

Improving Nash Social Welfare Approximations of Indivisible Goods

Jugal Garg

Peter McGlaughlin

University of Illinois Urbana-Champaign,
Urbana, IL 61801 USA

JUGAL@ILLINOIS.EDU

MCGLGHL2@ILLINOIS.EDU

Abstract

We consider the problem of fairly allocating a set of indivisible goods among n agents. Various fairness notions have been proposed within the rapidly growing field of fair division, but the Nash social welfare (NSW) serves as a focal point. In part, this follows from the ‘unreasonable’ fairness guarantees provided, in the sense that a max NSW allocation meets multiple other fairness metrics simultaneously, all while satisfying a standard economic concept of efficiency, Pareto optimality. However, existing approximation algorithms fail to satisfy all of the remarkable fairness guarantees offered by a max NSW allocation, instead targeting only the specific NSW objective. We address this issue by presenting a 2 max NSW, Prop-1, $1/(2n)$ MMS, and Pareto optimal allocation in strongly polynomial time. Our techniques are based on a market interpretation of a fractional max NSW allocation. We present novel definitions of fairness concepts in terms of market prices, and design a new scheme to round a market equilibrium into an integral allocation in a way that provides most of the fairness properties of an integral max NSW allocation.

1. Introduction

We study the age old problem of fair allocation of goods among agents with additive valuations. The formal study of fair division of goods dates back to the cake cutting problem introduced by Steinhaus (1948). In the simplest version, two agents must agree on a way to split a heterogeneous cake. The well known cut and choose protocol achieves two standard notions of fairness: envy-freeness where every agent weakly prefers their allocation over any other agents’ allocation, and proportionality where every agent receives at least a $1/n$ share of the goods.

Another fairness criterion, max Nash social welfare (NSW), i.e., geometric mean of the agents’ valuations, also yields envy-free and proportional allocations. NSW strikes a balance between the perfectly efficient utilitarian objective, maximizing the sum of agents’ valuations, and the undoubtedly fair egalitarian objective, maximizing the minimum valued bundle. Further, an allocation which maximizes NSW (max NSW) corresponds to a certain equilibrium of a linear Fisher market which defines a *fractional* allocation and a set of *prices*. This equilibrium, known as a competitive equilibrium of equal incomes (CEEI), yields an allocation that is not only envy-free and proportional (Varian, 1974), but also Pareto optimal (PO), a standard notion of economic efficiency. Then, *fractional* max NSW allocations are both *fair* and *efficient*.

When goods are indivisible, no algorithm can ensure either envy-freeness or proportionality, e.g., consider allocating a single good between two agents. This motivates the need

for alternate concepts of fairness. One choice is to provide additive approximations. In an envy-free up to one good (EF-1) allocation each agent weakly prefers her bundle after *removing* some good from another agent’s bundle. Similarly, an allocation is proportional up to one good (Prop-1), if each agent values her bundle at least a $1/n$ share of the goods after *adding* some good not allocated to her. A maximin share (MMS) offers an intriguing alternative relaxation to proportionality. The idea is a natural generalization of the cut and choose protocol from the cake cutting problem. Suppose we allow one agent to choose a partition of the goods into n bundles (one for each agent), with the caveat that the other agents get to choose a bundle before her. In the worst case, she receives her least preferred bundle. Clearly, the agent must select a partition that maximizes the value of her least preferred bundle. We call the value of this bundle the agent’s maximin share (MMS). In an α MMS allocation, $\alpha \in [0, 1]$, each agent receives a bundle she values at least α MMS.

NSW also serves as a major focal point in fair division of indivisible goods. Caragiannis et al. (2016) present a compelling argument in favor of the ‘unreasonable’ fairness of max NSW allocations. Analogous to the case of divisible goods, an *integral* max NSW allocation is EF-1 (and therefore Prop-1, Conitzer, Freeman, and Shah, 2017), $\Omega(n^{-1/2})$ MMS, and PO.

1.1 Our Contribution

As a testament to its inherent fairness, NSW has been rediscovered in multiple contexts and employed in wide variety of applications to ensure fair outcomes. Much of its ‘unreasonable’ fairness stems from the fact a max NSW allocation also meets multiple, sometimes conflicting, fairness notions simultaneously, all while still satisfying the standard economic notion of efficiency, Pareto optimality.

However, finding an integral max NSW allocation is APX-hard (Lee, 2017). Moreover, existing approximations only target the NSW objective, i.e., α max NSW, for $\alpha > 1$, falling short of the array of fairness properties found in a true max NSW allocation. Thus, current literature leaves a clear gap between *all* of the strong fairness guarantees offered by a max NSW allocation and its approximations. This begs the question: Does an approximation exist that more closely resembles all the remarkable properties of a true max NSW allocation? In this paper, we definitively answer yes. Specifically, we provide an algorithm which yields a 2 max NSW, Prop-1, $1/(2n)$ MMS, and PO allocation. We also note that none of the previous algorithms achieves these bounds, and we provide counterexamples in Section 5.

Our techniques rely on a market interpretation of *fractional* max NSW allocations. First, we present novel bounds on any agent’s Prop and MMS guarantees in terms of market prices. Then, we design a new scheme to round a market equilibrium into an integral allocation in a way that provides most fairness properties of an integral max NSW allocation. We note that a preliminary version of this paper appeared in the proceedings of IJCAI 2019 (Garg & McGlaughlin, 2019).

1.2 Related Work

Within the rapidly growing literature on fair division of indivisible goods, various authors proposed a myriad of fairness metrics. We focus on some of the most common notions

of a fair allocation: EF-1, Prop-1, NSW, and MMS. We refer the reader to Amanatidis, Birmpas, and Markakis (2018) for a more thorough survey of the current landscape of fair division and the relationship between various fairness concepts.

EF-1 and Prop-1. Budish (2011) first introduced EF-1, although the concept was implied by the earlier work of Lipton, Markakis, Mossel, and Saberi (2004) which offered a local search algorithm to find such an allocation. More recently, Conitzer et al. (2017) introduced Prop-1 in the context of fair division of public goods, i.e., where *all* agents receive the same allocation. We note that this setting generalizes the private goods case we consider. The authors also demonstrate that a round robin procedure achieves a Prop-1 allocation in polynomial time. However, they leave the problem of finding a Prop-1 and PO allocation in polynomial time as an open question.

The recent work of Barman and Krishnamurthy (2019) addresses another issue concerning the relationship between integral max NSW allocations and Fisher market equilibria. It is well known that a CEEI, i.e., a Fisher market equilibrium where all agents have the same budget, results in a max NSW, envy-free, and proportional allocation. However, in general, the CEEI allocation is fractional not integral. Barman and Krishnamurthy (2019) show that in any problem instance, there exists a ‘nearby’ approximate CEEI with integral allocation where the agents’ budgets are ‘perturbed’ by no more than the maximum market price. Further, they use this result to design a strongly polynomial time algorithm which yields a Prop-1 and PO allocation.

NSW. Cole and Gkatzelis (2015) give the first polynomial time constant factor (2.89) approximation to the NSW objective by rounding a fractional allocation into an integral one. Cole et al. (2017) refine the analysis of the algorithm, showing it yields a 2 max NSW allocation. More recently, Barman, Krishnamurthy, and Vaish (2018) design a purely combinatorial algorithm that yields a 1.45 max NSW, PO, and a $(1 + \epsilon)$ EF-1 allocation for any $\epsilon > 0$. We note that both works heavily inspired notable subsequent papers (Anari, Mai, Gharan, and Vazirani, 2018; Garg, Hoefer, and Mehlhorn, 2018; Chaudhury et al., 2018) which extend these techniques beyond the case of additive valuations.

Caragiannis, Gravin, and Huang (2019) propose a new interesting approach to find a constant factor approximation to max NSW. The authors primarily consider EF-X allocations, a stronger form of EF-1 wherein each agent prefers her bundle after removing *any* good from another agent’s bundle. Starting with an α max NSW allocation, they show how to obtain an EF-X allocation after *removing* a ‘small’ number of goods from the problem instance. Here, a ‘small’ number of goods means that α max NSW is reduced by no more than half the original instance, i.e., in the reduced instance the allocation is 2α max NSW. Further, the local search algorithm of Lipton et al. (2004) provides a means to ‘add’ the removed goods back to the instance to generate a 2α max NSW and EF-1 allocation in polynomial time, the first result of it’s kind.

