

Fair Division with Social Impact

Michele Flammini^{1,2}, Gianluigi Greco², Giovanna Varricchio²

¹Gran Sasso Science Institute, L'Aquila, Italy

²University of Calabria, Rende, Italy

michele.flammini@gssi.it, gianluigi.greco@unical.it, giovanna.varricchio@unical.it

Abstract

In this paper, we consider the problem of fair division of indivisible goods when the allocation of goods impacts society. Specifically, we introduce a second valuation function for each agent, determining the social impact of allocating a good to the agent. Such impact is considered desirable for the society – the higher, the better. Our goal is to understand how to allocate goods fairly from the agents' perspective while maintaining society as happy as possible. To this end, we measure the impact on society using the utilitarian social welfare and provide both possibility and impossibility results. Our findings reveal that achieving good approximations, better than linear in the number of agents, is not possible while ensuring fairness to the agents. These impossibility results can be attributed to the fact that agents are completely unconscious of their social impact. Consequently, we explore scenarios where agents are socially aware, by introducing related fairness notions, and demonstrate that an appropriate definition of fairness aligns with the goal of maximizing the social objective.

Introduction

Fair division is a fundamental research area at the intersection of economics and computer science. In the last decade, it turned out to be a particularly active field due to its applications to several real-world contexts. It focuses on distributing resources (goods) among individuals (agents) in a fair manner according to their valuations. To achieve fairness, various criteria have been studied, including *envy-freeness* (EF) (Foley 1966), where every agent weakly prefers her bundle over the bundle received by any other, and *proportionality* (PROP) (Steinhaus 1948), where every agent values her bundle at least as much as her proportional share, that is, the value she attributes to the whole set of goods normalized by the number of agents. Unfortunately, in indivisible settings, EF and PROP allocations do not always exist. Consequently, several relaxations have been considered. These include, but are not limited to, the *up to one* (Budish 2011), *up to any* (Caragiannis et al. 2019), *epistemic* (Aziz et al. 2018; Caragiannis et al. 2023), and *randomized* (Aziz et al. 2023) variants of EF, PROP, and related notions. The fair division model has also been extended to deal

with more complex scenarios, e.g., endowing agents with entitlements (Suksompong 2025; Chakraborty et al. 2021; Aziz, Ganguly, and Micha 2023; Hoefer, Schmalhofer, and Varricchio 2024) or introducing additional constraints (Bild et al. 2022; Barman et al. 2023; Bredereck, Kaczmarczyk, and Niedermeier 2022).

In most of this literature, the focus has been put on achieving fairness for the involved agents, possibly subject to efficiency, without considering any kind of externality. However, there are many scenarios in which allocating resources has an impact on society. We can imagine, for example, there are some green strategies that the government must assign to companies. A company will profit from implementing the strategy (as companies more sensitive to green aspects are increasingly appealing to customers), while the implementation will positively impact society by reducing emissions. Companies judge the outcome based solely on the profit they and other companies can attain, whereas society aims to maximize the overall reduction of emissions.

With this paper, we expand the scope of fair division of indivisible goods incorporating the concept of social impact. Our goal is twofold: To ensure fairness among agents and guarantee a high social impact.

Our Contribution. We refer to the preliminaries for the formal definition of the considered fairness criteria.

In this work, we aim at determining the loss in efficiency (expressed by the utilitarian welfare w.r.t. the social impact of agents, i.e., the sum of the social impact of agents) while pursuing fairness. In particular, we estimate the *Price of Fairness* (PoF), that is, the worst-case ratio between the value of the maximum utilitarian welfare and the maximum utilitarian welfare achievable by a fair solution. Besides determining suitable bounds for several fairness notions, we also provide efficient algorithms for computing the corresponding approximately optimal and fair allocations. An overview of our results for n agents is given in Table 1.

For all the fairness notions we considered, an approximation better than linear in the number of agents is not possible. This can be attributed to the fact that agents are unconscious of their social impact. We therefore investigate which kind of consciousness on the agents' side allows the existence of fair and optimal outcomes. To this end, we introduce the notion of *socially aware envy-freeness* (sEF) and establish that

Price of	EF1	EFX	EF2	PROP1	epistemic EF1
UB	$O(n^2)$ (Thm. 4) n^\dagger (Thm. 2)	n^* (Thm. 3)	n (Thm. 5)	n (Thm. 6)	n (Thm. 6)
LB	n (Thm. 1)	n (Cor. 1)	$n - 1$ (Thm. 1)	n (Cor. 2)	n (Cor. 2)

Table 1: Lower (LB) and upper (UB) bounds to the PoF w.r.t. the maximum utilitarian social impact. The results marked with \dagger and $*$ hold for ordered and identical valuations, respectively. All the LBs have been proven even under identical valuations.

an allocation maximizing the utilitarian social impact and sEF *up to one good* always exists. Due to space constraints, we refer to the full paper for the details.

