns out to be negative for PPS and RRS, assuming $\mathcal{P} \neq \mathcal{NP}$. For division of private goods, however, we show that there exists a polynomial time algorithm that satisfies PPS and PO.

1.2 Related Work

Two classic fair division mechanisms — the leximin mechanism and the maximum Nash welfare (MNW) solution — play an important role in this paper. Both mechanisms have been extensively studied in the literature on private goods division. In particular, Kurokawa et al. [23] (Section 3.2) show that the leximin mechanism satisfies envy-freeness, proportionality, Pareto optimality, and a strong game-theoretic notion called group strategyproofness, which prevents even groups of players from manipulating the outcome to their benefit by misrepresenting their preferences, in a broad fair division domain with private goods and a specific form of non-additive utilities. On the other hand, the MNW solution has been well studied in the realm of additive utilities [14, 31]. For divisible goods, the MNW solution coincides with another well-known solution concept called *competitive equilibrium from equal incomes* (CEEI) [34], which also admits an approximate version for indivisible goods [11]. For indivisible goods, the MNW solution satisfies envy-freeness up to one good, Pareto optimality, and approximations to other fairness guarantees. One line of research aims to approximate the optimum Nash welfare [16, 25], although it is unclear if this achieves any of the appealing fairness guarantees of the MNW solution.

Our model is closely related to that of voting in combinatorial domains (see [24] for an overview). However, this literature focuses on the case where there is dependency between decisions on different issues. In contrast, our model remains interesting even though the issues are independent, and incorporating dependency is an interesting future direction. Although there is a range of work in the voting literature that focuses on fairness [2, 6, 15, 27], especially in the context of *representation* in multi-winner elections, it focuses on ordinal, rather than cardinal, preferences.¹

¹That said, there is a recent line of work on *implicit utilitarian voting* that attempts to maximize an objective with respect to the cardinal utilities underlying the ordinal preferences [7, 29], and is therefore closer to our work.

Another difference is that fairness concepts in voting apply most naturally when $n \gg m$, whereas our notions of fairness are most interesting when $m \geq n$.

Our work is also reminiscent of the *participatory budgeting* problem [12], in which there are multiple public projects with different costs, and a set of projects need to be chosen based on preferences of the participants over the projects, subject to a budget constraint. Recently, researchers in computational social choice have addressed this problem from an axiomatic viewpoint [20], including fairness considerations [17], and from the viewpoint of *implicit utilitarian voting* [5]. However, they assume access only to ordinal preferences (that may stem from underlying cardinal utilities), while we assume a direct access to cardinal utilities, as is common in the fair division literature. Also, we do not have a budget constraint that binds the outcomes on different issues.

2 MODEL

For $k \in \mathbb{N}$, define $[k] \triangleq \{1, \dots, k\}$. Before we introduce the problem we study in this paper, let us review the standard fair division setting with private goods.

Private goods division. A private goods division problem consists of a set of *players* $N = [n]$ and a set of m *goods* M . Each player $i \in N$ is endowed with a utility function $u_i : M \rightarrow \mathbb{R}_+$ such that $u_i(g)$ denotes the value player i derives from good $g \in M$. A standard assumption in the literature is that of additive valuations, i.e., (slightly abusing the notation) $u_i(S) = \sum_{g \in S} u_i(g)$ for $S \subseteq M$. An *allocation* \mathbf{A} is a partition of the set of goods among the set of players, where A_i denotes the bundle of goods received by player i . Importantly, players only derive utility from the goods they receive, i.e., the goods private to them. The utility of player i under allocation \mathbf{A} is $u_i(\mathbf{A}) = u_i(A_i)$.

Public decision making. A *public decision making* problem also has a set of players $N = [n]$, but instead of private goods, it has a set of *issues* $T = [m]$. Each issue $t \in T$ has an associated set of alternatives $A^t = \{a_1^t, \dots, a_{k_t}^t\}$, exactly one of which must be chosen. Each player i is endowed with a utility function $u_i^t : A^t \rightarrow \mathbb{R}_+$ for each issue t , and derives utility $u_i^t(a_j^t)$ if alternative a_j^t is chosen for issue t . In contrast to private goods division, a single alternative can provide positive utility to multiple players.

An outcome $\mathbf{c} = (c_1, \dots, c_m)$ of a public decision making problem is a choice of an alternative for every issue, i.e., it consists of an outcome $c_t \in A^t$ for each issue $t \in T$. Let C denote the space of possible outcomes. Slightly abusing the notation, let $u_i^t(\mathbf{c}) = u_i^t(c_t)$ be the utility player i derives from the outcome of issue t . We also assume additive valuations: let $u_i(\mathbf{c}) = \sum_{t \in T} u_i^t(\mathbf{c})$ be the utility player i derives from outcome \mathbf{c} .

In this work, we study deterministic outcomes, and in Section 5, discuss the implications when randomized outcomes are allowed. Further, we study the *offline* problem in which we are presented with the entire problem up front, and need to choose the outcomes on all issues simultaneously. One can also define an online version of the problem [19], in which we must commit to the outcome of issue t before observing issues t' with $t' > t$ for all t , but we do not consider that version here.

Private goods versus public decisions. To see why public decision making generalizes private goods division, take an instance of private goods division, and create an instance of public decision making as follows. Create an issue t_g for each good g . Let there be n alternatives in A^{t_g} , where alternative $a_i^{t_g}$ gives player i utility $u_i(g)$ while giving zero utility to all other players. It is easy to see that choosing alternative $a_i^{t_g}$ is equivalent to allocating good g to player i . Hence, the constructed public decision making problem effectively mimics the underlying private goods division problem.

2.1 Efficiency and Fairness

In this paper, we not only adapt classical notions of efficiency and fairness defined for private goods division to our public decision making problem, but also introduce three fairness axioms that are novel for both public decision making and private goods division. First, we need additional notation that we will use throughout the paper.

Let $p \triangleq \lfloor m/n \rfloor$. For issue $t \in T$ and player $i \in N$, let $a_{max}^t(i) \in \arg \max_{a \in A^t} \{u_i^t(a)\}$ and $u_{max}^t(i) = u_i^t(a_{max}^t(i))$. That is, $a_{max}^t(i)$ is an alternative that gives player i the most utility for issue t , and $u_{max}^t(i)$ is the utility player i derives from $a_{max}^t(i)$. Let the sequence $\langle u_{max}^{(1)}(i), \dots, u_{max}^{(m)}(i) \rangle$ represent the maximum utilities player i can derive from different issues, sorted in a non-ascending order. Hence, $\{u_{max}^{(k)}(i)\}_{k \in [m]} = \{u_{max}^t(i)\}_{t \in T}$ and $u_{max}^{(k)}(i) \geq u_{max}^{(k+1)}(i)$ for $k \in [m-1]$.

Efficiency. In this paper, we focus on a popular notion of economic efficiency. We say that an outcome \mathbf{c} is *Pareto optimal* (PO) if there does not exist another outcome \mathbf{c}' that can provide at least as much utility as \mathbf{c} to every player, i.e., $u_i(\mathbf{c}') \geq u_i(\mathbf{c})$ for all $i \in N$, and strictly more utility than \mathbf{c} to some player, i.e., $u_{i^*}(\mathbf{c}') > u_{i^*}(\mathbf{c})$ for some $i^* \in N$.

Fairness. For private goods division, perhaps the most prominent notion of fairness is envy-freeness [18]. An allocation \mathbf{A} is called *envy-free* (EF) if every player values her bundle at least as much as she values any other player's bundle, i.e., $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$. Because envy-freeness cannot in general be guaranteed, prior work also focuses on its relaxations. For instance, an allocation \mathbf{A} is called *envy-free up to one good* (EF1) if no player envies another player after removing at most one good from the latter player's bundle, i.e., for all $i, j \in N$, either $u_i(A_i) \geq u_i(A_j)$ or $\exists g_j \in A_j$ such that $u_i(A_i) \geq u_i(A_j \setminus \{g_j\})$.