We note that max NSW has also been studied from the perspective of social choice; see Darmann and Schauer (2015), or Baumeister et al. (2017).

MMS. Budish (2011) defined maximin share which gives an intuitive local measure of fairness of an allocation, i.e., a specific objective for each agent. Bouveret and Lemaître (2014) showed that an MMS allocation exists in certain special settings. Procaccia and Wang (2014) and Kurokawa, Procaccia, and Wang (2016) obtained the surprising result that MMS al-

Algorithm	max NSW	Prop-1	MMS	$(1 + \epsilon)$ EF-1	PO
Our work	2	✓	$O(1/n)$	X	✓
(Cole et al., 2017)	2	X	$O(1/n)$	X	✓
(Barman, Krishnamurthy et al., 2018)	1.45	X	$O(1/n)$	✓	✓
(Barman & Krishnamurthy, 2019)	X	✓	X	X	✓
(Caragiannis et al., 2019)	2.9	✓	$O(1/n)$ ¹	✓	X

Table 1: A comparison of fairness guarantees provided.

locations might not exist, and showed the existence of a $2/3$ MMS allocation. Amanatidis, Markakis, Nikzad, and Saberi (2017) obtained a PTAS which finds a $(2/3 - \epsilon)$ MMS allocation. Later, Barman and Krishnamurthy (2017) and Garg, McGlaughlin, and Taki (2019) obtained simple algorithms to find a $2/3$ MMS allocation. More recently, Ghodsi, Hajjaghayi, Seddighin, Seddighin, and Yami (2018) gave a PTAS to find a $(3/4 - \epsilon)$ MMS allocation, and Garg and Taki (2019) obtained an algorithm to find a state of the art $3/4$ MMS allocation. We note that a few works consider generalizations of the MMS concept, e.g., (Barman, Biswas, Krishnamurthy, and Narahari, 2018; and Farhadi et al., 2019).

We note that none of the above algorithms yields a constant factor approximation to max NSW, Prop-1, $1/(2n)$ MMS, and PO allocation simultaneously in polynomial time. Table 1 summarizes the capabilities of existing algorithms, and compares them to our result.

2. Preliminaries

We consider the fair allocation of a set $M = \{1, \dots, m\}$ of m indivisible goods among a set $N = \{1, \dots, n\}$ of n agents. We assume that $m > n$, otherwise we can find an exact max NSW allocation in polynomial time using a maximum weight matching (Nguyen & Rothe, 2014). A fractional allocation $x = (x_i)_{i \in N}$ is an assignment of goods to agents. Agent i 's bundle is $x_i = (x_{ij})_{j \in M}$ where $x_{ij} \in [0, 1]$ is the fraction of good j allocated to i . An allocation is integral, corresponding to the case of indivisible goods, if $x_{ij} \in \{0, 1\}, \forall i, j$. We use the notation $A = (A_i)_{i \in N}$ to denote integral allocations to avoid confusion between fractional and integral allocations. That is, $A_i = \{j \in M : x_{ij} = 1\}$. Let $\mathcal{X} = \{x : \sum_{i \in N} x_{ij} \leq 1, \forall j \in M\}$ be the set of feasible fractional allocations, and \mathcal{A} be the set of feasible integral allocations.

Let v_{ij} denote agent i 's valuation of good j , and for a bundle x_i , i 's valuation is $v_i(x_i) = \sum_{j \in M} v_{ij} x_{ij}$. When x is an integral allocation, we overload the definition of v_i as follows. For any subset of goods $S \subseteq M$, define $v_i(S) = \sum_{j \in S} v_{ij}$.

NSW of an allocation is defined as

$$NSW(x) = \left(\prod_{i=1}^n v_i(x_i) \right)^{1/n}.$$

A max NSW allocation x^* maximizes NSW over all feasible allocations, i.e., $x^* \in \arg \max_{x \in \mathcal{X}} NSW(x)$. An allocation x is an α max NSW allocation if $\alpha NSW(x) \geq NSW(x^*)$ for some $\alpha \geq 1$.

1. Segal-Halevi and Suksompong (2018) show that any EF-1 allocation is also $1/n$ MMS.

MMS of agent $i \in N$ is defined as

$$MMS_i = \max_{A \in \mathcal{A}} \min_{A_k \in A} v_i(A_k).$$

A maximin share allocation A satisfies $v_i(A_i) \geq MMS_i, \forall i \in N$. For any $\alpha \in (0, 1)$, an α MMS allocation gives each agent a bundle with value $v_i(A_i) \geq \alpha MMS_i$.

EF-1 Envy-free (EF) requires that each agent weakly prefers her bundle over any other, i.e., $v_i(A_i) \geq v_i(A_k), \forall k \in N$. A EF-1 allocation is envy-free after removing some item from another agent's bundle. That is,

$$\forall i, k \in N, \exists j \in A_k, v_i(A_i) \geq v_i(A_k \setminus j).$$

An α -EF1 allocation, a relaxation of EF-1 for $\alpha > 1$, is defined as

$$\forall i, k \in N, \exists j \in A_k, \alpha v_i(A_i) \geq v_i(A_k \setminus j).$$

Prop-1 Proportionality (Prop) requires that each agent receives a bundle they value at least as much as an equal share of the goods, i.e., $v_i(A_i) \geq \frac{1}{n} v_i(M), \forall i$. A Prop-1 allocation is proportional after adding some good not allocated to the agent. That is,

$$\forall i \in N, \exists j \in M \setminus A_i, v_i(A_i) + v_{ij} \geq v_i(M)/n.$$

PO An allocation x' Pareto dominates another allocation x if $v_i(x'_i) \geq v_i(x_i), \forall i \in N$, and $v_k(x'_k) > v_k(x_k)$ for some agent k , i.e., if at least one agent's utility improves without sacrificing the utility of any other agent. An allocation x is PO if no allocation x' dominates x .

3. Linear Fisher Markets

Our algorithm relies on rounding a fractional allocation x obtained from a market equilibrium. We note that this approach was used by Cole and Gkatzelis (2015) to obtain the first constant factor max NSW approximation. In this section, we summarize the relationship between the fractional NSW, the Eisenberg-Gale (EG) program, and linear Fisher markets. Although the agents in our problem do not spend any money, the fictitious market interpretation provides useful information to aid in rounding the fractional allocation.

Consider maximizing the NSW over fractional allocations, which is equivalent to maximizing the sum of logarithms of the agent valuations. This turns out to be a special case of well known Eisenberg-Gale program (Eisenberg & Gale, 1959) where $e_i = 1, \forall i$.

$$\max_{x \in \mathcal{X}} \sum_{i=1}^n e_i \log v_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^n x_{ij} \leq 1, \text{ and } x_{ij} \geq 0, \forall i, j. \quad (\text{EG})$$

3.1 A Market Interpretation

The central concept in the design of our algorithm plays off an interpretation of solutions to (EG) as the equilibrium of a linear Fisher market. A Fisher market is the tuple $\langle N, M, V, e \rangle$, where N is a set of n agents, M is a set of m goods, $V = (v_i)_{i \in N}$ defines the agents valuations,

and $e = (e_i)_{i \in N}$ gives the agents' budgets. In this setting, each agent seeks to spend their budget on fractions of goods. The utility of agent i for receiving x_{ij} amount of each good j is $\sum_{j \in M} v_{ij} x_{ij}$. A fractional allocation x and a price vector $p = (p_j)_{j \in M}$ correspond to primal and dual variables of (EG) respectively.

Definition 1. The pair (\bar{x}, \bar{p}) is an equilibrium if: (i) all goods are fully allocated $\sum_{i=1}^n \bar{x}_{ij} = 1, \forall j \in M$, (ii) all agents spend their budget $\sum_{j \in M} \bar{x}_{ij} \bar{p}_j = e_i, \forall i \in N$, and (iii) all agents only purchase maximum bang per buck (MBB) goods, i.e., if $\bar{x}_{ij} > 0$, then $j \in \arg \max_{k \in M} v_{ik} / \bar{p}_k$.