Related Work. The literature on fair division is wide; we refer the interested reader to some recent surveys (Amanatidis et al. 2023; Suksompong 2021, 2025). We next discuss the work directly related to our.

The goal of achieving fairness while maximizing the utilitarian welfare has been extensively studied when $s_i = v_i$. Allocations having maximum agents’ utilitarian social welfare do not necessarily guarantee fairness, even in divisible settings. For this reason, the PoF with respect to such a welfare has been extensively studied (Nicosia, Pacifici, and Pferschy 2017; Barman, Bhaskar, and Shah 2020; Bei et al. 2021; Li et al. 2024). In this stream of research, several fairness notions as well as restrictions on the valuations have been exploited. Furthermore, the problem of determining whether there exists a fair allocation that is also socially optimal has been considered in (Aziz, Moulin, and Sandmirskiy 2020; Bu et al. 2022). In a nutshell, they show that for all the considered criteria, both determining the existence of a maximum utilitarian allocation which is also fair and computing an allocation that maximizes the utilitarian welfare among fair solutions are strongly NP-hard problems.

The idea of employing two distinct valuation functions for each agent has been examined in prior research (Bu et al. 2023) albeit from a different viewpoint. In this approach, one valuation function captures the agent’s preferences while the other represents the allocator’s interpretation of the agent’s preferences – what the allocator thinks a bundle of goods should value for a certain agent. The main objective of Bu et al. was to achieve fairness in both the agents’ and the allocator’s perspectives. In our model, society does not care if the allocation is fair according to the social impact functions; rather, society would like the allocation to maximize its own interests. In (Bu et al. 2023), the authors briefly discuss the problem of maximizing the allocator’s utilitarian welfare by providing a simple m -approximation algorithm, where m is the number of goods. With our contribution, we improve this result achieving an approximation polynomial in the number of agents which gives a significant improvement as the number of goods might be extremely larger than, even exponential in, the number of agents.

Preliminaries

In this section, we provide the basics of fair division as well as a formal description of the framework with social impact. Hereafter, we denote by $[t]$, for $t \in \mathbb{N}$, the set $\{1, \dots, t\}$.

In fair division of indivisible goods, we are given a set of n agents \mathcal{N} and a set of m goods \mathcal{G} , also called *items*, that have to be allocated to agents in a fair manner. An allocation \mathcal{A} consists of disjoint bundles (A_1, \dots, A_n) , where each A_i is the set of goods, the bundle, allocated to agent i . Unless stated otherwise, we assume allocations to be complete, meaning that all goods have been allocated, i.e., $\cup_{i \in \mathcal{N}} A_i = \mathcal{G}$. If an allocation is not complete, we call it *partial*. The fairness of an allocation is determined according to the agents’ preferences; for each agent $i \in \mathcal{N}$ there exists a *valuation function* $v_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$ mapping bundles of goods to non-negative real values. A fair division instance is defined by the triple $\mathcal{I} = (\mathcal{N}, \mathcal{G}, \{v_i\}_{i \in \mathcal{N}})$. In this paper, we focus on *additive valuations*, i.e., for each bundle $A \subseteq \mathcal{G}$, $v_i(A) = \sum_{g \in A} v_i(\{g\})$. For the sake of simplicity, we write $v_i(g)$ to denote $v_i(\{g\})$. Moreover, we assume $v_i(\emptyset) = 0$.

The most relevant and intuitive notion of fairness is the one of *envy-freeness* (EF), which postulates that, in a given allocation \mathcal{A} , no agent envies the bundle of any other agent, i.e., for each $i, j \in \mathcal{N}$, $v_i(A_i) \geq v_i(A_j)$. Unfortunately, such an allocation might not exist even in simple cases: Consider an instance with two agents and one good, no matter who will receive the good, the other will be envious. Thus, several relaxations have been proposed. Among the others:

An allocation \mathcal{A} is said to be *envy-free up to one good* (EF1) if for each pair i, j of agents, either $A_j = \emptyset$, or $\exists g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

An allocation \mathcal{A} is said to be *envy-free up to any good* (EFX) if for each pair of agents i, j either $A_j = \emptyset$ or, for each $g \in A_j$, $v_i(A_i) \geq v_i(A_j \setminus \{g\})$.