Unfortunately, as argued in Section 1, the notion of envy is not well defined for public decisions. Hence, for public decision making, we focus on another fairness axiom, Proportionality, and its relaxations. For private goods division, proportionality is implied by envy-freeness.²

Proportionality (Prop). At a high level, proportionality requires that each player must receive at least her "proportional share", which is a $1/n$ fraction of the utility she would derive if she could act as the dictator. For a public decision making problem, the *proportional share* of player i (Prop_i) is $1/n$ times the sum of the maximum utilities the player can derive across all issues, i.e.,

$$\text{Prop}_i = \frac{1}{n} \sum_{t \in T} u_{max}^t(i).$$

For $\alpha \in (0, 1]$, we say that an outcome \mathbf{c} satisfies α -*proportionality* (α -Prop) if $u_i(\mathbf{c}) \geq \alpha \cdot \text{Prop}_i$ for all players $i \in N$. We refer to 1-Prop simply as Prop.

Proportionality up to one issue (Prop1). We introduce a novel relaxation of proportionality (more generally, of α -proportionality) in the same spirit as envy-freeness up to one good, which is a relaxation of envy-freeness. For $\alpha \in (0, 1]$, we say that an outcome \mathbf{c} satisfies α -*proportionality up to one issue* (α -Prop1) if for every player $i \in N$, there exists an issue $t \in T$ such that, ceteris paribus, changing the outcome of t from c_t to $a_{max}^t(i)$ ensures that player i achieves an α fraction of her proportional share, i.e., if

$$\forall i \in N \exists t \in T \text{ s.t. } u_i(\mathbf{c}) - u_i^t(\mathbf{c}) + u_{max}^t(i) \geq \alpha \cdot \text{Prop}_i.$$

We refer to 1-Prop1 simply as Prop1.

Round robin share (RRS). Next, we introduce another novel fairness axiom that is motivated from the classic round robin method that, for private goods, lets players take turns and in each turn, pick a single most favorite item left unclaimed. For public decision making, we instead let

²This assumes non-wastefulness, i.e., that all goods are allocated. We make this assumption throughout the paper.

players make a decision on a single issue in each turn. The utility guaranteed to the players by this approach is captured by the following fairness axiom.

Recall that the sequence $\langle u_{max}^{(1)}(i), \dots, u_{max}^{(m)}(i) \rangle$ represents the maximum utility player i can derive from different issues, sorted in a non-ascending order. Then, we define the *round robin share* of player i (RRS_i) as

$$RRS_i = \sum_{k=1}^p u_{max}^{(k \cdot n)}(i).$$

This is player i 's utility from the round robin method, if she is last in the ordering and all issues she does not control give her utility 0. For $\alpha \in (0, 1]$, we say that an outcome \mathbf{c} satisfies α -*round robin share* (α -RRS) if $u_i(\mathbf{c}) \geq \alpha \cdot RRS_i$ for all players $i \in N$. Again, we refer to 1-RRS simply as RRS.

Pessimistic proportional share (PPS). We introduce another novel fairness axiom that is a further relaxation of round robin share. Note that the round robin method, by letting players make a decision on a single issue per turn, allows each player to make decisions on at least $p = \lfloor m/n \rfloor$ issues. The following axiom captures the utility players would be guaranteed if each player still made decisions on a "proportional share" of p issues, but if these issues were chosen pessimistically.

We define the *pessimistic proportional share* of player i (PPS_i) to be the sum of the maximum utilities the player can derive from a set of p issues, chosen adversarially to minimize this sum:

$$PPS_i = \sum_{k=m-p+1}^m u_{max}^{(k)}(i).$$

For $\alpha \in (0, 1]$, we say that an outcome \mathbf{c} satisfies α -*pessimistic proportional share* (α -PPS) if $u_i(\mathbf{c}) \geq \alpha \cdot PPS_i$ for all players $i \in N$. Again, we refer to 1-PPS simply as PPS.

Connections among fairness properties. Trivially, proportionality (Prop) implies proportionality up to one issue (Prop1). In addition, it can also be checked that the following sequence of logical implications holds: Prop \implies MMS \implies RRS \implies PPS.

Here, MMS is the *maximin share guarantee* [11, 30]. Adapting the definition naturally from private goods division to public decision making, the maximin share of a player is the utility the player can guarantee herself by dividing the set of issues into n bundles, if she gets to make the decisions best for her on the issues in an adversarially chosen bundle. The maximin share (MMS) guarantee requires that each player must receive utility that is at least her maximin share. We do not focus on the maximin share guarantee in this paper.

2.2 Mechanisms

A mechanism for a public decision making problem (resp. a private goods division problem) maps each input instance of the problem to an outcome (resp. an allocation). We say that a mechanism satisfies a fairness or efficiency property if it always returns an outcome satisfying the property. There are three prominent mechanisms that play a key role in this paper.

Round robin method. As mentioned earlier, the round robin method first fixes an ordering of the players. Then the players take turns choosing their most preferred alternative on a single issue of their choice whose outcome has not yet been determined.

The leximin mechanism. The leximin mechanism chooses an outcome which maximizes the utility of the worst off player, i.e., $\min_{i \in N} u_i(\mathbf{c})$. Subject to this constraint, it maximizes the utility of the second least well off player, and so on. Note that the leximin mechanism is trivially Pareto optimal because if it were possible to improve some player's utility without reducing that of any other, it would improve the objective that the leximin mechanism optimizes.

Maximum Nash welfare (MNW). The *Nash welfare* of an outcome \mathbf{c} is the product of utilities to all players under \mathbf{c} : $NW(\mathbf{c}) = \prod_{i \in N} u_i(\mathbf{c})$. When there exists an outcome \mathbf{c} with $NW(\mathbf{c}) > 0$, then the MNW solution chooses an arbitrary outcome \mathbf{c} that maximizes the Nash welfare. When all outcomes have zero Nash welfare, it finds a largest cardinality set S of players that can be given non-zero utility, and selects an outcome maximizing the product of their utilities, i.e., $\prod_{i \in S} u_i(\mathbf{c})$.

2.3 Examples

We illustrate the fairness properties through two examples.

Example 2.1. Consider a public decision making problem with two players ($N = [2]$) and two issues ($T = [2]$). Each issue has two alternatives ($|A^1| = |A^2| = 2$). The utilities of the two players for the two alternatives in both issues are as follows.

$$\begin{array}{c|cc} & a_1^t & a_2^t \\ \hline u_1^t & 1 & 0 \\ u_2^t & 0 & 1 \end{array} \quad \text{for } t \in [2].$$

The various fair shares of the two players are $\text{Prop}_1 = \text{RRS}_1 = \text{PPS}_1 = \text{Prop}_2 = \text{RRS}_2 = \text{PPS}_2 = 1$. Now, outcome $\mathbf{c} = (a_1^1, a_2^2)$ gives utilities $u_1(\mathbf{c}) = 2$ and $u_2(\mathbf{c}) = 0$, and therefore violates Prop, RRS, and PPS. It satisfies Prop1 because switching the decision on either issue in favor of player 2 makes her achieve her proportional share. On the other hand, outcome $\mathbf{c} = (a_1^1, a_2^1)$ gives utility 1 to both players, and thus satisfies Prop (as well as Prop1, RRS, and PPS, which are relaxations of Prop).