For any vector of prices \bar{p} , we define the set of MBB goods for agent i as

$$MBB_i = \arg \max_{k \in M} v_{ik} / \bar{p}_k. \tag{1}$$

Spending Graph. Our algorithm exploits some structural properties of equilibria. Namely, an equilibrium (\bar{x}, \bar{p}) admits a representation as a forest. First, we define a bipartite graph using the vertices $N \cup M$, with the set of agents N on one side and the set of goods M on the other. Using the fractional allocation \bar{x} , we create the spending graph $Q(\bar{x}, \bar{p})$ by adding a weighted edge (i, j) between agent i and good j with weight $\bar{x}_{ij} \bar{p}_j$, if $\bar{x}_{ij} > 0$. Due to equilibrium condition (iii), $j \in MBB_i$ for each edge (i, j) of the spending graph.

Lemma 1. (*Duan, Garg, & Mehlhorn, 2016*) *It is always possible to rearrange the agents' spending to ensure the spending graph is a forest.*

Theorem 2. (*Mas-Colell, Whinston, & Green, 1995*) *Every Fisher market equilibrium is Pareto optimal.*

3.2 Spending Restricted Equilibrium

Equilibria of linear Fisher markets define fractional allocations \bar{x} . Naturally, one hopes to devise a way to round \bar{x} into an integral allocation A that, at least approximately, preserves its strong fairness guarantees. However, Cole and Gkatzelis (2015) demonstrated that no rounding procedure yields a meaningful approximation to the fractional max NSW, and they gave a clever modification. Specifically, they relax the constraint that all goods need to be fully allocated, and instead require that the total spending on any good does not exceed 1, the budget of an agent. That is, for all goods $j \in M$, the total spending on good j is $\sum_{i=1}^n \bar{x}_{ij} \bar{p}_j = \min(1, \bar{p}_j)$. The modified constraint slightly changes the equilibrium conditions.

Definition 2. The pair (\bar{x}, \bar{p}) is a spending restricted (SR) equilibrium if: (i') the total spending on each good j is $\sum_{i=1}^n \bar{x}_{ij} \bar{p}_j = \min(1, \bar{p}_j)$, and conditions (ii) and (iii) of Definition 1 hold.

The condition (i') of Definition 2 is the defining characteristic of an SR equilibrium, and the only difference between SR and Fisher equilibria. Notably, in an SR equilibrium all agents still purchase only MBB goods, and we can rearrange agents' spending so that the resulting spending graph is a forest. Further, Cole and Gkatzelis (2015) provide a strongly polynomial time algorithm to compute an SR equilibrium. Note that we always compute SR equilibria using agents' budgets $e_i = 1, \forall i \in N$.

4. Approximation Algorithm

We present an algorithm that takes as input an SR equilibrium (\bar{x}, \bar{p}) , and outputs an integral allocation $A = (A_i)_{i \in N}$ that is 2 max NSW, $1/(2n)$ MMS, Prop-1, and PO.

We start with the important observation that the NSW, MMS, and Prop-1 of an agent are scale invariant, i.e., for any agent i and any $c \in \mathbb{R}^+$, scaling i 's valuations as $v_{ij} \leftarrow cv_{ij}, \forall j \in M$, doesn't change the problem. Define i 's MBB ratio $\alpha_i := \max_{j \in M} v_{ij}/\bar{p}_j$. It is useful to rescale agents' valuations, $v_{ij} \leftarrow v_{ij}/\alpha_i$. Notice that this implies $v_{ij} = \bar{p}_j, \forall j \in MBB_i$, and $v_{ij} < \bar{p}_j$ otherwise.

Definition 3. For any price vector \bar{p} , we say that an agent's valuations are scaled to prices if: $v_{ij} = \bar{p}_j$ for all $j \in MBB_i$, and $v_{ij} < \bar{p}_j$ otherwise.

4.1 Basic Bounds

Before discussing specifics of the rounding scheme, we establish upper bounds on the various fairness metrics. These results rely on a few properties of SR equilibria which provide some insight on how to approach rounding the fractional allocation \bar{x} . To avoid repetition, let (\bar{x}, \bar{p}) be an SR equilibrium and assume agents' valuations are scaled to prices.

Lemma 3. (Cole & Gkatzelis, 2015) *Let x^* be an integral max NSW allocation. Then, $\prod_{i=1}^n v_i(x_i^*) \leq \prod_{\bar{p}_j > 1} \bar{p}_j$.*

Next, we provide upper bounds on any agents' proportionality, $Prop_i = v_i(M)/n$, and MMS_i guarantees in an SR equilibrium. Let $H = \{j \in M : \bar{p}_j > 1\}$ be the set of high priced goods, and $L = M \setminus H$ be the set of low priced goods. Since the budget of each agent is 1, the total spending on all goods in an equilibrium is n . Also, the spending on any good is capped at 1, so that the total spending on high valued goods H is $|H|$. Then, clearly $|H| \leq n$, due to SR equilibrium conditions, see Definition 2. Moreover, we can assume that $|H| < n$, otherwise that would imply $m = n$, meaning we can compute an exact max NSW allocation in polynomial time, see Section 2.

Lemma 4. *For agent $i \in N$, $Prop_i(M) = v_i(M)/n \leq \max(1, \max_{j \in M} v_{ij})$.*

Proof. Following the above discussion and the fact that agents' valuations are scaled to prices

$$\sum_{j \in L} v_{ij} \leq \sum_{j \in L} \bar{p}_j = n - |H|, \forall i \in N. \quad (2)$$

Fix an agent i and let $v_i^* = \max_{j \in M} v_{ij}$, then since agents' valuations are scaled to prices $v_i(M) = \sum_{j \in L} v_{ij} + \sum_{j \in H} v_{ij} \leq n - |H| + v_i^*|H|$. By considering the two cases: $v_i^* > 1$ or $v_i^* \leq 1$, the result follows easily. \square

Lemma 5. *For any agent $i \in N$, $MMS_i \leq 1$.*

Proof. Let $MMS_i^n(M)$ be agent i 's MMS given the goods M and n agents, i.e., i must partition the goods of M into n bundles in a way that maximizes the minimum valued bundle. We first show that for any agent i , if we remove any other agent l and any good j , then $MMS_i^{n-1}(M \setminus j) \geq MMS_i^n(M)$.

Let $B = (B_i)_{i \in N}$ be the bundles i selects when calculating her $MMS_i^n(M)$, i.e., $v_i(B_k) \geq MMS_i^n(M), \forall B_k \in B$. Wlog, assume B_n contains j . Suppose we redistribute the goods of $B_n \setminus j$ arbitrarily to the other bundles of B to create a partition of $M \setminus j$ into $n - 1$ bundles $B' = (B'_1, \dots, B'_{n-1})$. Then, $v_i(B'_k) \geq v_i(B_k) \geq MMS_i^n(M), \forall k \leq n - 1$.

Clearly, the above argument generalizes to the case of removing k agents and k goods, i.e., $MMS_i^{n-k}(M \setminus \{j_1, \dots, j_k\}) \geq MMS_i^n(M)$, for any $k < n$ and $\{j_l\}_{l=1}^k \subseteq M$. Applying this to the set of high valued goods H yields $MMS_i^n(M) \leq MMS_i^{n-|H|}(M \setminus H) = MMS_i^{n-|H|}(L)$. Further, $MMS_i^{n-|H|}(L) \leq v_i(L)/(n-|H|)$. Then, using (2), $MMS_i^n(M) \leq \frac{v_i(L)}{n-|H|} \leq \frac{1}{n-|H|} \sum_{j \in L} \bar{p}_j \leq 1$. \square

4.2 Rounding Algorithm

We explain the steps of our algorithm below and provide a formal description in Algorithm 1, which we refer to as ALG. Let (\bar{x}, \bar{p}) be an SR equilibrium. We assume agents' valuations are scaled to prices \bar{p} , and that the spending graph is a forest T , by Lemma 1. In a tree $t \in T$, let M_t and k_t be the set of goods and the number of agents respectively. We note that the initial steps are similar to the ones in Cole and Gkatzelis (2015).

In Step 0, for each tree $t \in T$, we pick an arbitrary agent to serve as the root, which we refer to as the *root agent*. Notice that this gives the levels of each tree a particular structure. All vertices at depth 1 are goods that the root agent spends on in the fractional allocation \bar{x} , vertices at depth 2 correspond to agents purchasing some fraction of the goods at depth 1, and so on. Since the tree t is part of the spending graph, all goods $j \in M_t$ connected to agent i are in fact MBB_i .