While EF1 allocations are known to always exist, even for monotone valuations, the existence of EFX allocations remains a fundamental open question for additive valuations. The EF1 definition has been further extended allowing the hypothetical removal of at most k goods from the envied bundle, in this case, we talk about *envy-free up to k goods* (EF k).

Besides comparison-based fairness notions like EF, shared-based ones have also been introduced. Among these, *proportionality* (PROP) guarantees to each agent at least her proportional share; formally, for each $i \in \mathcal{N}$, $v_i(A_i) \geq \text{PS}_i := \frac{1}{n} \cdot v_i(\mathcal{G})$. Achieving proportionality, as for EF, might not be possible. Hence, *proportional up to one good* (PROP1), the analogous of EF1 for proportionality, has been proposed. \mathcal{A} is said to be PROP1 if for each i there exists $g \notin A_i$ such that $v_i(A_i \cup \{g\}) \geq \text{PS}_i$. For additive valuations, it is known $\text{EF} \Rightarrow \text{PROP}$ and $\text{EF1} \Rightarrow \text{PROP1}$.

We finally define the epistemic variant of EF1; \mathcal{A} is said to be *epistemic EF1* if, for each $i \in \mathcal{N}$, there exists an al-

location \mathcal{A}^* such that $A_i^* = A_i$ and \mathcal{A}^* is EF1 in i 's perspective. Such an allocation \mathcal{A}^* is also called the epistemic EF1 certificate (for agent i). In other words, for each agent, there exists an allocation where the bundle A_i makes her EF1 satisfied. Because of the existence of this certificate for each agent i , being EF1 \Rightarrow PROP1, we know that A_i makes i PROP1 (as share-based notions solely depend on the received bundle); hence, epistemic EF1 \Rightarrow PROP1.

In the sequel, we focus on additive valuations. An interesting subclass is the one of *ordered valuations*, where it is further assumed that there exists a fixed ordering of the goods g^1, \dots, g^m such that, for each agent i and $h, k \in [m]$ with $h < k$, it holds $v_i(g^h) \geq v_i(g^k)$. A special case of ordered instances is the one of *identical valuations*, where there exists an additive function v s.t. $v_i = v$, for each $i \in \mathcal{N}$.

To understand our algorithms, we first overview standard approaches to compute EF1 allocations.

The Envy-Cycle Elimination. The envy-cycle elimination algorithm, introduced by (Lipton et al. 2004), starts with an empty allocation \mathcal{A} . At each round, one available good g is allocated to an agent i who is not envied by any other agent. After allocating the good g , if there exists an agent j envying agent i , the EF1 is satisfied as agent j was not envious before the allocation of g . To ensure the existence of such an agent i , there is a subroutine, the CYCLEELIMINATION, which takes as input the current partial allocation \mathcal{A} and the corresponding *envy-graph* G where nodes are the agents and there is a directed edge from i to j if i envies j in \mathcal{A} . The CYCLEELIMINATION finds a cycle in G and trades bundles along the cycle, that is, for each directed edge i, j in the cycle, agent i receives the bundle A_j , and it is run until there is at least one agent who is not envied by anyone else. CYCLEELIMINATION preserves the EF1 condition among agents, even for monotone valuations.

Sequential Allocation Algorithms. To describe sequential allocation algorithms, we first introduce the concept of *picking sequence*. A picking sequence $\sigma = (\sigma(1), \dots, \sigma(m))$ is a sequence of m entries where $\sigma(k) \in \mathcal{N}$, for each $k \in [m]$. A *prefix* σ^k , for $k \in [m]$, is the subsequence $(\sigma(1), \dots, \sigma(k))$. A sequence σ is said to be *recursively balanced* (RB) if for any prefix σ^k and any pair of agents i, j , denoted by p_i and p_j the number of occurrences of i and j in σ^k , respectively, it holds $|p_i - p_j| \leq 1$.

A sequential allocation algorithm takes as input a fair division instance and a picking sequence of length m , where m is the number of goods to be allocated, and proceeds in m steps. It starts with an empty allocation. At step k , agent $\sigma(k)$ choose the most preferred available good which is put in her bundle and removed from the set of available ones.