Example 2.2. Consider a public decision making problem with two players ($N = [2]$) and eight issues ($T = [8]$). Once again, each issue has two alternatives, for which the utilities of the two players are as follows.

$$\begin{array}{c|cc} & a_1^t & a_2^t \\ \hline u_1^t & 1 & 0 \\ u_2^t & 0 & 1 \end{array} \quad \text{for } t \in \{1, 2, 3, 4\}, \quad \text{and} \quad \begin{array}{c|cc} & a_1^t & a_2^t \\ \hline u_1^t & 1 & 0 \\ u_2^t & 0 & 0 \end{array} \quad \text{for } t \in \{5, 6, 7, 8\}.$$

In this case, we have $\text{Prop}_1 = \text{RRS}_1 = \text{PPS}_1 = 4$, whereas $\text{Prop}_2 = \text{RRS}_2 = 2$ and $\text{PPS}_2 = 0$. Consider outcome $\mathbf{c} = (a_1^1, a_2^2, a_3^3, a_4^4, a_5^5, a_6^6, a_7^7, a_8^8)$. Then, we have $u_1(\mathbf{c}) = 8$ while $u_2(\mathbf{c}) = 0$, which satisfies PPS but violates RRS. Further, \mathbf{c} also violates Prop1 because switching the outcome of any single issue can only give player 2 utility at most 1, which is less than $\text{Prop}_2 = 2$. On the other hand, outcome $\mathbf{c} = (a_2^2, a_2^2, a_3^3, a_4^4, a_5^5, a_6^6, a_7^7, a_8^8)$ achieves $u_1(\mathbf{c}) = u_2(\mathbf{c}) = 4$, and satisfies Prop (and thus its relaxations Prop1, RRS, and PPS).

3 (APPROXIMATE) SATISFIABILITY OF AXIOMS

If we are willing to sacrifice Pareto optimality, then we can easily achieve both RRS (and therefore PPS) and Prop1 simultaneously with the round robin mechanism. This is not a surprising result. RRS is defined based on the guarantee provided by the round robin mechanism, and PPS is a relaxation of RRS. The round robin mechanism is also known to satisfy EF1 for private goods division, which is similar in spirit to Prop1.

THEOREM 3.1. *The round robin mechanism satisfies RRS (and therefore PPS) and Prop1, and runs in polynomial time.*

PROOF. The round robin mechanism clearly runs in polynomial time (note that it is easy for a player to choose the next issue on which to determine the outcome). To see why it satisfies RRS, note that the mechanism allows every player to make a decision on one issue once every n turns. Thus, for each $k \in [p]$, every player gets to make decisions on at least k of her “top” $k \cdot n$ issues,

when issues are sorted in the descending order of the utility her favorite alternative in the issue gives her. It is easy to see that this implies every player i gets utility at least RRS_i . Because RRS implies PPS, the mechanism also satisfies PPS. It remains to show that it satisfies Prop1 as well.

Fix a player i and let \mathbf{c} be the outcome produced by the round robin mechanism for some choosing order of the players. Because the round robin mechanism satisfies RRS, player i gets utility at least

$$u_i(\mathbf{c}) \geq \sum_{k=1}^p u_{\max}^{(k \cdot n)}(i).$$

For $k \in [m]$, let the k^{th} favorite issue of player i be the issue t for which $u_{\max}^t(i)$ is the k^{th} highest. Let $\ell \in \mathbb{N} \cup \{0\}$ be the largest index such that for every $k \in [\ell]$, outcome \mathbf{c} chooses player i 's most preferred alternative on her k^{th} favorite issue. Let t^* be her $(\ell + 1)^{\text{th}}$ favorite issue. To show that \mathbf{c} satisfies Prop1, we construct outcome \mathbf{c}' from \mathbf{c} by only changing the outcome of issue t^* to $a_{\max}^{t^*}(i)$, and show that $u_i(\mathbf{c}') \geq \text{Prop}_i$. Note that if $\ell \geq p$, then

$$u_i(\mathbf{c}') \geq \sum_{k=1}^{\ell+1} u_{\max}^{(k)}(i) \geq \sum_{k=1}^{p+1} u_{\max}^{(k)}(i) \geq \frac{1}{n} \sum_{k=1}^m u_{\max}^{(k)}(i) = \text{Prop}_i.$$

Let $\ell < p$. Then, using the fact that the round robin mechanism lets player i choose her most preferred alternative for at least k of her favorite $k \cdot n$ issues for every $k \leq p$ (and her $(\ell + 1)^{\text{th}}$ favorite issue was not one of these), we have

$$u_i(\mathbf{c}') \geq \sum_{k=1}^{\ell+1} u_{\max}^{(k)}(i) + \sum_{k=\ell+1}^p u_{\max}^{(k \cdot n)}(i) \geq \frac{1}{n} \sum_{k=1}^{(\ell+1) \cdot n} u_{\max}^{(k)}(i) + \frac{1}{n} \sum_{k=(\ell+1)n+1}^m u_{\max}^{(k)}(i) = \text{Prop}_i.$$

Therefore, the round robin mechanism satisfies Prop1. ■

While this result seems to reflect favorably upon the round robin mechanism, recall that it violates Pareto optimality even for private goods division. For public decision making, a simple reason for this is that the round robin mechanism, for each issue, chooses an alternative that is some player's favorite, while it could be unanimously better to choose compromise solutions that make many players happy. Imagine there are two players and two issues, each with two alternatives. The "extreme" alternative in each issue $i \in \{1, 2\}$ gives utility 1 to player i but 0 to the other, while the "compromise" alternative in each issue $i \in \{1, 2\}$ gives utility $2/3$ to both players. It is clear that both players prefer choosing the compromise alternative in both issues to choosing the extreme alternative in both issues. Because such "Pareto improvements" which make some players happier without making any player worse off are unanimously preferred by the players, the round robin outcome becomes highly undesirable. We therefore seek mechanisms that provide fairness guarantees while satisfying Pareto optimality.

A natural question is whether there exists a mechanism that satisfies RRS, Prop1, and PO. An obvious approach is to start from an outcome that already satisfies RRS and Prop1 (e.g., the round robin outcome), and make Pareto improvements until no such improvements are possible. While Pareto improvements preserve RRS as the utilities to the players do not decrease, Prop1 can be lost as it depends on the exact alternatives chosen and not only on the utilities to the players. We leave it as an important open question to determine if RRS, Prop1, and PO can be satisfied simultaneously.

We therefore consider satisfying each fairness guarantee individually with PO. One can easily find an outcome satisfying RRS and PO by following the aforementioned approach of starting with an outcome satisfying RRS, and making Pareto improvements while possible. There is also a more direct approach to satisfying RRS and PO. Recall that the leximin mechanism chooses the outcome which maximizes the minimum utility to any player, subject to that maximizes the second

minimum utility, and so on. It is easy to see that this mechanism is always Pareto optimal. Now, let us normalize the utilities of all players such that $RRS_i = 1$ for every player $i \in N$.³ Because the round robin mechanism gives every player i utility at least $RRS_i = 1$, it must be the case that the leximin mechanism operating on these normalized utilities must also give every player utility at least 1, and thus produce an outcome that is both RRS and PO.

THEOREM 3.2. *The leximin mechanism satisfies RRS, PO, and (1/2)-Prop1.*

That leximin satisfies (1/2)-Prop1 follows directly from the following lemma, and noting that leximin satisfies RRS.

LEMMA 3.3. *RRS implies (1/2)-Prop1.*

PROOF. Note that $RRS_i = \sum_{k=1}^p u_{max}^{(k \cdot n)}(i) \geq \frac{1}{n} \sum_{t=n+1}^m u_{max}^{(t)}(i)$ and $u_{max}^{(1)}(i) \geq \frac{1}{n} \sum_{t=1}^n u_{max}^{(t)}(i)$. Summing the two equations, we get

$$RRS_i + u_{max}^{(1)}(i) \geq \frac{1}{n} \sum_{t=1}^m u_{max}^{(t)}(i) = \text{Prop}_i.$$

Therefore, $\max\{RRS_i, u_{max}^{(1)}(i)\} \geq \frac{1}{2} \text{Prop}_i$.

