

Improving Nash Social Welfare Approximations

Jugal Garg and Peter McGlaughlin

University of Illinois at Urbana-Champaign

{jugal, mcglgh12}@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, yields an allocation that is not only envy-free and proportional [Varian, 1976], 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 *et al.*, 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, 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.

1.2 Related Work

We focus on three of the most common notions of a fair allocation: NSW, MMS, and Prop-1.¹

NSW. [Cole and Gkatzelis, 2015] gave the first constant factor (2.89) approximation algorithm to the NSW objective. Later, [Cole *et al.*, 2017] refined the analysis of the algorithm, showing it yields a 2 max NSW allocation. More recently, [Barman *et al.*, 2018] designed an FPTAS that yields a 1.45 max NSW, PO, and a $(1 + \epsilon)$ EF-1 allocation.

Prop-1. [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. This generalizes the private goods case we consider. The authors demonstrate a polynomial time procedure to achieve a Prop-1 allocation, but leave the problem of finding a Prop-1 and PO allocation in polynomial time. The recent work of [Barman and Krishnamurthy, 2019] resolves this issue for private good setting by providing an algorithm that yields a Prop-1 and PO allocation.

MMS. [Budish, 2011] defined maximin share which gives an intuitive local measure of fairness of an allocation. [Bouveret and Lemaître, 2014] showed that an MMS allocation exists in certain special settings. [Procaccia and Wang, 2014; Kurokawa *et al.*, 2016] obtained the surprising result that MMS allocations might not exist, and showed the existence of a $2/3$ MMS allocation. [Amanatidis *et al.*, 2017] obtained a PTAS which finds a $(2/3 - \epsilon)$ MMS allocation. Later, [Barman and Krishna Murthy, 2017] and [Garg *et al.*, 2019] obtained simple algorithms to find a $2/3$ MMS allocation. More recently, [Ghodsi *et al.*, 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 none of the above algorithms yields a constant factor approximation to max NSW, Prop-1, $1/(2n)$ MMS, and PO allocation simultaneously in poly-time.

¹[Amanatidis *et al.*, 2018] survey the current landscape of fair division and the relationship between various fairness concepts.

2 Preliminaries

We consider the fair allocation of a set M of m indivisible goods among a set N of n agents. 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 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 $\alpha \geq 1$.

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

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

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 $v_i(A_i) \geq \alpha MMS_i$.

Prop-1 Proportionality (Prop) requires that each agent receive a bundle their value as 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 a good not allocated to the agent.

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

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 market equilibrium. 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 [1959] where $e_i = 1, \forall i$.

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

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 $\sum_{j=1}^m v_{ij} x_{ij}$.

A fractional allocation x and a price vector $p = (p_j)_{j \in M}$ correspond to primal and dual variables of (EG) respectively. A pair (\bar{x}, \bar{p}) which solve (EG) defines an equilibrium.

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

We define the set of MBB goods for agent i as

$$MBB_i = \{j \in M : j \in \arg \max_{k \in M} v_{ik} / \bar{p}_k\}. \quad (1)$$

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 et al., 2016] *It is always possible to rearrange the agents' spending to ensure the spending graph is a forest.*

Theorem 2. [Mas-Colell et al., 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. A spending restricted (SR) equilibrium is a fractional allocation \bar{x} and prices \bar{p} such that: (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 equilibrium 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 $v_{ij} \leftarrow c v_{ij}, \forall j \in M$ doesn't change the problem. Define i 's MBB ratio $\alpha_i := \max_j v_{ij} / \bar{p}_j$. 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. Given an equilibrium 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 provides 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 and 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 $H = \{\bar{p}_j \in M : \bar{p}_j > 1\}$ be the set of high priced goods, and $L = M \setminus H$ be the set of low priced goods.

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

Proof. 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 $|H|$. Therefore, since 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 $MM S_i^{n-1}(M \setminus j) \geq MM S_i^n(M)$.

Let $B = (B_i)_{i \in N}$ be the bundles i selects when calculating her $MM S_i^n(M)$, i.e., $v_i(B_k) \geq MM S_i^n(M), \forall B_k \in B$. 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, $v_i(B'_k) \geq v_i(B_k) \geq MM S_i^n(M), \forall k \leq n-1$.