Fact 1 (Aziz et al. 2019). For additive valuations, if σ is RB, then, the sequential algorithm returns an EF1 allocation.

Given an ordering of the agents $1, \dots, n$, a well-known RB picking sequence is $\sigma = (1, \dots, n, 1, \dots, n, 1, \dots)$ whose corresponding sequential algorithm is the famous *round-robin* (RR) algorithm.

Social Impact. In this paper, we assume there exists an additive *social impact* function $s_i : 2^{\mathcal{G}} \rightarrow \mathbb{R}_{\geq 0}$, which rep-

resent the happiness of the society for allocating a bundle of goods to an agent $i \in \mathcal{N}$. We assume, the higher the social impact, the better is for society. We pay particular attention to the problem of maximizing the *utilitarian* social welfare, given by the sum of the agents' social impact, that is, $SW(\mathcal{A}) = \sum_i s_i(A_i)$. We denote by MAXUT the problem of maximizing the utilitarian welfare, by OPT any optimal allocation for a given instance of MAXUT and by $\text{opt} = SW(\text{OPT})$ its social welfare.

Impossibility Results

Let us start by observing that it is possible to efficiently compute an optimal solution of MAXUT if the EF1 condition is not a constraint. It is sufficient to allocate each good $g \in \mathcal{G}$ to an agent i^* who has the highest social impact for it, that is, to $i^* \in \arg \max_{i \in \mathcal{N}} s_i(g)$. Nonetheless, an approximation better than n is not possible for MAXUT while requiring EF1. This is a direct consequence of the following result.

Theorem 1. *An approximation better than $n - k + 1$ to opt is not possible when requiring EF k , even for identical agents.*

Proof. Consider an instance with n agents having identical valuation functions and $m = h \cdot n + k - 1$ goods, where $h = n - k$. All agents value 1 every good, agent 1 has a social impact of 1 for each good, and all remaining ones social impact 0. In such an instance $\text{opt} = m$ and it is obtained by assigning all goods to agent 1.

Concerning EF k allocations, no agent can receive more than $h + k - 1$ goods without violating the EF k condition, otherwise, there exist two bundles differing by more than k goods – a contradiction to EF k .

In conclusion, agent 1 cannot receive more than $h + k - 1$ goods in an EF k allocation, and hence, no approximation better than $\frac{m}{h+k-1} = \frac{h \cdot n + k - 1}{h+k-1}$ is possible, where the last equality holds as $h = n - k$. \square

Since in the just described instance all agents value 1 all goods, EF1 is equivalent to EFX showing the following.

Corollary 1. *An approximation better than n to MAXUT is not possible when requiring EFX.*

Moreover, in the instance described in the proof of Theorem 1, if $k = 1$, the only way to get a PROP1 or epistemic EF1 allocation is to allocate to each agent the same number of goods – in fact, for such instance, PROP1 is equivalent to PROP and PROP1 is a necessary condition for epistemic EF1. These observations together imply the following.

Corollary 2. *An approximation better than n to MAXUT is not possible when requiring PROP1 or epistemic EF1.*

Approximating MAXUT Subject to Fairness

Given the above impossibility results, we now investigate what is the best achievable approximation to opt when requiring fairness. We first focus on ordered and identical valuations, and then turn the attention to the broader class of additive valuations.

A General Approximation Algorithm for Ordered Valuations

In this section, we present a general approximation algorithm providing EF1 allocations in the case of ordered instances, proving the following theorem.

Theorem 2. *For ordered valuations, there exists a poly-time algorithm computing an EF1 allocation that is an n -approximation of opt .*

Assume w.l.o.g. $m = q \cdot n$, otherwise, we can add dummy goods of value and social impact 0 for all agents. The dummy goods will be ranked last from all agents maintaining the ordered valuations property. Let g^1, g^2, \dots, g^m the common ordering of the goods according to the ordered valuations assumption. We refer to g^k as the k -th good. Every sequential allocation algorithm we employ will assign g^k to the k -th agent in the picking sequence order. For our purposes, we also partition \mathcal{G} , according to their ordering, into blocks of cardinality n . In particular, we denote by $\mathcal{G}^h = \{g^{(h-1)n+1}, \dots, g^{hn}\}$, for $h \in [q]$. In other words, block \mathcal{G}^1 contains the n most preferred goods; block \mathcal{G}^2 contains the most n preferred goods among $\mathcal{G} \setminus \mathcal{G}^1$, and so forth.