Suppose that $u_i(\mathbf{c}) \geq RRS_i$ for some outcome \mathbf{c} . Then either i already receives her most valued item, in which case she receives utility at least $\max\{RRS_i, u_{max}^{(1)}(i)\} \geq \frac{1}{2} \text{Prop}_i$, or she does not receive her most valued item. If she does not, then after giving it to her, she receives utility at least $\frac{1}{2} \text{Prop}_i$. Therefore, \mathbf{c} satisfies (1/2)-Prop1. ■

Next, we study whether we can achieve Prop1 and PO simultaneously. Neither of the previous approaches seems to work: we already argued that following Pareto improvements could lose Prop1, and the normalization trick is difficult to apply because Prop1 is not defined in terms of any fixed share of utility.

One starting point to achieving Prop1 and PO is the maximum Nash welfare (MNW) solution, which, for private goods division, is known to satisfy the similar guarantee of EF1 and PO [14]. It turns out that the MNW solution is precisely what we need.

THEOREM 3.4. *The MNW solution satisfies proportionality up to one issue (Prop1) and Pareto optimality (PO).*

Before we prove this, we need a folklore result, which essentially states that if the sum of n terms is to be reduced by a fixed quantity δ that is less than each term, then their product reduces the most when δ is taken out of the lowest term. The following lemma proves this result when all initial terms are 1, which is sufficient for our purpose. The proof of the lemma appears in the full version of the paper.⁴

LEMMA 3.5. *Let $\{x_1, \dots, x_n\}$ be a set of n non-negative real numbers such that $\sum_{i=1}^n \max\{0, 1 - x_i\} \leq \delta$, where $0 < \delta < 1$. Then, $\prod_{i=1}^n x_i \geq 1 - \delta$.*

PROOF OF THEOREM 3.4. Fix an instance of the public decision making problem. Let $S \subseteq N$ be the set of players that the MNW outcome \mathbf{c} gives positive utility to. Then, by the definition of the MNW outcome, S must be a largest set of players that can simultaneously be given positive utility, and \mathbf{c} must maximize the product of utilities of players in S .

First, we show that \mathbf{c} is PO. Note that a Pareto improvement over \mathbf{c} must either give a positive utility to a player in $N \setminus S$ or give more utility to a player in S , without reducing the utility to any

³Players with zero round robin share can be incorporated via a simple extension to the argument.

⁴The full version is available at the authors' websites.

player in S . This is a contradiction because it violates either optimality of the size of S or optimality of the product of utilities of players in S . Hence, MNW satisfies PO.

We now show that MNW also satisfies Prop1. Suppose for contradiction that Prop1 is violated for player i under \mathbf{c} . First, note that we must have $\text{Prop}_i > 0$. Further, it must be the case that $u_{max}^t(i) > 0$ for at least $n + 1$ issues. Were this not the case, Prop1 would be trivially satisfied for player i since we can give her utility

$$u_{max}^{(1)}(i) \geq \frac{1}{n} \sum_{t=1}^n u_{max}^{(t)}(i) = \frac{1}{n} \sum_{t=1}^m u_{max}^{(t)}(i) = \text{Prop}_i$$

by changing the outcome on a single issue.

We now show that $u_i(\mathbf{c}) > 0$ (i.e., $i \in S$). For contradiction, suppose otherwise. For each of the (at least) $n + 1$ issues with $u_{max}^t(i) > 0$, there must exist another player $j \neq i$ that gets positive utility *only* from that issue under \mathbf{c} (otherwise we could use that issue to give positive utility to i while not reducing any other agents' utility to zero, contradicting the maximality of S). But this is impossible, since there are at least $n + 1$ issues and only $n - 1$ agents (other than i).

Because MNW outcomes and the Prop1 property are invariant to individual scaling of utilities, let us scale the utilities such that $\text{Prop}_i = 1$ and $u_j(\mathbf{c}) = 1$ for all $j \in S \setminus \{i\}$. Select issue $t^* \in T$ as

$$t^* \in \arg \min_{t \in T} \frac{\sum_{j \in N \setminus \{i\}} u_j^t(\mathbf{c})}{u_{max}^t(i) - u_i^t(\mathbf{c})}.$$

Note that t^* is well defined because $u_{max}^t(i) > u_i^t(\mathbf{c})$ for at least one $t \in T$, otherwise Prop1 would not be violated for player i .

We now show that outcome \mathbf{c}' such that $c_{t^*}' = a_{max}^{t^*}(i)$ and $c_t' = c_t$ for all $t \in T \setminus \{t^*\}$ achieves strictly greater product of utilities of players in S than outcome \mathbf{c} does, which is a contradiction as \mathbf{c} is an MNW outcome. First, note that

$$\frac{\sum_{j \in N \setminus \{i\}} u_j^{t^*}(\mathbf{c})}{u_{max}^{t^*}(i) - u_i^{t^*}(\mathbf{c})} \leq \frac{\sum_{t \in T} \sum_{j \in N \setminus \{i\}} u_j^t(\mathbf{c})}{\sum_{t \in T} (u_{max}^t(i) - u_i^t(\mathbf{c}))} = \frac{\sum_{j \in N \setminus \{i\}} u_j(\mathbf{c})}{n \text{Prop}_i - u_i(\mathbf{c})} \leq \frac{(n-1)}{(n-1) \text{Prop}_i} = 1, \quad (1)$$

where the penultimate transition follows because we normalized utilities to achieve $u_j(\mathbf{c}) = 1$ for every $j \in S \setminus \{i\}$, every $j \in N \setminus S$ satisfies $u_j(\mathbf{c}) = 0$, and player i does not receive her proportional share. The final transition holds due to our normalization $\text{Prop}_i = 1$.

Let $\delta = \sum_{j \in S \setminus \{i\}} u_j^{t^*}(\mathbf{c})$. Then, Equation (1) implies $u_i(\mathbf{c}') - u_i(\mathbf{c}) = u_{max}^{t^*}(i) - u_i^{t^*}(\mathbf{c}) \geq \delta$. Thus,

$$u_i(\mathbf{c}) + \delta \leq u_i(\mathbf{c}') < 1, \quad (2)$$

where the last inequality follows because the original outcome \mathbf{c} violated Prop1 for player i . In particular, this implies $\delta < 1$. Our goal is to show that $\prod_{j \in S} u_j(\mathbf{c}') > u_i(\mathbf{c}) = \prod_{j \in S} u_j(\mathbf{c})$, where the last equality holds due to our normalization $u_j(\mathbf{c}) = 1$ for $j \in S \setminus \{i\}$ and because $i \in S$. This would be a contradiction because \mathbf{c} maximizes the product of utilities of players in S . Now,

$$\sum_{j \in S \setminus \{i\}} \max\{0, 1 - u_j(\mathbf{c}')\} = \sum_{j \in S \setminus \{i\}} \max\{0, u_j^{t^*}(\mathbf{c}) - u_j^{t^*}(\mathbf{c}')\} \leq \sum_{j \in S \setminus \{i\}} u_j^{t^*}(\mathbf{c}) = \delta,$$

where the first transition follows from setting $1 = u_j(\mathbf{c})$ (by our normalization) and noting that \mathbf{c} and \mathbf{c}' are identical for all issues except t^* , and the second because all utilities are non-negative.

Hence, Lemma 3.5 implies that $\prod_{j \in S \setminus \{i\}} u_j(\mathbf{c}') \geq 1 - \delta$. Thus,

$$\prod_{j \in S} u_j(\mathbf{c}') \geq (1 - \delta) \cdot (u_i(\mathbf{c}) + \delta) = u_i(\mathbf{c}) + \delta \cdot (1 - u_i(\mathbf{c})) - \delta^2 > u_i(\mathbf{c}) + \delta^2 - \delta^2 = u_i(\mathbf{c}),$$

where the inequality holds because $u_i(\mathbf{c}) + \delta < 1$ from Equation (2). ■

For private goods division, this result can be derived in a simpler fashion. Caragiannis et al. [14] already show that MNW satisfies PO. In addition, they also show that MNW satisfies EF1, which implies Prop1 due to our next result. To be consistent with the goods division literature, we use proportionality up to one good (rather than one issue) in the private goods division context.