Observe that, if good $j \in M_t$ is a leaf node with parent agent i , then i is the only agent buying j in \bar{x} . Hence, we assign every leaf good to its parent agent in Step 1.

Step 2 assigns all goods $j \in M_t$ such that $\bar{p}_j < 1/2$ to their parent agent i . For any child agent k of good j , we make k the root agent of the new tree formed by cutting the edge between k and j in the spending graph, i.e., k becomes the root of the subtree containing all goods and agents below her. Notice that, since agents spend their entire budget of 1 and $\bar{p}_j < 1/2$, the newly created root agent k loses no more than half her budget in this step.

In Step 3, for any $j \in M_t$ with two or more child agents, we cut the edges of all child agents *except* the one spending the most on j , i.e., the agent i with the largest \bar{x}_{ij} . Similar to Step 2, any other child agent k of j becomes the root agent of a new tree. Also, notice that the root agent k of the newly formed tree loses no more than half of her budget in this step since the total spending on any good is capped at 1. Steps 2 and 3 endow the remaining trees with a specific structure. Each tree $t \in T$ contains k_t agents and exactly $k_t - 1$ unassigned goods from M_t . Further, each unassigned good has exactly one parent agent and one child agent. In addition, all unassigned goods are worth a least $\bar{p}_j \geq 1/2$ to *both* the parent and child agents. We offer a recursive rounding scheme to meet all stated fairness notions in Step 4.

In the recursion, we select the agent receiving the highest value allocation before the assignment of any remaining goods in the tree, i.e., $i \in \arg \max_{k \in t} v_k(A_k)$ (ties are broken arbitrarily). We test if adding either the agent's parent good or any of its child goods satisfies her Prop-1 guarantee. Formally, let $N_i = \{j \in M_t : j \text{ is parent or child of } i\}$ be the neighbors of i in the tree. Then, we check if $v_i(A_i) + \max_{j \in N_i} v_{ij} \geq 1$. If this condition

Algorithm 1: Rounding Algorithm

Input : SR equilibrium (\bar{x}, \bar{p})

Output: A 2 max NSW, Prop-1, $\frac{1}{2n}$ MMS, and PO allocation

- 1 **Initialize:** $A_i = \emptyset$ for all $i \in N$.
 - 2 **Step 0:** Select a root agent in each tree.
 - 3 **Step 1:** Assign all leaf goods to their parent agents.
 - 4 **Step 2:** Assign agents all child goods j s.t. $\bar{p}_j < 1/2$.
 - 5 **Step 3:** For all goods j s.t. $\bar{p}_j \geq 1/2$, cut edges between j and all child agents except the one spending the most on j , i.e., the largest \bar{x}_{ij} .
 - 6 **Step 4:** For tree t with $k_t \geq 2$ agents, let N_i be the neighbors of agent i in the tree t (parent and child goods). Sort and relabel agents in decreasing order of current total value, i.e., $v_i(A_i)$. Let $i^* = \min\{i \in t : v_i(A_i) + \max_{j \in N_i} v_{ij} \geq 1\}$, and remove i^* from the tree.
 - 7 **Step 5:** Assign all remaining goods in the tree.
-

is satisfied, then we remove i from the tree. Otherwise, we test the same condition for the second highest spending agent in the tree, and so on until we find an agent that values her current allocation enough so that she can meet the Prop-1 guarantee without receiving any item from the matching tree (refer to Lemma 8).

At the start of Step 4, there are matchings that assign $k_t - 1$ agents a good in the matching tree. Removing one agent ensures that a perfect matching between remaining agents and goods exists. In Step 5, we allocate all unassigned goods.

Theorem 6. *ALG returns a 2 max NSW, Prop-1, $1/(2n)$ MMS, PO allocation in strongly polynomial time.*

4.3 Proof of Theorem 6

Observe that each step of ALG runs in strongly polynomial time, and therefore so does ALG. In the following lemmas we verify that all fairness concepts are satisfied. Let $A = (A_1, \dots, A_n)$ be the final integral allocation output by ALG.

Lemma 7. *A is Pareto optimal.*

Proof. We define a modified Fisher market instance (N, M, V, e') using the same set of agents, goods and agent valuations, but with different budgets. For each agent i , we set their new budget as $e'_i = \sum_{j \in A_i} \bar{p}_j$. By Theorem 2, it is enough to check the allocation A and prices \bar{p} satisfy the equilibrium conditions (Definition 1). Obviously, all goods are allocated, and each agent spends their entire budget by construction. Observe that agent i 's allocation is a subset of goods she spent on in the SR equilibrium, i.e., i purchases only MBB goods. It follows that A and prices \bar{p} are market equilibrium for the Fisher market instance (N, M, V, e') . \square

The following observations play a key role in the remaining analysis. First, in the fractional allocation each agent receives a bundle they value at 1. This follows from equilibrium

condition (ii) and the fact that valuations are scaled to prices, i.e., $v_i(\bar{x}_i) = \sum_{j \in M} \bar{p}_j \bar{x}_{ij} = e_i = 1$. Second, after completing Steps 0-3, all remaining trees t have exactly k_t agents and $k_t - 1$ unassigned goods. We call these matching trees. Further, every unassigned good in a matching tree has one parent agent and one child agent. Finally, in each matching tree $t \in T$, there exists *at most* one good $j \notin M_t$ such that some agent i was spending on j in \bar{x} , i.e., $\bar{x}_{ij} > 0$. To see this, notice that we only cut edges of the spending graph in Steps 2 and 3. If the edge between i and j is cut in either step, then i was a child of j . After cutting the edge, i becomes the root of a new tree. Since edges between root agents and their child goods are never cut, then only the root agent can lose one edge from the original spending graph.

The following lemma lower bounds the value of the bundle any agent receives in a matching tree. This lemma is crucial to proving the various fairness properties of ALG.

Lemma 8. *In each matching tree t , the agent i removed in Step 4 of ALG receives a bundle such that $v_i(A_i) \geq 1/(2k_t)$.*

Proof. Let $A' = (A'_1, \dots, A'_n)$ be the agents' allocations at the beginning of Step 4. Also, let N_i be the neighbors of agent i in the matching tree (parent and child goods). Recall agent i is removed from the tree if $v_i(A'_i) + \max_{j \in N_i} v_{ij} \geq 1$.

First, we show that we can always find an agent satisfying this condition. Observe that, at the beginning of Step 4, any leaf agent i retains all edges in the original spending graph. Also, she values her fractional allocation $v_i(\bar{x}_i) = 1$, and she has exactly one parent good, say j , in the matching tree. Then, $v_i(A'_i) = v_i(\bar{x}_i) - \bar{x}_{ij} v_{ij}$. Since good j is not integrally assigned to i , $\bar{x}_{ij} < 1$. It follows that $v_i(A'_i) + v_{ij} = v_i(\bar{x}_i) + (1 - \bar{x}_{ij})v_{ij} > v_i(\bar{x}_i) = 1$.

We show that i^* removed from the tree in Step 4 receives a bundle such that $v_{i^*}(A_{i^*}) \geq 1/(2k_t)$. Observe that there is at most one good $j \notin M_t$ that the root agent was spending on in \bar{x} , and she spent no more than $\bar{x}_{ij} p_j \leq 1/2$. Therefore, the total spending of agents in the tree on the goods of M_t , together with the goods already assigned is at least $k_t - 1/2$, since all agents spend their entire budget in the equilibrium. Also, the agents spend no more than $k_t - 1$ on the unassigned goods in the matching tree, since (\bar{x}, \bar{p}) is an SR equilibrium. This implies that some agent must receive at least $1/(2k_t)$ before Step 4. Recall that we check the agents in decreasing order of the value of their bundle at the beginning of Step 4. If the first agent is removed, then clearly she receives a bundle worth at least $1/(2k_t)$.