Lemma 1. *Given an ordered fair division instance, if \mathcal{A} is an allocation in which each agent receives a bundle containing exactly one good in each \mathcal{G}^h , $\forall h \in [q]$, then, \mathcal{A} is EF1.*

This lemma can be proven noticing that an allocation satisfying such a condition is attainable by running a sequential algorithm with an appropriate RB picking sequence.

We next build our approximation algorithm, Algorithm 1, around the conditions of Lemma 1. The algorithm will take as input an allocation \mathcal{A} and transform it into an EF1 allocation \mathcal{A}' while maintaining suitable properties w.r.t. the social impact of \mathcal{A} . Specifically, Algorithm 1 proceeds into q rounds and at round $h \in [q]$ assigns the goods in \mathcal{G}^h , one for each agent. Firstly, for each agent i with $\mathcal{G}^h \cap A_i \neq \emptyset$, it puts in A'_i the best good according to the social impact of i among the ones in $\mathcal{G}^h \cap A_i$, that is, a good in $\arg \max_{g \in \mathcal{G}^h \cap A_i} s_i(g)$. The remaining agents with $\mathcal{G}^h \cap A_i = \emptyset$ will receive one of the remaining goods in \mathcal{G}^h . A more formal description is presented in Algorithm 1. Notice, the algorithm solely depends on the goods ordering and not on the agents' valuations.

Proposition 1. *If \mathcal{A} is the input of Algorithm 1 and \mathcal{A}' is the output, then, \mathcal{A}' is EF1. Moreover, $n \cdot s_i(A'_i) \geq s_i(A_i)$.*

Sketch. Algorithm 1, by design, computes an allocation satisfying the hypothesis of Lemma 1, and thus, \mathcal{A}' is EF1.

About the approximation, at each round, agent i receives the best, according to s_i , remaining goods of A_i while the other agents receive at most one good from A_i . Thus, the social impact of the good agent i receives is, in the worst case, an n -approximation to the goods of A_i allocated in this round. Summing up for all rounds the thesis follows. \square

Theorem 2 follows by Proposition 1 as the utilitarian welfare is the sum of the agents' utilities. Proposition 1 has broader implications on the egalitarian, Nash welfare, and in general on p -means, where an n -approximation can be guaranteed as well.

An n -Approx. Subject to EFX for Identical Agents

In this section, we further assume that agents have identical valuations and strengthen Theorem 2 as follows:

Theorem 3. *For identical valuations, there exists a poly-time algorithm computing an EFX allocation that is an n -approximation of opt .*

Sketch. We recall that, for identical valuations, an EFX allocation can be computed in polynomial time (Barman, Krishnamurthy, and Vaish 2018). Let \mathcal{A} be such an allocation. Since agents have identical valuations, any permutation of bundles of \mathcal{A} remains EFX. We can therefore compute a maximum matching between agents and bundles of \mathcal{A} , where the weight of an agent-bundle edge is the social impact the agent would have by receiving that bundle. Let us assume w.l.o.g. agent i is matched to bundle A_i in the maximum matching. Clearly, any other matching cannot provide a better social impact. Consider now, n distinct matchings such that, for each i and j , i gets A_j in exactly one of these matchings. Such matchings always exist as they can be obtained by shifting bundles among agents. Therefore, $n \cdot \sum_{i=1}^n s_i(A_i) \geq \sum_{i=1}^n \sum_{j=1}^n s_i(A_j) \geq \text{opt}$, where the last inequality holds as $s_i(\mathcal{G}) = \sum_{j=1}^n s_i(A_j)$.

In conclusion, the allocation corresponding to such maximum matching guarantees an n -approximation to opt . \square

An $O(n^2)$ -Approximation Subject to EF1

In this section, we tackle the problem of finding an EF1 allocation with good approximation to opt for additive valuations. The paper (Bu et al. 2023) provides an easy m -approximation to opt subject to EF1. Their algorithm, that we call SIMPLEEF1APPROX, proceeds as follows: Finds a good g^* of maximum social impact, that is, $\max_i s_i(g^*) \geq \max_{i,g} s_i(g)$. Then, it assigns g^* to the agent realizing the maximum social impact for it, say i^* . The remaining goods are allocated in an RR fashion w.r.t. an ordering where i^* is the last picking. Clearly, optimally allocating only one good cannot guarantee an approximation to opt better than m . Considering that the number of goods might be significantly larger n , we now improve the approximation as follows.