LEMMA 3.6. *For private goods division, envy-freeness up to one good (EF1) implies proportionality up to one good (Prop1).*

PROOF. Take an instance of private goods division with a set of players N and a set of goods M . Let \mathbf{A} be an allocation satisfying EF1. Fix a player $i \in N$.

Due to the definition of EF1, there must exist⁵ a set of goods $X = \{g_j\}_{j \in N \setminus \{i\}}$ such that $u_i(A_i) \geq u_i(A_j) - u_i(g_j)$ for every $j \in N \setminus \{i\}$. Summing over all $j \in N \setminus \{i\}$, we get

$$\begin{aligned} (n-1) \cdot u_i(A_i) &\geq \left(\sum_{j \in N \setminus \{i\}} u_i(A_j) \right) - u_i(X) \implies n \cdot u_i(A_i) \geq u_i(M) - u_i(X) \\ &\implies u_i(A_i) + \frac{u_i(X)}{n} \geq \frac{u_i(M)}{n}. \end{aligned} \quad (3)$$

Note that X has less than n goods. Suppose player i receives good $g^* \in \arg \max_{g \in X} u_i(g)$. Note that $g^* \notin A_i$. Then, we have $u_i(A_i \cup \{g^*\}) \geq u_i(M)/n = \text{Prop}_i$, which implies that Prop1 is satisfied with respect to player i . Because player i was chosen arbitrarily, we have that EF1 implies Prop1. ■

Equation 3 in the proof of Lemma 3.6 directly implies the following lemma because the set X in the equation contains at most $n-1$ goods.

LEMMA 3.7. *Let \mathbf{A} be an allocation of private goods that satisfies EF1. Then, for every player i ,*

$$u_i(A_i) \geq \text{Prop}_i - \frac{1}{n} \sum_{t=1}^{n-1} u_{\max}^{(t)}(i),$$

where $u_{\max}^{(t)}(i)$ is the utility player i derives from her t^{th} most valued good.

Next, we turn our attention to RRS and PPS. While MNW does not satisfy either of them, it approximates both.

THEOREM 3.8. *The MNW solution satisfies 1/n-RRS (and therefore 1/n-PPS). The approximation is tight for both RRS and PPS up to a factor of $O(\log n)$.*

PROOF. We first show the lower bound. Fix an instance of public decision making, and let \mathbf{c} denote an MNW outcome. Let $S \subseteq N$ denote the set of players that achieve positive utility under \mathbf{c} .

Without loss of generality, let us normalize the utilities such that $u_j(\mathbf{c}) = 1$ for every $j \in S$. Suppose for contradiction that for some player i , $u_i(\mathbf{c}) < (1/n) \cdot \text{RRS}_i$. First, this implies that $\text{RRS}_i > 0$, which in turn implies that player i must be able to derive a positive utility from at least n different issues. By an argument identical to that used to argue that $u_i(\mathbf{c}) > 0$ in the proof of Theorem 3.4, it can be shown that we must also have $u_i(\mathbf{c}) > 0$ in this case (i.e., $i \in S$).

Now, recall that the sequence $\langle u_{\max}^{(1)}(i), \dots, u_{\max}^{(m)}(i) \rangle$ contains the maximum utility player i can derive from different issues, sorted in a non-ascending order. For every $q \in [p]$, let

$$t_q = \arg \min_{(q-1)n+1 \leq t \leq qn} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}).$$

⁵If $A_j = \emptyset$, we can add a dummy good g_j that every player has utility 0 for, and make $A_j = \{g_j\}$.

That is, we divide the public decision making into sets of n issues, grouped by player i 's maximum utility for them, and for each set of issues, we let t_q be the one that the remaining players derive the lowest total utility from. Note that $t_q \leq qn$ for each $q \in [p]$, and therefore $u_{max}^{(t_q)}(i) \geq u_{max}^{(qn)}(i)$.

We will show that outcome \mathbf{c}' , where $c'_{t_q} = a_{max}^{t_q}(i)$ for all $q \in [p]$ and $c'_t = c_t$ for all other issues t , achieves a higher product of utilities to players in S than \mathbf{c} does, which is a contradiction because \mathbf{c} is an MNW outcome. First, note that

$$u_i(\mathbf{c}') \geq \sum_{q=1}^p u_{max}^{(t_q)}(i) \geq \sum_{k=1}^p u_{max}^{(k \cdot n)}(i) = \text{RRS}_i > n.$$

Further, we have

$$\begin{aligned} \sum_{j \in S \setminus \{i\}} \max\{0, 1 - u_j(\mathbf{c}')\} &= \sum_{j \in S \setminus \{i\}} \max\{0, u_j(\mathbf{c}) - u_j(\mathbf{c}')\} \\ &= \sum_{j \in S \setminus \{i\}} \sum_{q=1}^p \max\{0, u_j^{t_q}(\mathbf{c}) - u_j^{t_q}(\mathbf{c}')\} \leq \sum_{j \in S \setminus \{i\}} \sum_{q=1}^p u_j^{t_q}(\mathbf{c}), \end{aligned}$$

where the first equality follows from our normalization, the second because \mathbf{c} and \mathbf{c}' only differ on issues $\{t_q\}_{q \in [p]}$, and the last because all utilities are non-negative.

Reversing the order of the summation and further manipulating the expression, we have

$$\sum_{q=1}^p \sum_{j \in S \setminus \{i\}} u_j^{t_q}(\mathbf{c}) \leq \sum_{q=1}^p \frac{1}{n} \sum_{t=(q-1)n+1}^{qn} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}) = \frac{1}{n} \sum_{t=1}^{pn} \sum_{j \in S \setminus \{i\}} u_j^t(\mathbf{c}) \leq \frac{n-1}{n},$$

where the first transition follows from the definition of t_q . By Lemma 3.5, we have

$$\prod_{j \in S} u_j(\mathbf{c}') = u_i(\mathbf{c}') \prod_{j \in S \setminus \{i\}} u_j(\mathbf{c}') > n \cdot \left(1 - \frac{n-1}{n}\right) = 1 = \prod_{j=1}^n u_j(\mathbf{c}),$$

where the inequality holds because $u_i(\mathbf{c}') \geq \text{RRS}_i > n \cdot u_i(\mathbf{c}) = n$, as player i receives her round robin share under \mathbf{c}' but did not even receive a $1/n$ fraction of it under \mathbf{c} . Hence, outcome \mathbf{c}' achieves a higher product of utilities to players in S than \mathbf{c} does, which is a contradiction.

Due to space considerations, the proof of the upper bound appears in the full version. ■

For private goods, we can show that the MNW solution provides much better approximations to both RRS and PPS, as a result of its strong fairness guarantee of EF1.

LEMMA 3.9. *For private goods division, envy-freeness up to one good (EF1) implies PPS and $n/(2n-1)$ -RRS, but does not imply $n/(2n-2)$ -RRS.*

PROOF. Let \mathbf{A} be an allocation of private goods that satisfies EF1. First, we show that \mathbf{A} must also satisfy PPS. Suppose for contradiction that it violates PPS. Then, there exists a player i such that $u_i(A_i) < \text{PPS}_i$, which in turn implies that $|A_i| < p$. Because the average number of goods per player is $\frac{m}{n} \geq p$, there must exist a player j such that $|A_j| > p$. Hence, for any good $g \in A_j$, player j has at least p goods even after removing g from A_j , which implies $u_i(A_j \setminus \{g\}) \geq \text{PPS}_i > u_i(A_i)$. However, this contradicts the fact that \mathbf{A} is EF1.