Suppose that the agent receiving the most at the beginning of Step 4, say i , does not satisfy the condition. Then, $v_i(A'_i) + v_{ij} = \sum_{k \in A'_i} \bar{p}_k + \bar{p}_j < 1, \forall j \in N_i$, since valuations are scaled to prices. Therefore, if i and *any* good $j \in N_i$ are removed from the tree, then the remaining $k_t - 1$ agents spend at least $k_t - 1/2 - (\sum_{k \in A'_i} \bar{p}_k + \bar{p}_j) \geq k_t - 3/2$ on the goods of $M_t \setminus j$, and no more than $k_t - 2$ is spent on the other $k_t - 2$ unassigned goods of $M_t \setminus j$. Therefore, some agent must receive a bundle worth at least $1/[2(k_t - 1)]$. Noting that all leaf agents satisfy $v_i(A'_i) + \max_{j \in N_i} v_{ij} \geq 1$, we can iterate this argument to ensure the agent removed in Step 4 receives a bundle worth at least $1/(2k_t)$. \square

Lemma 9. *ALG returns a $1/(2n)$ MMS allocation.*

Proof. By Lemma 5, $MMS_i \leq 1$ for all agents. Observe that if agent i was never part of a matching tree, then she receives a bundle worth at least $1/2$. In a matching tree, every agent, except the one removed in Step 4, receives one of the unassigned goods they value

$v_{ij} = \bar{p}_j \geq 1/2$. Lemma 8 shows that i^* removed from the tree receives bundle worth at least $1/(2k_t)$. Obviously, the size of any tree $k_t \leq n$. It follows that all agents receive at least $1/(2n)MMS_i$. \square

Lemma 10. *ALG returns a Prop-1 allocation.*

Proof. Recall that $Prop_i \leq \max(1, \max_{j \in M} v_{ij})$, by Lemma 4. Therefore, for any i , if $\max_{j \in M} v_{ij} > 1$, then the Prop-1 condition is trivially satisfied. For this reason, assume that $\max_{j \in M} v_{ij} \leq 1, \forall i \in N$.

The agent i^* removed from a matching tree t during Step 4 clearly satisfies Prop-1. Suppose agent i receives good j in a matching tree, i.e., $v_{ij} = \bar{p}_j \geq 1/2$. We consider two cases. If i has two or more child goods, then at least one child good, say k , is not in her integral allocation A_i , and $v_{ik} = \bar{p}_k \geq 1/2$. Therefore, $v_i(A_i) + v_{ik} \geq v_{ij} + v_{ik} \geq 1$. If i has at most one child good (or is in a tree of size $k_t = 1$), then there is at most one good k in her fractional allocation that is not in her integral allocation, i.e., $v_i(A_i) + v_{ik} \geq v_i(\bar{x}_i) = 1$. \square

The next part is an adaptation of the analysis given in Cole et al. (2017). Notice that all but one agent in each matching tree receives a good with value $\bar{p}_j \geq 1/2$. The remaining agent, say r , was removed from the tree in Step 4, and she receives at least $v_r = v_r(A_r) \geq 1/(2k_t)$, by Lemma 8.

Lemma 11. *In any matching tree t with $k_t \geq 2$ agents, let $r \in N$ be the agent removed in Step 4 of ALG and let $v_r = v_r(A_r)$ be the value of the bundle she receives. The minimum product of agents' valuations is achieved when at least $k_1 = \lfloor (k_t - 2v_r)/(1 + 2v_r) \rfloor$ agents receive a bundle worth at least 1.*

Proof. Let A be an allocation with minimum product of the agents' valuations, and let $A' = (A'_i)_{i \in N}$ be the agents' bundles (partial allocations) before Step 4 which eventually lead to A . Recall that the agent r removed from the tree satisfies $v_r(A'_r) + v_{rj} \geq 1$ for some $j \in N_r$, and we check this condition in decreasing order of $v_i(A'_i)$. Therefore, for any i such that $v_i(A'_i) > v_r(A'_r)$, we have that her final allocation A_i satisfies $v_i(A_i) < 1$, since her final allocation $A_i = A'_i + v_{ij}$ for some good $j \in N_i$. Also, for any agent i receiving a final bundle worth $v_i(A_i) \geq 1$, the total spending on the goods of A_i is no more than $1 + v_r$, i.e., $\sum_{j \in A_i} \bar{p}_j \leq 1 + v_r$, since i receives only one additional good.

The proof follows by contradiction. Let $q_j \leq 1$ denote the total spending on good $j \in M_t$, and let M'_t be goods allocated to the k_1 agents receiving at least 1 and agent r . Then

$$\begin{aligned} \sum_{j \in M_t \setminus M'_t} q_j &= k_t - 1/2 - [k_1(1 + v_r) + v_r] \\ &= \frac{1}{2}(k_t - k_1) + \frac{1}{2}k_t - \left(\frac{1}{2} + v_r\right)(k_1 + 1). \end{aligned}$$

Suppose, for contradiction, $k_1 + 1 \leq \lfloor (k_t - 2v_r)/(1 + 2v_r) \rfloor$. Then the remaining $k_t - k_1 - 1$ agents have value at least $(k_t - k_1)/2$. Further, at least two of these agents receive a bundle worth $v_i(A_i) > 1/2$, otherwise the remaining agent receives a bundle worth at least $(k_t - k_1)/2 - (k_t - k_1 - 2)/2 = 1$, contradicting the definition of k_1 . However, if two agents, say 1 and 2, receive bundles they value $v_1, v_2 \in (1/2, 1)$, then clearly giving agent 1 a bundle

worth $1/2$ and agent 2 a bundle worth $v_1 + v_2 - 1/2$ produces a lower product of the agents' valuations, a contradiction. \square

Lemma 12. (Cole et al., 2017) For any tree t with k_t agents

$$\prod_{i \in t} v_i(A_i) \geq \frac{1}{2^{k_t}} \prod_{j \in M_t: \bar{p}_j > 1} \bar{p}_j. \quad (3)$$

Proof. This follows from Lemma 11 and the argument of Lemma 26 in Cole et al. (2017). We include a proof in the Appendix for the sake completeness. \square

Lemma 13. ALG returns a 2 max NSW allocation.

Proof. Using the upper bound on the optimal integral NSW from Lemma 3, we want to show that

$$2 \left(\prod_{i=1}^n v_i(A_i) \right)^{1/n} \geq \left(\prod_{\bar{p}_j > 1} \bar{p}_j \right)^{1/n}.$$

From Lemma 12, it follows that

$$\left(\prod_{i \in N} v_i(A_i) \right)^{1/n} = \left(\prod_{t \in T} \prod_{i \in t} v_i(A_i) \right)^{1/n} \geq \frac{1}{2} \left(\prod_{\bar{p}_j > 1} \bar{p}_j \right)^{1/n},$$

since $\sum_{t \in T} k_t = n$. \square

Proof. (Theorem 6) The approximation guarantees follow immediately from Lemmas 7, 9, 10, and 13. \square

5. Examples

In this section we provide explicit examples to demonstrate that existing approximation algorithms fail to satisfy either: Prop-1, or a constant factor max NSW and MMS allocation.

Example 1. We show that the algorithm of Barman and Krishnamurthy (2019), which gives a Prop-1 and PO allocation, does not yield a α max NSW or β MMS allocation for any $\alpha, \beta > 0$. The algorithm relies on rounding a Fisher market equilibrium. Consider n identical agents, i.e., $v_i(\cdot) = v_j(\cdot), \forall i, j \in N$. We let v denote their common valuation function. There are n goods where $v(1) = n - 1 + 1/n$, and $v(j) = 1/n, \forall j \geq 2$. Clearly, the max NSW is positive. Also, $MMS_i = 1/n$ for each agent i , since $|M| = n$.

Consider the fractional allocation x , where $x_1 = (1/[n(n - 1 + 1/n)], 1, \dots, 1)$, and $x_i = (1/[n - 1 + 1/n], 0, \dots, 0), \forall i \geq 2$. In words, agents 2 through n spend their budget on good 1 and receive nothing else, while agent 1 receives a smaller portion of good 1 and all goods 2 through n . Define the prices $p_j = v(j)$, i.e., price of every good is equal to the agents' valuation. If all agents have a budget of 1, then it is easily verified that equilibrium conditions (Definition 1) hold. Therefore, (x, p) defines an equilibrium. However, regardless of rounding procedure, $n - 2$ agents receive nothing in the integral allocation.