Algorithm 1: Transformation into an EF1 allocation

Input: An allocation \mathcal{A} for a fair division instance with ordered valuations and with $m = q \cdot n$

Output: An EF1 allocation \mathcal{A}'

```

1 for  $k = 1, \dots, q$  do
2    $X \leftarrow \mathcal{N}$  and  $Y \leftarrow \mathcal{G}^k$ 
3   for  $i = 1, \dots, n$  do
4     if  $\mathcal{G}^k \cap A_i \neq \emptyset$  then
5        $g \leftarrow g \in \arg \max_{g \in \mathcal{G}^k \cap A_i} s_i(g)$ 
6        $X = X \setminus \{i\}$  and  $Y = Y \setminus \{g\}$ 
7        $A'_i = A'_i \cup \{g\}$ 
8   Each good in  $Y$  is given to a distinct agent in  $X$ 
9 return  $\mathcal{A}'$ 

```

Theorem 4. *There exists a poly-time algorithm computing an EF1 allocation providing a $O(n^2)$ -approximation to opt .*

To this aim, our algorithm will take into account the real optimum to compute the EF1 solution. Let OPT be the optimum for MAXUT that will be used by our algorithm, and OPT_i be the bundle of i in OPT . Recall, opt is the optimal value $\text{SW}(\text{OPT})$. Let m_i be the number of goods in OPT_i . We sort the social value of such goods from the highest to the lowest, namely, $\text{opt}_i^1 \geq \text{opt}_i^2 \geq \dots \geq \text{opt}_i^{m_i}$, and denote the good of value opt_i^k by o_i^k . Clearly, $\text{opt} = \sum_{i \in \mathcal{N}} \sum_{k=1}^{m_i} \text{opt}_i^k$. In what follows, given any i such that $m_i < n$, we add dummy objects o_i^k valued 0 by all the agents and with $\text{opt}_i^k = 0$, for every $m_i + 1 \leq k \leq n$; removing such goods at the end will not affect EF1 nor the approximation. Moreover, whenever we talk about the social value of a good, we mean the social value it has in OPT .

Our algorithm is based on a case distinction depending on how the critical mass of opt is distributed among the bundles $\text{OPT}_1, \dots, \text{OPT}_n$. Let us denote by $\Delta_1 = \{o_i^k \mid i \in \mathcal{N}, k \in [n]\}$, the n socially best goods in OPT_i , for each i . Note that Δ_1 contains exactly n^2 goods. The social value of Δ_1 is given by $\delta_1 = \sum_{i=1}^n \sum_{k=1}^n \text{opt}_i^k$. Furthermore, we indicate by $\Delta_2 = \mathcal{G} \setminus \Delta_1$ and denote by δ_2 its social value. Hence, $\text{opt} = \delta_1 + \delta_2$. We distinguish the following two cases:

Case 1) $\delta_1 \geq \delta_2$;

Case 2) $\delta_2 < \delta_1$.

A $2n^2$ -Approximation for Case 1. This case is the easiest; it suffices to use SIMPLEEF1APPROX to get a $2n^2$ -approximation of opt . In fact, we can allocate the good g^* , of highest social impact, to the agent who owns it in OPT – this means we are correctly allocating the good of highest social impact in OPT . The social impact of such good provides an upper bound to every good in Δ_1 ; since $|\Delta_1| = n^2$, the outcome of SIMPLEEF1APPROX is trivially an n^2 -approximation of δ_1 . Being $\delta_1 \geq \frac{\delta_1 + \delta_2}{2} = \frac{\text{opt}}{2}$, this ensures an $2n^2$ -approximation.

Despite the easy approach and analysis, Case 1 constitutes the bottleneck for the approximation of opt . In what follows, we focus on Case 2 which will require a more interesting algorithm providing a $2n$ -approximation of opt .