We now show that \mathbf{A} also satisfies $1/2$ -RRS. By Lemma 3.7, we have

$$u_i(A_i) \geq \frac{1}{n} \sum_{t=n}^m u_{max}^{(t)}(i) \geq \frac{1}{n} u_{max}^{(n)}(i) + \sum_{k=2}^p u_{max}^{(k \cdot n)}(i). \quad (4)$$

Further, since \mathbf{A} satisfies EF1, it must be the case that

$$u_i(A_i) \geq u_{\max}^{(n)}(i). \quad (5)$$

To see this, suppose for contradiction that $u_i(A_i) < u_{\max}^{(n)}(i)$, which implies that player i is not allocated any of her n most valued goods. Therefore, by the pigeonhole principle, there must exist a player $j \in N \setminus \{i\}$ that is allocated at least two of these goods. Hence, for any $g \in A_j$, we have $u_i(A_j \setminus \{g\}) \geq u_{\max}^{(n)}(i) > u_i(A_i)$, which violates EF1. Finally, adding n times Equation (4) with $n - 1$ times Equation (5), we obtain

$$(2n - 1) \cdot u_i(A_i) \geq n \cdot u_{\max}^{(n)}(i) + n \cdot \sum_{k=2}^p u_{\max}^{(k \cdot n)}(i) = n \cdot \text{RRS}_i,$$

which implies the desired $n/(2n - 1)$ -RRS guarantee.

For the upper bound, consider an instance with n players and n^2 goods, and define player 1's utility function to be

$$u_1(g_j) = \begin{cases} 1 & 1 \leq j \leq n, \\ \frac{1}{n-1} & n+1 \leq j \leq n^2. \end{cases}$$

Note that $\text{RRS}_1 = 1 + (n - 1) \frac{1}{n-1} = 2$. Consider the allocation \mathbf{A} with $A_1 = \{g_{n+1}, \dots, g_{2n}\}$, $A_2 = \{g_1, g_2\}$, and $A_i = \{g_i, g_{(i-1)n+1}, \dots, g_{i \cdot n}\}$ for all players $i > 2$. Let the utilities of players 2 through n be positive for the goods they receive and zero for the remaining goods. Hence, they clearly do not envy any players. For player 1, we have $u_1(A_1) = \frac{n}{n-1}$, $u_1(A_2 \setminus \{g_2\}) = 1$, and $u_1(A_i \setminus \{g_i\}) = \frac{n}{n-1}$ for all $i > 2$. That is, player 1 does not envy any other player up to one good. Hence, \mathbf{A} satisfies EF1, and player 1 obtains a $\frac{n}{2n-2}$ fraction of her RRS share, as required. ■

As a corollary of Lemma 3.9, EF1 implies $1/2$ -RRS, and this approximation is asymptotically tight. Further, because the MNW solution satisfies EF1, Lemma 3.9 immediately provides guarantees (lower bounds) for the MNW solution. However, the upper bound in the proof of Lemma 3.9 does not work for the MNW solution. Next, we establish a much weaker lower bound, leaving open the possibility that the MNW solution may achieve a constant approximation better than $1/2$ to RRS.

THEOREM 3.10. *For private goods division, the MNW solution satisfies PPS and $n/(2n - 1)$ -RRS. For every $\varepsilon > 0$, the MNW solution does not satisfy $(2/3 + \varepsilon)$ -RRS.*

PROOF. The lower bounds follow directly from Lemma 3.9 and the fact that the MNW solution satisfies EF1. For the upper bound, consider an instance with two players and four goods. Player 1 has utilities $(1 - \delta, 1 - \delta, 1/2, 1/2)$ and player 2 has utilities $(1, 1, 0, 0)$ for goods (g_1, g_2, g_3, g_4) , respectively. Note that $\text{RRS}_1 = 3/2 - \delta$. The MNW allocation \mathbf{A} is given by $A_1 = \{g_3, g_4\}$ and $A_2 = \{g_1, g_2\}$. Thus, $\frac{u_1(A_1)}{\text{RRS}_1} = \frac{2}{3-2\delta}$. The upper bound follows by setting δ sufficiently small. ■

4 COMPUTATIONAL COMPLEXITY

In Section 3, we showed that without requiring Pareto optimality, we can achieve both RRS (thus PPS) and Prop1 in polynomial time using the round robin method (Theorem 3.1). In contrast, when we require PO, the leximin mechanism (with an appropriate normalization of utilities) provides RRS (thus PPS) and PO, while the MNW solution provides Prop1 and PO. However, both these solutions are \mathcal{NP} -hard to compute [4, 31]. This raises a natural question whether we can efficiently find outcomes satisfying our fairness guarantees along with PO. For PPS, the answer is negative.

THEOREM 4.1. *It is \mathcal{NP} -hard to find an outcome satisfying PPS and PO.*

Note that it is the *search problem* of finding an outcome (any outcome) satisfying PPS and PO for which we prove computational hardness; the *decision problem* of testing the existence of such an outcome is trivial as we know it always exists. Before we prove this result, we need to introduce a new (to our knowledge) decision problem and show that it is \mathcal{NP} -complete.

Exact Triple-Cover by 3-sets (X33C): An instance (Y, \mathcal{T}) of X33C is given by a set Y of r vertices and a set $\mathcal{T} = \{T_1, T_2, \dots, T_m\}$, where each T_i is a set of three vertices. The decision problem is to determine whether it is possible to choose r sets, with repetition allowed, such that every vertex v is contained in exactly three of the chosen sets (an exact triple-cover).

Let us contrast this with the definition of the popular NP-complete problem, Exact Cover by 3-sets (X3C): An instance (X, \mathcal{S}) of X3C is given by a set X of $3q$ vertices and a set $\mathcal{S} = \{S_1, \dots, S_n\}$, where each S_i is a set of three vertices. The decision problem is to determine if there exists a subset of \mathcal{S} of size q that covers every vertex $x \in X$ exactly once (an exact cover).

LEMMA 4.2. *X3CC is NP-complete.*

The proof of Lemma 4.2 proceeds via a reduction from X3C, and is presented in the full version. Using this lemma, we can now show that finding an outcome satisfying PPS and PO is \mathcal{NP} -hard through a reduction from X33C.

PROOF OF THEOREM 4.1. Let (Y, \mathcal{T}) be an instance of X33C, with $|Y| = r$. Let $\varepsilon \in (0, 1/(3r))$. We define a public decision making problem as follows. There are r players, one corresponding to each vertex $v \in Y$, and r issues. For each issue, there are $|Y| + |\mathcal{T}|$ alternatives. For each issue t and each player i , there is an alternative $a_{t,i}$ which is valued at $1 - \varepsilon$ by player i , and 0 by all other players. The remaining $|\mathcal{T}|$ alternatives correspond to the 3-sets from the X33C instance. For a set $T_j \in \mathcal{T}$, the corresponding alternative is valued at $\frac{1}{3}$ by players $i \in T_j$, and valued at 0 by all other players. Note that $\text{PPS}_i = 1 - \varepsilon$ for each player i , because there are exactly as many issues as players, and each player values its most preferred alternative for each issue at $1 - \varepsilon$.