Example 2. We show that the algorithm of Barman, Krishnamurthy et al. (2018) may not give a Prop-1 allocation. Recall, their algorithm yields an $(1 + \epsilon)$ EF-1 allocation, i.e., $(1 + \epsilon)v_i(A_i) \geq v_i(A_j \setminus g)$ for some $g \in A_j$. Fix $\epsilon > 0$, and let $k = \lceil 2/\epsilon \rceil$.




Figure 1: Spending graph for example 3.




Figure 2: Spending graph for NSW approximation lower bound.

Consider an instance with two identical agents and $|M| = k + 4$ goods, where $v(j) = 1/k$ for all $j \in \{1, \dots, k + 3\}$, and $v(k + 4) = 1$. Let the allocation be $A_1 = \{1, \dots, k + 3\}$, and $A_2 = \{k + 4\}$. Then $(1 + \epsilon)v(A_2) \geq 1 + 2/k = (k + 3 - 1)/k = v(A_1 \setminus j)$, $\forall j \in A_1$, so that the allocation is $(1 + \epsilon)$ EF-1. However, $v(M) = 2 + 3/k$. Then, $Prop_2 = 1 + 3/(2k) > 1 + 1/k = v_2(A_2) + v_{ij}$, $\forall j \in A_1$. We note that if ϵ is chosen small enough based on instance size, then the Barman, Krishnamurthy et al. (2018) algorithm yields an EF-1, hence Prop-1 allocation. However, the runtime may not be polynomial.

Example 3. We show that the algorithm of Cole and Gkatzelis (2015) may fail to yield a Prop-1 allocation. Figure 1 shows the spending graph of an SR equilibrium. Agents are shown as circles, and goods as squares. Directed arrows show the spending of each agent. All agents, except b , share the same valuations, shown below of each good. Agent b values all goods, except 7, the same as other agents. However, b values good 7 at $v_{b7} = 0.5$ as shown in parenthesis. Their algorithm starts by selecting a root agent for each tree. If agent a is picked, then goods 1, 2, and 3 are assigned to a , cutting the edge between b and good 3, and making b the root of the tree shown in the dashed box.

First all leaf goods are assigned to their parent agent's, creating a partial allocation A' . Then the algorithm computes a maximum weight matching to assign the remaining goods in the dashed box. Weights in the matching are set as $\log(v_i(A'_i) + v_{ij})$ between i and j s.t. $\bar{x}_{ij} > 0$. It is easily verified that all matchings have maximum weight, therefore b may not be assigned to any additional goods. In this case, b 's final allocation is good 4 with value $v_b(A'_b) = 0.4$. We compute $Prop_b = 11/12 > v_b(A_b) + v_{bj}$, $\forall j \in M$.

6. Tightness of Max NSW Approximation

In this section, we show that the 2 max NSW approximation is tight for our algorithm. We note that the following example is similar to the one found in Cole et al. (2017).

We describe a sequence of problem instances based on a parameter $z \in \mathbb{N}$. There are $n = z + 1$ agents and $m = 2z$ goods. Agent $i \in \{1, \dots, z\}$ values good i at $1/2$, good $z + i$ at $1/2 + 1/z$, and all other goods at 0. Agent $n = z + 1$ values goods 1 through z at 1, and all other goods at 0. Figure 2 shows the SR equilibrium, where agents are shown as circles,

goods as squares, and the directed edges show the spending of each agent. Note that goods 1 to z have price $1/2$, while goods $z + 1$ through $2z$ have price $1/2 + 1/z$.

Suppose ALG selects agent n as the root of the tree. Observe that all goods 1 through z are assigned to her. Each agent $i \in \{1, \dots, z\}$ receives good $z + i$, leading to the product of valuations $(1/2 + 1/z)^z$. However, consider an allocation in which agent n receives only one good $i \in \{1, \dots, z\}$, agent i receives good $z + i$, and all other agents $i' \in \{1, \dots, z\}$ get both goods i' and $z + i'$. In this allocation, the product of valuations is at least $1/2$. Therefore, the ratio of NSW between these allocations converges to 2 as $z \rightarrow \infty$.

7. Relating Prop-1 to Other Fairness Criteria

The large and ever growing number of fairness notions has spurred interest in understanding the relationships between these concepts, e.g., Amanatidis et al. (2018) and Caragiannis et al. (2016). To the best of our knowledge, the only work of this type concerning Prop-1 is that EF-1 implies Prop-1 (Conitzer et al., 2017). In this section, we provide a few additional results regarding Prop-1, and consider a few ways to strengthen this fairness metric.

Proposition 14. *An α -MMS allocation is not necessarily Prop-1.*

Proof. For clarity, we show a special case when $\alpha = 0.9$. The basic concept is straightforward to generalize to arbitrary $\alpha \in (0, 1)$, but requires additional details. Let $\alpha = 0.9$, and set $\delta = 0.05$. The instance contains $n \geq 2$ identical agents, i.e., $v_i(\cdot) = v_k(\cdot)$, $\forall i, k \in N$. Let $v(\cdot)$ denote their common valuation function. There are $m = n/\delta$ goods, each with value $v(j) = \delta$. Clearly, creating n bundles with $1/\delta = 20$ goods each gives an MMS allocation with $MMS_i = 1 = Prop_i$, $\forall i \in N$. Consider the allocation where agent 1 receives $18 = 1/\delta - 2$ goods, and the remaining goods are split among the other agents such that each agent receives at least $20 = 1/\delta$ goods. Agent 1 values her bundle at $v(A_1) = 18\delta = 0.9 = \alpha$ MMS. However, adding any good not allocated to her only gives $v(A_1) + v_{1j} = 19\delta = 0.95 < Prop_i$. Observe that the above example easily generalizes for any α of the form $\alpha = 1 - 10^{-k}$, for any $k \in \mathbb{N}$, by setting $\delta = 10^{-k}/2$. \square

Proposition 15. *Prop-1 does not imply α MMS for any $\alpha > 0$.*

Proof. Consider an instance with n identical agents and n goods, each with value $v(j) = 1$, $\forall j \in M$. Clearly, $Prop_i = 1$ for all agents, and $MMS_i = 1$ since $n = m$. Consider the allocation which assigns: agent 1 two goods, agents 2 through $n - 1$ one good each, and agent n no goods. It is easy verified that the allocation is Prop-1. However, it is not α MMS allocation for any $\alpha > 0$, since agent n receives a bundle worth 0. \square

Notice that the above allocation is neither EF-1, nor does it satisfy α max NSW for any $\alpha > 0$. Despite this, it is PO since the agents have identical valuations. This illustrates that even a Prop-1 and PO allocation fails to meet other common fairness metrics.

7.1 Strengthening Prop-1

The example of Proposition 15 highlights the potential for unfair Prop-1 allocations. We note that, other fairness notions might lead to unfair outcomes as well. For example, it is straightforward to construct EF-1 allocations that are only fair in the sense that they are

uniformly disliked by all agents.² Strengthening a fairness metric provides one possible way to address such issues. For instance, Caragiannis et al. (2016) proposed a stronger version of EF-1, called EF-X, in which each agent weakly prefers her bundle after removing *any* good from another agent's bundle. Next, we examine the possibility of stronger versions of Prop-1.

α **Prop-X.** Let $A = (A_i)_i^n$ be an integral allocation. We define a Proportional up to least valued good (Prop-X) as

$$v_i(A_i) + \min_{j \in M \setminus A_i} v_{ij} \geq v_i(M)/n.$$

In words, adding *any* good not allocated to the agent gives the agent at least her $Prop_i$ value. For any $\alpha \in (0, 1]$, define an α Prop-X allocation as

$$v_i(A_i) + \min_{j \in M \setminus A_i} v_{ij} \geq \alpha v_i(M)/n.$$

Note that the most crucial difference between α Prop-X and Prop-1 is the use of $\min_{j \in M \setminus A_i} v_{ij}$ as opposed to some good in the definition of Prop-1. This simple change has significant consequences.