A $2n$ -Approximation for Case 2. In this case, $\exists i \in \mathcal{N}$ such that $|\text{OPT}_i| > n$, otherwise, $\Delta_2 = \emptyset$ and $\delta_2 = 0 \leq \delta_1$. Hence, we partition the goods of each OPT_i into groups of n goods as long as it is possible. More precisely, if $m_i = k_i \cdot n + r_i$, with $r_i < n$, we create a group $B_i^{k_i}$ containing the goods $o_i^{(k_i-1)n+1}, \dots, o_i^{k_i n}$, for $k \in [k_i]$. The remaining goods, if any, are stored in $B_i^{k_i+1}$, otherwise, we let $B_i^{k_i+1} = \emptyset$. In the social impact perspective, any good in $B_i^{k_i}$ values at least as much as any good in $B_i^{k_i+1}$ when assigned to agent i . This observation is the key ingredient of the following.

Lemma 2. *In Case 2, any allocation that assigns to each agent $i \in \mathcal{N}$ one good in $B_i^{k_i}$, for each $k = 1, \dots, k_i$, guarantees an $2n$ -approximation to opt .*

Algorithm 2: $O(n)$ -approx. for MAXUT in Case 2

Input: $\text{OPT}, \mathcal{G}, \mathcal{N}, \{v_i\}_{i \in \mathcal{N}}, \{s_i\}_{i \in \mathcal{N}}$

Output: An allocation $\mathcal{A} = (A_1, \dots, A_n)$

/* Phase 1: Preprocessing */

1 Partition OPT_i in $B_i^1, \dots, B_i^{k_i+1}$ for each $i \in \mathcal{N}$

2 $\mathcal{A} \leftarrow (\emptyset, \dots, \emptyset)$

/* Phase 2: Compute a partial allocation satisfying Lemma 2 */

3 Build the envy-graph G corresponding to \mathcal{A}

4 **for** $i \in \mathcal{N}$ and $k = 1, \dots, k_i$ **do**

5 $\sigma \leftarrow \text{TOPORD}(G)$

6 **for** $j = 1, \dots, n$ **do**

7 $g \leftarrow g \in \arg \max_{g \in B_i^k} v_{\sigma(j)}(g)$

8 $A_{\sigma(j)} \leftarrow A_{\sigma(j)} \cup \{g\}$ and $B_i^k \leftarrow B_i^k \setminus \{g\}$

9 Build the envy-graph G corresponding to \mathcal{A}

10 **while** $\text{TOPORD}(G) = \text{False}$ **do**

11 $\mathcal{A} \leftarrow \text{CYCLEELIMINATION}(\mathcal{A}, G)$

12 Build the envy-graph G corresponding to \mathcal{A}

/* Phase 3: Allocate remaining goods w/o violating Lemma 2 */

13 $\mathcal{B} \leftarrow \cup_{i \in \mathcal{N}} B_i^{k_i+1}$

14 Use envy-cycle elimination to allocate the goods \mathcal{B}

15 **return** \mathcal{A}

Lemma 2 determines a sufficient condition for an allocation to be a $2n$ -approximation to opt . Algorithm 2 will provide an EF1 solution satisfying such a condition. This algorithm uses the same approach of (Bu et al. 2023) of partitioning \mathcal{G} in packages of n goods. Their purpose was to achieve EF1 both in the allocator and the agents perspectives, under the assumption that for all i $s_i = s$ for some allocator valuation s . With our approach, we use a different partition of the goods into packages and our results hold for arbitrary s_i and concerns the approximation factor of the resulting EF1 allocation.

Informally, Algorithm 2 takes an optimum allocation OPT of MAXUT and in the *preprocessing (Phase 1)* partitions the bundle OPT_i into $B_i^1, \dots, B_i^{k_i+1}$ as we previously described. In *Phase 2*, the algorithm proceeds into rounds and ensures that, at the beginning of each round, the envy-graph corresponding to the current allocation is acyclic. We can therefore compute a topological order of the agents – an agent in this ordering may envy only agents coming afterward – which will be used in the current round to assign goods. In particular, the algorithm allocates in an RR fashion the goods of a certain $B_i^{k_i}$ with respect to the just computed topological order, for $k \leq k_i$. So no two goods in B_i^k are assigned to the same agent. At the end of the round, the EF1 condition is finally satisfied (interestingly, this might not be true in the middle of a round) and all the existing cycles in the envy-graph are eliminated. Finally, in Phase 3, the remaining goods \mathcal{B} are allocated via the standard envy-cycle elimination algorithm.

The algorithm makes use of the $\text{TOPORD}(G)$ procedure

which determines whether a directed graph G admits a topological order, and if so, it outputs one.