We now show that there exists an exact triple-cover by 3-sets if and only if all outcomes to the public decision making problem that satisfy PPS and PO have $u_i(\mathbf{c}) = 1$ for all i . First, suppose that there exists an exact triple-cover by 3-sets. We need to show that all outcomes satisfying PPS and PO have $u_i(\mathbf{c}) = 1$ for all i . So suppose otherwise – that there exists an outcome satisfying PPS and PO with $u_i(\mathbf{c}) \neq 1$ for some player i . In particular, some player must have $u_i(\mathbf{c}) > 1$, otherwise \mathbf{c} is not PO (because it is possible to choose an outcome corresponding exactly to an exact triple-cover, which gives each player utility 1). But players only derive utility in discrete amounts of $1 - \varepsilon$ or $\frac{1}{3}$, which means that any player with $u_i(\mathbf{c}) > 1$ has $u_i(\mathbf{c}) \geq \frac{4}{3} - \varepsilon$. Since \mathbf{c} satisfies PPS, the social welfare must satisfy

$$\sum_{i=1}^r u_i(\mathbf{c}) \geq \frac{4}{3} - \varepsilon + \sum_{i=1}^{r-1} (1 - \varepsilon) = \frac{1}{3} + r - r\varepsilon > r,$$

where the last inequality holds because $\varepsilon < 1/(3r)$. However, this is a contradiction because each alternative in each of the r issues contributes at most 1 to the social welfare. Therefore, every outcome satisfying PPS and PO has $u_i(\mathbf{c}) = 1$ for all i .

Next, suppose that there does not exist an exact triple-cover by 3-sets. So if we choose an alternative corresponding to a 3-set for every issue, it is not possible for every player to derive utility 1. Therefore, some player must derive utility $\frac{2}{3}$ (or lower), which violates PPS. Thus, every outcome that satisfies PPS must include at least one issue where the chosen alternative is one that corresponds to a player, not to a 3-set. Such an alternative only contributes $1 - \varepsilon$ to social welfare. Therefore, the social welfare is strictly less than r , which means that some player gets utility strictly

less than 1. Therefore, there is no outcome satisfying PPS (either with or without PO) such that $u_i(\mathbf{c}) = 1$ for all i . Since the set of outcomes satisfying PPS is always non-empty, it is therefore not the case that all outcomes satisfying PPS and PO have $u_i(\mathbf{c}) = 1$ for all i . ■

Because every outcome satisfying RRS also satisfies PPS, we have the following corollary.

COROLLARY 4.3. *It is \mathcal{NP} -hard to find an outcome satisfying RRS and PO.*

For private goods division, we show, in stark contrast to Theorem 4.1, that we can find an allocation satisfying PPS and PO in polynomial time. This is achieved using Algorithm 1. Interestingly, it produces not an arbitrary allocation satisfying PPS and PO, but an allocation that assigns at least $p = \lfloor m/n \rfloor$ goods to every player – implying PPS, and maximizes weighted (utilitarian) social welfare according to some weight vector – implying PO.

ALGORITHM 1: Polynomial time algorithm to achieve PPS and PO for private goods

Input: The set of players N , the set of private goods M , and players' utility functions $\{u_i\}_{i \in N}$

Output: A deterministic allocation \mathbf{A} satisfying PPS and PO

```

1  $\mathbf{w} \leftarrow (1/n, \dots, 1/n) \in \mathbb{R}^n$ ;
2  $\mathbf{A} \leftarrow \arg \max_{\mathbf{A}'} \sum_{i \in N} w_i \cdot u_i(\mathbf{A}')$ ;
3 while  $\exists i \in N, |A_i| < p$  do           /* Until every player receives at least  $p = \lfloor m/n \rfloor$  goods */
4    $GT \leftarrow \{i \in N : |A_i| > p\}$ ;      /* Partition players by the number of goods they receive */
5    $EQ \leftarrow \{i \in N : |A_i| = p\}$ ;
6    $LS \leftarrow \{i \in N : |A_i| < p\}$ ;
7    $DEC = GT$ ;                               /* Players whose weights we will decrease */
8   while  $DEC \cap LS = \emptyset$  do
9     /* Minimally reduce weights of players in  $DEC$  so a player in  $DEC$  loses a good */
10     $(i^*, j^*, g^*) \leftarrow \arg \min_{i \in DEC, j \in N \setminus DEC, g \in A_i} (w_i \cdot v_{i,g}) / (w_j \cdot v_{j,g})$ ;
11     $r \leftarrow (w_{i^*} \cdot v_{i^*,g^*}) / (w_{j^*} \cdot v_{j^*,g^*})$ ;
12     $\forall i \in DEC, w_i \leftarrow w_i / r$ ;
13     $DEC \leftarrow DEC \cup \{j^*\}$ ;
14     $D(j^*) \leftarrow (i^*, g^*)$ ;           /* Bookkeeping:  $j^*$  can receive  $g^*$  from  $i^*$  */
15  end
16   $j^* \leftarrow DEC \cap LS$ ;             /* Player from  $LS$  who receives a good */
17  while  $j^* \notin GT$  do
18     $(i^*, g^*) \leftarrow D(j^*)$ ;
19     $A_{i^*} \leftarrow A_{i^*} \setminus \{g^*\}$ ;
20     $A_{j^*} \leftarrow A_{j^*} \cup \{g^*\}$ ;
21     $j^* \leftarrow i^*$ ;
22  end
23 return  $\mathbf{A}$ ;

```

At a high level, the algorithm works as follows. It begins with an arbitrary weight vector \mathbf{w} , and an allocation \mathbf{A} maximizing the corresponding weighted (utilitarian) social welfare. Then, it executes a loop (Lines 3-22) while there exists a player receiving less than p goods, and each iteration of the loop alters the allocation in a way that one of the players who received more than p goods loses a good, one of the players who received less than p goods gains a good, and every other player retains the same number of goods as before.

Each iteration of the loop maintains a set DEC of players whose weight it reduces. Initially, DEC consists of players who have more than p goods (Line 7). When the weights are reduced enough so that a player in DEC is about to lose a good to a player, necessarily outside DEC (Lines 9-11), the latter player is added to DEC (Line 12) before proceeding further. When a player who has less than p goods is added to DEC , this process stops and the algorithm leverages the set of ties it created along the way to make the aforementioned alteration to the allocation (Lines 16-21).

We now formally state that this produces an allocation satisfying PPS and PO, and that it runs in polynomial time; the proof appears in the full version.

THEOREM 4.4. *For private goods division, PPS and PO can be satisfied in polynomial time.*

The complexity of finding an allocation (of private goods) satisfying the stronger guarantee RRS along with PO in polynomial time remains open, as does the complexity of finding an allocation satisfying Prop1 and PO.

We note that the convenient approach of weighted welfare maximization we use in Theorem 4.4 cannot be used for finding an outcome satisfying RRS and PO, as the following example shows. This leads us to conjecture that it may be \mathcal{NP} -hard to find such an outcome.

Example 4.5. Consider a private goods division problem with two players and four goods. Player 1 has utilities $(4, 4, 1, 1)$ and player 2 has utilities $(3, 3, 2, 2)$ for goods (g_1, g_2, g_3, g_4) , respectively. Note that $RRS_1 = RRS_2 = 5$. Consider assigning weights w_1 and w_2 to players 1 and 2, respectively. If $4w_1 > 3w_2$, i.e., $w_1 > 3w_2/4$ then player 1 receives both g_1 and g_2 , which means that player 2 receives utility less than her RRS share. On the other hand, if $3w_2 > 4w_1$, i.e., $w_1 < 3w_2/4$ then player 2 receives both g_1 and g_2 , which means that player 1 receives utility less than her RRS share.

The only remaining possibility is that $w_1 = 3w_2/4$, but in that case, player 2 receives both g_3 and g_4 . Regardless of how we divide goods g_1 and g_2 , one of the two players still receives utility less than her RRS share.