Proposition 16. *An α Prop-X allocation might not exist for any $\alpha \in (0, 1]$.*

Proof. Fix $\alpha \in (0, 1]$. Consider an instance with three identical agents and three goods: a diamond (d), and two stones (s); with values: $v(d) = C = 9/\alpha$, and $v(s) = 1$, respectively. Suppose each agent receives one good. Observe that the allocation is EF-1, and therefore Prop-1. Now consider an agent receiving a stone. Her $Prop_i = (C + 2)/3 > 3/\alpha$, so that $\alpha Prop_i > 3$. Thus, the allocation is not α Prop-X since $v_i(s) + \min_{j \in M \setminus s} v_{ij} = 2$. Further, it is easily verified through similar computations that *no* allocation of these goods yields an α Prop-X allocation. \square

Notice that the primary issue with the α Prop-X definition lies in the requirement that adding *any* good not allocated to the agent gives at least $\alpha Prop_i$, since $\alpha \in (0, 1]$. In light of Proposition 16, we consider a different way to strengthen the Prop-1 guarantee.

GProp-1. Let $A = (A_i)_i^n$ be an integral allocation. For any subset of agents $S \subseteq N$, define $A(S) = \cup_{i \in S} A_i$, i.e., the goods allocated to the agents of S . We define a group Prop-1 (GProp-1) allocation as

$$\forall S \subseteq N, \quad v_i(A_i) + \max_{j \in A(S) \setminus A_i} v_{ij} \geq v_i(A(S))/|S|, \quad \forall i \in S.$$

In words, for any subset of agents a GProp-1 allocation satisfies the Prop-1 property if the instance consisted of only those agents and the goods allocated to them.

By definition, a GProp-1 allocation satisfies Prop-1. As a first result, we show that GProp-1 is a stronger notion than Prop-1.

Proposition 17. *Prop-1 does not imply GProp-1.*

2. Consider n identical agents and n goods. Suppose agent i only likes good i , but the allocation assigns this good to agent $j \neq i$. It is straightforward to verify that this allocation is in fact EF-X.

Proof. We use modified version of the example in Proposition 15. Consider an instance with $n > 2$ identical agents and n goods, each with value $v(j) = 1, \forall j \in M$. Note that, $Prop_i = 1$ for all agents. Consider the allocation which assigns all goods to one agent, say $i = 1$. Then, the allocation is Prop-1. Indeed, for any agent $i \in [2, \dots, n]$, we have $v(A_i) + v_{ij} = 1 = Prop_i, \forall j \in A_1$. However, the allocation is not GProp-1. For example, consider the subset $S = 1 \cup 2$. For this subset, $A(S) = M$. We compute $Prop_i = v(A(S))/|S| = n/2$ for $i = 1, 2$. However, for agent 2 we have $v(A_2) + v_{ij} = 1, \forall j \in A_1$. \square

Next, we consider the existence and computation of a GProp-1 allocation.

Proposition 18. *An EF-1 allocation is also GProp-1. Thus, GProp-1 allocations always exist, and can be found in polynomial time.*

Proof. We note that the following proof is similar to that of EF-1 implies Prop-1 found in Conitzer et al. (2017).

Let A be an EF-1 allocation. Choose a subset $S \subseteq N$, and pick an agent $i \in S$. By definition of EF-1

$$v_i(A_i) \geq v_i(A_k) - \max_{j \in A_k} v_{ij}, \forall k \in S.$$

Let $v_{i^*} = \max_{j \in A(S) \setminus A_i} v_{ij}$. Summing over all $k \in S$ yields

$$|S|v_i(A_i) \geq v_i(A(S)) - \sum_{j \in S \setminus i} \max_{k \in A_j} v_{ij} \geq v_i(A(S)) - |S|v_{i^*}.$$

Therefore, the allocation is GProp-1.

Since an EF-1 allocation can be found in polynomial time, e.g., using the algorithm of Lipton et al. (2004), we have that a GProp-1 allocation can be found in polynomial time. \square

We note that the allocation described in the example of Proposition 15 satisfies GProp-1. Therefore, a GProp-1 allocation is not necessarily α MMS for any $\alpha \in (0, 1)$, or α max NSW for any $\alpha > 0$, nor is it EF-1. In addition, since every Prop-1 allocation is GProp-1, Proposition 14 shows α MMS does not imply GProp-1 for any $\alpha \in (0, 1)$.

Finally, we show that our algorithm fails to give a GProp-1 allocation.

Example 4. Figure 3 shows the spending graph of an SR equilibrium. As before, we show agents as circles and items as squares. The agents have identical valuations, shown below each good. Directed edges show the spending of each agent. Note that agent 3 on the right purchases $1/(m-3)$ identical goods with value $1/(m-3)$ each. Setting prices equal to valuations gives an SR equilibrium where each agent spends her budget of 1. Further, since all prices are less 1, this is also a Fisher market equilibrium. Therefore, the example also applies to the algorithm of Barman and Krishnamurthy (2019) as well.

Note that $Prop_i = 1$. Clearly, the allocation is Prop-1 regardless of the assignment of good 2. However, observe that agent 3, assigned many items with low value, blocks the possibility of rounding the fractional allocation into a GProp-1 allocation. For example, consider assigning good 2 to agent 1. Taking the subset of agents as $S = 2 \cup 3$, gives $A(S) = \{3, \dots, m\}$. Thus, the allocation fails to meet the GProp-1 condition, since all items of agent 3 have relatively small value for sufficiently large m .




Figure 3: Spending graph for the counter example to GProp-1.

8. Discussion

We used a market based approach to approximate max NSW allocations. First, we gave novel bounds on any agent’s Prop and MMS guarantees in terms of market prices. Then, we designed a new scheme to round an SR equilibrium into a 2 max NSW, Prop-1, $1/(2n)$ MMS, and PO allocation. Thus, our algorithm makes progress towards an approximate max NSW allocation which more closely resembles all of the remarkable fairness properties of a true max NSW allocation.

The most significant difference between our approximation and a true max NSW allocation is the fact that the latter is EF-1. We note that computing an EF-1 and PO allocation in polynomial time remains a difficult open problem in fair division of indivisible goods, even though it is straightforward to find an allocation satisfying either condition independently³. The major challenge lies in verifying Pareto optimality. Indeed, how does one confirm that *no other allocation* increases the total value received by one agent without harming another? Most existing work relies on the celebrated First Welfare Theorem, Theorem 2, which states that *fractional* market equilibria are efficient. While rounding a market equilibrium preserves PO, e.g. Lemma 7, Barman, Krishnamurthy et al. (2018) show that there exists SR equilibria where *no rounding procedure* yields an EF-1 allocation. Thus, the classical approximation algorithm approach of solving a convex relaxation of the problem, e.g. an SR equilibrium, followed by a rounding step seems insufficient for obtaining an EF-1 and PO allocation.

The recent work of Caragiannis et al. (2019) obtains a constant factor approximation to max NSW and EF-1 allocation by relaxing the efficiency requirement. As mentioned in Section 1.2, the authors propose an algorithm which starts with an α max NSW allocation, *removes* a ‘small’ subset of goods from the problem, and then gives an EF-X and 2α max NSW allocation in the reduced instance. As discussed in Caragiannis et al. (2019), one can ‘add’ the removed goods back using the local search algorithm of Lipton et al. (2004) to yield a 2α max NSW and EF-1 allocation. Unfortunately, this technique degrades the allocation’s max NSW approximation, and loses Pareto optimality. It is not clear how to add the removed items back in a way that preserves PO.

Aside from EF-1 and PO guarantees, another avenue to pursue is obtaining an allocation which gives a constant factor approximation to both max NSW and MMS. Observe that α max NSW serves as a global objective, i.e., places requirements on all agents’ valuations of their bundles, while α MMS provides local guarantees, i.e., an agent specific objective. Intuitively, this gives conflicting optimization objectives, though we lack an explicit example

3. In the case of additive valuations, simply assigning goods to the agent with the highest valuation yields a PO allocation.

to illustrate the difficulty of the problem. While a true max NSW allocation is only $\Omega(n^{-1/2})$, we still find this an interesting and challenging problem. In addition, we note that to the best of our knowledge, no constant factor MMS approximation algorithm is PO. Thus, obtaining a constant factor max NSW, α MMS, and PO allocation together seems extremely difficult.

Finally, we address the issue of finding a GProp-1 and PO allocation. As shown in Example 4 in Section 7, there exist instances for which no rounding algorithm (of either an SR or Fisher market equilibrium) yields a GProp-1 allocation. While GProp-1 is a weaker notion than EF-1, e.g., Proposition 14, it seems nearly as difficult to obtain in conjunction with PO.