Clearly, the above argument generalizes to the case of removing k agents and k goods, i.e., $MM S_i^{n-k}(M \setminus \{j_1, \dots, j_k\}) \geq MM S_i^n(M)$, for any $k < n$ and $\{j_l\}_{l=1}^k \subseteq M$. Applying this to the set of high valued goods H yields $MM S_i^n(M) \leq MM S_i^{n-|H|}(M \setminus H) = MM S_i^{n-|H|}(L)$. Further, $MM S_i^{n-|H|}(L) \leq v_i(L)/(n-|H|)$. Then, using (2), $MM S_i^n(M) \leq \frac{v_i(L)}{n-|H|} \leq \frac{1}{n-|H|} \sum_{j \in L} p_j \leq 1$. \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 for the prices \bar{p} , and that the spending graph is a forest, by Lemma 1. Let T denote the forest in the spending graph. 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 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 edges connecting an agent i to goods $j \in M_t$ are in fact $MM B_i$.

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.

Algorithm 1: Rounding Algorithm

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

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

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 satisfy 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 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 MMB 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 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 the root 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, 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 before Step 4. 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 $v_i(A_i) < 1$ since i receives a good $j \in N_i$. It follows that for any of the agents receiving a bundle $v_i(A_i) \geq 1$, the total spending on the goods of A_i is no more than $1 + v_r$, i.e., $\sum_{j \in A_i} \bar{p}_j \leq 1 + v_r$.

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]. \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 $x_1 = (1/\lceil n(n - 1 + 1/n) \rceil, 1, \dots, 1)$, and $x_i = (1/\lceil n - 1 + 1/n \rceil, 0, \dots, 0), \forall i \geq 2$. 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. 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.




Figure 1: Spending graph for Example 3.

Example 2. We show that the algorithm of [Barman 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 $k = \lceil 2/\epsilon \rceil$.

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 [Barman et al., 2018] algorithm yields an EF-1, hence Prop-1 allocation. However, the runtime may not be polynomial.

Example 3. We show that [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 Conclusions and Future Work

We presented 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. One interesting direction for future work is improving the approximation to meet EF-1 guarantees. Another avenue to pursue is obtaining a constant factor NSW and MMS allocation.

Acknowledgments

Work on this paper supported by NSF CRII Award 1755619.

References

- [Amanatidis *et al.*, 2017] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms (TALG)*, 13(4):52, 2017.
- [Amanatidis *et al.*, 2018] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. Comparing approximate relaxations of envy-freeness. In *IJCAI*, pages 42–48, 2018.
- [Barman and Krishna Murthy, 2017] Siddharth Barman and Sanath Kumar Krishna Murthy. Approximation algorithms for maximin fair division. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 647–664. ACM, 2017.
- [Barman and Krishnamurthy, 2019] Siddharth Barman and Sanath Kumar Krishnamurthy. On the proximity of markets with integral equilibria. In *AAAI*, 2019.
- [Barman *et al.*, 2018] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 557–574. ACM, 2018.
- [Bouveret and Lemaître, 2014] Sylvain Bouveret and Michel Lemaître. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2):259–290, 2014.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 305–322. ACM, 2016.
- [Cole and Gkatzelis, 2015] Richard Cole and Vasilis Gkatzelis. Approximating the Nash social welfare with indivisible items. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 371–380. ACM, 2015.
- [Cole *et al.*, 2017] Richard Cole, Nikhil R Devanur, Vasilis Gkatzelis, Kamal Jain, Tung Mai, Vijay V Vazirani, and Sadra Yazdanbod. Convex program duality, Fisher markets, and Nash social welfare. In *Proceedings of the Eighteenth ACM Conference on Economics and Computation*, pages 459–460, 2017.
- [Conitzer *et al.*, 2017] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 629–646. ACM, 2017.
- [Duan *et al.*, 2016] Ran Duan, Jugal Garg, and Kurt Mehlhorn. 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*, pages 90–106, 2016.
- [Eisenberg and Gale, 1959] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The parimutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [Garg and Taki, 2019] Jugal Garg and Setareh Taki. An improved approximation algorithm for maximin shares. *CoRR*, abs/1903.00029, 2019.
- [Garg *et al.*, 2019] Jugal Garg, Peter McGlaughlin, and Setareh Taki. Approximating maximin share allocations. In *2nd Symposium on Simplicity in Algorithms*, 2019.
- [Ghodsí *et al.*, 2018] Mohammad Ghodsí, Mohammadtaghi Hajiaghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 2018 ACM Conference on Economics and Computation, EC ’18*, pages 539–556, New York, NY, USA, 2018. ACM.
- [Kurokawa *et al.*, 2016] David Kurokawa, Ariel D Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *AAAI*, volume 16, pages 523–529, 2016.
- [Lee, 2017] Euiwoong Lee. Apx-hardness of maximizing Nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.
- [Mas-Colell *et al.*, 1995] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- [Procaccia and Wang, 2014] Ariel D Procaccia and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 675–692. ACM, 2014.
- [Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [Varian, 1976] Hal R Varian. Two problems in the theory of fairness. *Journal of Public Economics*, 5(3-4):249–260, 1976.