In contrast, a simple modification of Algorithm 1 seems to quickly find an allocation satisfying Prop1 and PO in hundreds of thousands of randomized simulations. At each iteration of this version, the set DEC initially consists of players who attain their proportional share (it is easy to show using the Pigeonhole principle that this set is non-empty for any weighted welfare maximizing allocation), and ends when a player is added to DEC that is not currently achieving Prop1. Thus, at every loop, a player that was receiving her proportional share may lose a good (but will still achieve at least Prop1), the player added to DEC that was not achieving Prop1 gains a good, and some players that were achieving Prop1 but not their proportional share may lose a good, but only if they gain one too. These three classes of players are therefore analogous to players with more than p goods, less than p goods, and exactly p goods in Algorithm 1. Unfortunately, we are unable to prove termination of this algorithm because it is possible that a player who achieves Prop1 but not her proportional share loses a high-valued good while gaining a low-valued good, thus potentially sacrificing Prop1. Thus we do not get a property parallel to the key property of Algorithm 1, that no player's utility ever drops below her PPS share, after she attains it. However, our algorithm always seems to terminate quickly and finds an allocation satisfying Prop1 and PO in our randomized simulations, which leads us to conjecture that it may be possible to find an allocation satisfying Prop1 and PO in polynomial time, either from our algorithm directly or via some other utilization of weighted welfare maximization.

5 DISCUSSION

We introduced several novel fairness notions for public decision making and considered their relationships to existing mechanisms and fairness notions. Throughout the paper, we highlighted

various open questions including the existence (and complexity) of a mechanism satisfying RRS, Prop1, and PO, the complexity of finding an outcome satisfying Prop1 and PO (for public decisions and private goods), the complexity of finding an outcome satisfying RRS and PO (for private goods), and whether MNW provides a constant approximation to RRS better than 1/2.

So far we only considered deterministic outcomes. If randomized outcomes are allowed (an alternative interpretation in the private goods case is that the goods are infinitely divisible), then the MNW solution satisfies Prop as a direct consequence of it satisfying Prop1 for deterministic outcomes (Theorem 3.4).⁶ To see this, consider replicating each issue K times and dividing utilities by K . The relative effect of granting a single player control of a single issue becomes negligible. Thus, as K approaches infinity, the utility of each player i in an MNW outcome approaches a value that is at least their proportional share Prop_i . The fraction of copies of issue t in which outcome a_j^t is selected can be interpreted as the weight placed on a_j^t in the randomized outcome. Because RRS, PPS, and Prop1 are relaxations of Prop, the randomized MNW outcome also satisfies all of them.

For private goods division, this can be seen as a corollary of the fact that the randomized MNW outcome satisfies envy-freeness, which is strictly stronger than proportionality. This hints at a very interesting question: Is there a stronger fairness notion than proportionality in the public decision making framework that generalizes envy-freeness in private goods division? Although such a notion would not be satisfiable by deterministic mechanisms, it may be satisfied by randomized mechanisms, or it could have novel relaxations that may be of independent interest.

At a high level, our work provides a framework bringing together two long-studied branches of social choice theory – fair division theory and voting theory. Both have at their heart the aim to aggregate individual preferences into a collective outcome that is fair and agreeable, but approach the problem in different ways. Fair division theory typically deals with multiple private goods, assumes access to cardinal utilities, and focuses on notions of fairness such as envy-freeness and proportionality. Voting theory, in contrast, typically deals with a single public decision (with the exception of combinatorial voting mentioned earlier), assumes access only to less expressive ordinal preferences, has the “one voter, one vote” fairness built inherently into the voting rules, and focuses on different axiomatic desiderata such as Condorcet consistency and monotonicity.

Of course, one can use a voting approach to fair division, since we can have players express preferences over complete outcomes, and this approach has been used successfully to import mechanisms from voting to fair division and vice versa [1, 3]. However, not only does this approach result in an exponential blowup in the number of alternatives, it also does not provide a convenient way to express fair-division-like axioms. Continuing to explore connections between the two fields remains a compelling direction for future work.

ACKNOWLEDGMENTS

Conitzer and Freeman are supported by NSF under awards IIS-1527434 and CCF-1337215, ARO under grants W911NF-12-1-0550 and W911NF-11-1-033, and a Guggenheim Fellowship.

REFERENCES

- [1] H. Aziz, F. Brandt, and P. Stursberg. 2013. On popular random assignments. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*. 183–194.
- [2] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [3] H. Aziz and P. Stursberg. 2014. A Generalization of Probabilistic Serial to Randomized Social Choice.. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 559–565.

⁶Of course, the *realization* may fail to satisfy Prop (and other desiderata), but the *lottery* is fair if we consider expected utilities.

- [4] N. Bansal and M. Sviridenko. 2006. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*. 31–40.
- [5] G. Benade, A. D. Procaccia, S. Nath, and N. Shah. 2017. Preference Elicitation for Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- [6] A. Bogomolnaia, H. Moulin, and R. Stong. 2005. Collective choice under dichotomous preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.
- [7] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. 2015. Optimal Social Choice Functions: A Utilitarian View. *Artificial Intelligence* 227 (2015), 190–213.
- [8] S. Bouveret and J. Lang. 2008. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. *Journal of Artificial Intelligence Research* 32 (2008), 525–564.
- [9] S. J. Brams, M. Kilgour, and C. Klamler. 2012. The undercut procedure: An algorithm for the envy-free division of indivisible items. *Social Choice and Welfare* 39, 2–3 (2012), 615–631.
- [10] F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia (Eds.). 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [11] E. Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [12] Y. Cabannes. 2004. Participatory budgeting: A significant contribution to participatory democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.
- [13] I. Caragiannis, C. Kaklamani, P. Kanellopoulos, and M. Kyropoulou. 2012. The efficiency of fair division. *Theory of Computing Systems* 50, 4 (2012), 589–610.
- [14] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*. 305–322.
- [15] J. R. Chamberlin and P. N. Courant. 1983. Representative Deliberations and Representative Decisions: Proportional Representation and the Borda Rule. *American Political Science Review* 77, 3 (1983), 718–733.
- [16] R. Cole and V. Gkatzelis. 2015. Approximating the Nash Social Welfare with Indivisible Items. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*. 371–380.
- [17] B. Fain, A. Goel, and K. Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*. 384–399.
- [18] D. Foley. 1967. Resource allocation and the public sector. *Yale Economics Essays* 7 (1967), 45–98.
- [19] R. Freeman, S.M. Zahedi, and V. Conitzer. 2017. Fair Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*. Forthcoming.
- [20] A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2016. Knapsack Voting: Voting mechanisms for Participatory Budgeting. Manuscript. (2016).
- [21] C.-J. Haake, M. G. Raith, and F. E. Su. 2002. Bidding for envy-freeness: A procedural approach to n -player fair-division problems. *Social Choice and Welfare* 19 (2002), 723–749.
- [22] F. Klijn. 2000. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare* 17 (2000), 201–215.
- [23] D. Kurokawa, A. D. Procaccia, and N. Shah. 2015. Leximin Allocations in the Real World. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*. 345–362.
- [24] J. Lang and L. Xia. 2016. Voting in combinatorial domains. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 8.
- [25] E. Lee. 2015. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. arXiv:1507.01159. (2015).
- [26] R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*. 125–131.
- [27] B. L. Monroe. 1995. Fully Proportional Representation. *American Political Science Review* 89, 4 (1995), 925–940.
- [28] J. Nash. 1950. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- [29] A. D. Procaccia and J. S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*. 317–331.
- [30] A. D. Procaccia and J. Wang. 2014. Fair Enough: Guaranteeing Approximate Maximin Shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*. 675–692.
- [31] S. Ramezani and U. Endriss. 2010. Nash social welfare in multiagent resource allocation. In *Proceedings of the 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC)*. 117–131.
- [32] J.H. Reijnierse and J. A. M. Potters. 1998. On finding an envy-free Pareto-optimal division. *Mathematical Programming* 83 (1998), 291–311.
- [33] H. Steinhaus. 1948. The problem of fair division. *Econometrica* 16 (1948), 101–104.
- [34] H. Varian. 1974. Equity, envy and efficiency. *Journal of Economic Theory* 9 (1974), 63–91.