Acknowledgments

Work on this paper supported by NSF CRII Award 1755619. We thank the anonymous IJCAI 2019 reviewers, as well as the JAIR reviewers for providing advice to improve our work.

Appendix: Proof of Lemma 12

We note that the following proof is Lemma 26 found in Cole et al. (2017). We include it here for the reader's convenience.

Proof. Let k_1 be the number of agents such that $v_i(A_i) \geq 1$. Suppose i is one of these agents, and let j be the good she is assigned from the matching tree. Observe that $v_i(A_i) \geq \max(1, \bar{p}_j)$. Therefore, the product of these agents valuations is at least $\prod_{j \in M_t: \bar{p}_j > 1} \bar{p}_j$, so it is enough to show that product of the remaining $k_2 = k_t - k_1$ agents' valuations is at least 2^{-k_t} .

Clearly, the product of valuations for agents receiving $v_i(A_i) \in [1/2, 1)$, is minimized when at most one of these agents receives a bundle worth more than $1/2$. Let v_α be the value of that agent's bundle. Since the total spending on goods in the tree t is at least $k_t - 1/2$, then

$$v_\alpha \geq k_t - 1/2 - (k_1(1 + v_r) + v_r + (k_t - k_1 - 2)/2).$$

Let $\hat{k}_1 = \frac{k_t - 2v_r}{1 + 2v_r}$, and $\delta = \hat{k}_1 - k_1 = \frac{k_t - 2v_r}{1 + 2v_r} - \lfloor \frac{k_t - 2v_r}{1 + 2v_r} \rfloor$, be the rounding error. It follows that $v_\alpha \geq (1 + \delta)/2$, and, therefore

$$\prod_{i \in t} v_i(A_i) \geq \frac{v_\alpha v_r}{2^{k_t - k_1 - 2}} \prod_{j \in M_t: \bar{p}_j > 1} \bar{p}_j \geq \frac{v_r}{2^{k_t - \hat{k}_1 - 1}} \frac{1 + \delta}{2^\delta} \prod_{j \in M_t: \bar{p}_j > 1} \bar{p}_j \geq \frac{v_r}{2^{k_t - \hat{k}_1 - 1}} \prod_{j \in M_t: \bar{p}_j > 1} \bar{p}_j,$$

where the final inequality uses the fact that $1 + \delta \geq 2^\delta$, for $\delta \in [0, 1]$. Thus, it is enough to show that

$$\frac{v_r}{2^{k_t - \hat{k}_1 - 1}} \geq \frac{1}{2^{k_t}}, \quad \text{or} \quad v_r 2^{\frac{k_t + 1}{1 + 2v_r}} \geq 1.$$

Since $v_r \in [1/(2k_t), 1/2]$, by Lemma 8, then

$$v_r 2^{\frac{k_t + 1}{1 + 2v_r}} \geq \frac{1}{2k_t} 2^{\frac{k_t + 1}{1 + 2v_r}} \geq \frac{1}{2k_t} 2^{\frac{k_t + 1}{2}} \geq 1, \quad \text{for } k \geq 7.$$

Next, observe that $v_r 2^{\frac{k_t+1}{1+2v_r}}$ is minimized at the same points as $\log v_r + \frac{k_t+1}{1+2v_r}$. Taking the derivative with respect to v_r gives

$$\left(\log v_r + \frac{k_t + 1}{1 + 2v_r} \right)' = \frac{1}{v_r \ln 2} - \frac{2(k_t + 1)}{(1 + 2v_r)^2}. \quad (4)$$

For $k \leq 4$, (4) is positive for all v_r . Thus, $v_r 2^{\frac{k_t+1}{1+2v_r}}$ is minimized at $v_r = 1/(2k_t)$, where it takes the value $\frac{1}{2k_t} 2^{\frac{k_t(k_t+1)}{k_t+1}} = \frac{2^{k_t}}{2k_t} \geq 1$. For the cases of $k_t = 5, 6$, one can substitute the value of k_t into (4), and verify that $v_r 2^{\frac{k_t+1}{1+2v_r}}$ is minimized at $v_r = 1/10$ and $v_r = 1/12$ respectively. Further, $v_r 2^{\frac{k_t+1}{1+2v_r}} \geq 1$ in both cases. \square

References

- Amanatidis, G., Birmpas, G., & Markakis, E. (2018). Comparing approximate relaxations of envy-freeness. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pp. 42–48.
- Amanatidis, G., Markakis, E., Nikzad, A., & Saberi, A. (2017). Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms (TALG)*, 13(4), 52.
- Anari, N., Mai, T., Gharan, S. O., & Vazirani, V. V. (2018). Nash social welfare for indivisible items under separable, piecewise-linear concave utilities. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2274–2290.
- Barman, S., Biswas, A., Krishnamurthy, S. K., & Narahari, Y. (2018). Groupwise maximin fair allocation of indivisible goods. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pp. 917–924.
- Barman, S., & Krishnamurthy, S. K. (2017). Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 647–664.
- Barman, S., & Krishnamurthy, S. K. (2019). On the proximity of markets with integral equilibria. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, pp. 1748–1755.
- Barman, S., Krishnamurthy, S. K., & Vaish, R. (2018). Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 557–574.
- Baumeister, D., Bouveret, S., Lang, J., Nguyen, N., Nguyen, T. T., Rothe, J., & Saffidine, A. (2017). Positional scoring-based allocation of indivisible goods. *Autonomous Agents and Multi-Agent Systems*, 31(3), 628–655.
- Bouveret, S., & Lemaître, M. (2014). Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2), 259–290.

- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 1061–1103.
- Caragiannis, I., Gravin, N., & Huang, X. (2019). Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pp. 527–545.
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., & Wang, J. (2016). The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 305–322.
- Chaudhury, B. R., Cheung, Y. K., Garg, J., Garg, N., Hofer, M., & Mehlhorn, K. (2018). On fair division for indivisible items. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, pp. 25:1–25:17.
- Cole, R., Devanur, N. R., Gkatzelis, V., Jain, K., Mai, T., Vazirani, V. V., & Yazdanbod, S. (2017). Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the Eighteenth ACM Conference on Economics and Computation*, pp. 459–460.
- Cole, R., & Gkatzelis, V. (2015). Approximating the Nash social welfare with indivisible items. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 371–380.
- Conitzer, V., Freeman, R., & Shah, N. (2017). Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 629–646.
- Darmann, A., & Schauer, J. (2015). Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2), 548–559.
- Duan, R., Garg, J., & Mehlhorn, K. (2016). An improved combinatorial polynomial algorithm for the linear Arrow-Debreu market. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pp. 90–106.
- Eisenberg, E., & Gale, D. (1959). Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1), 165–168.
- Garg, J., Hofer, M., & Mehlhorn, K. (2018). Approximating the Nash social welfare with budget-additive valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 2326–2340.
- Garg, J., & McGlaughlin, P. (2019). Improving Nash social welfare approximations. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 294–300.
- Garg, J., McGlaughlin, P., & Taki, S. (2019). Approximating maximin share allocations. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019*, pp. 20:1–20:11.
- Garg, J., & Taki, S. (2019). An improved approximation algorithm for maximin shares. *CoRR*, abs/1903.00029.
- Ghodsi, M., Hajiaghayi, M., Seddighin, M., Seddighin, S., & Yami, H. (2018). Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 539–556.

- Kurokawa, D., Procaccia, A. D., & Wang, J. (2016). When can the maximin share guarantee be guaranteed?. In *AAAI*, Vol. 16, pp. 523–529.
- Lee, E. (2017). Apx-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122, 17–20.
- Lipton, R. J., Markakis, E., Mossel, E., & Saberi, A. (2004). On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on Electronic commerce*, pp. 125–131.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. Oxford university press, New York.
- Nguyen, T. T., & Rothe, J. (2014). Minimizing envy and maximizing average nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics*, 179, 54–68.
- Procaccia, A. D., & Wang, J. (2014). Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 675–692.
- Segal-Halevi, E., & Suksompong, W. (2018). Democratic fair allocation of indivisible goods. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pp. 482–488.
- Steinhaus, H. (1948). The problem of fair division. *Econometrica*, 16, 101–104.
- Varian, H. (1974). Equity, envy, and efficiency. *Journal of Economic Theory*, 29(2), 217–244.