In order to prove the correctness of the algorithm, we first observe that the output of Algorithm 2 satisfies Lemma 2. In a nutshell, since it assigns to each bundle exactly one good from B_i^k , for every $i \in \mathcal{N}$ and $k \in [k_i]$, no matter how bundles are subsequently shuffled, this property remains true. Therefore, a linear approximation immediately follows. In the remaining, we show the EF1 condition.

Proposition 2. *Phase 2 of Algorithm 2 is well-defined and produces an EF1 allocation.*

Proof. We show at the end of each round the current allocation is EF1. At the first round, the agents will have one good each and the allocation is EF1.

Assume we are in a generic round of *Phase 2* and the partial allocation \mathcal{A} at the beginning of the round is EF1. We have obtained \mathcal{A} after repeatedly running, in the previous round, the procedure CYCLEELIMINATION as long as the envy-graph is not acyclic. Hence, TOPORD will now correctly output a topological order σ of the agents.

We first notice that the **for** loop starting at line 6 assigns one good x_i to each agent i in a RR fashion with respect to the permutation σ . Hence, for each agent $\sigma(j)$, $v_{\sigma(j)}(x_{\sigma(j)}) \geq v_{\sigma(j)}(x_{\sigma(j')})$ for each $j' > j$. Moreover, being σ a topological ordering of the agents for the partial allocation \mathcal{A} , we have $v_{\sigma(j)}(A_{\sigma(j)}) \geq v_{\sigma(j)}(A_{\sigma(j')})$ for each $j' < j$. At the end of the **for** loop, agent $\sigma(j)$ owns the bundle $A'_{\sigma(j)} = A_{\sigma(j)} \cup \{x_{\sigma(j)}\}$. We next show that this bundle makes $\sigma(j)$ EF1 at the end of the **for** loop.

Consider $j' < j$; being $v_{\sigma(j)}(A'_{\sigma(j)}) \geq v_{\sigma(j)}(A_{\sigma(j)})$ and

$$v_{\sigma(j)}(A_{\sigma(j)}) \geq v_{\sigma(j)}(A_{\sigma(j')}) = v_{\sigma(j)}(A'_{\sigma(j')} \setminus \{x_{\sigma(j')}\}),$$

agent $\sigma(j)$ turns out to be EF1 w.r.t. $\sigma(j')$, as it suffices to remove the last inserted good from $A'_{\sigma(j')}$ to make $\sigma(j)$ EF.

Consider now $j' > j$; since \mathcal{A} is EF1, $\exists g \in A_{\sigma(j')}$ s.t.

$$v_{\sigma(j)}(A_{\sigma(j)}) \geq v_{\sigma(j)}(A_{\sigma(j')} \setminus \{g\});$$

on the other hand $v_{\sigma(j)}(x_{\sigma(j)}) \geq v_{\sigma(j)}(x_{\sigma(j')})$ and therefore

$$v_{\sigma(j)}(A'_{\sigma(j)}) \geq v_{\sigma(j)}(A'_{\sigma(j')} \setminus \{g\})$$

making agent $\sigma(j)$ EF1 towards $\sigma(j')$ also in this case. \square

Since at the end of Phase 2 the current allocation \mathcal{A} is EF1, so will be the output of Algorithm 2 as Phase 3 simply applies the envy-cycle elimination algorithm.

The algorithms for Cases 1 and 2 imply Theorem 4.

An n -Approx. Subject to Weaker Fairness Criteria

The main obstacle in obtaining a linear approximation is the allocation of the best good o_i^1 , for each i , while ensuring fairness. In fact, this would provide a $O(n)$ -approximation of δ_1 and, hence, to **opt** in Case 1. This can be circumvented by relaxing the envy condition on the agents' side requiring EF2. In fact, assume we remove the goods o_1^1, \dots, o_n^1 , and run Algorithm 2 in the resulting instance. This provides a partial EF1 allocation \mathcal{A} that we are able to show to be n -approx.

to δ_2 . By simply assigning to each agent i the corresponding o_i^1 we trivially get an EF2 allocation. On the other hand, assigning each o_i^1 to agent i guarantees an n -approx. of δ_1 . These facts together imply:

Theorem 5. *There exists a poly-time algorithm computing an EF2 allocation providing a n -approximation to **opt**.*

The second relaxation we consider is epistemic EF1 which implies PROP1.

In the remainder of this section, we show the following.

Theorem 6. *There exists a poly-time algorithm computing an epistemic EF1 and PROP1 allocation providing a n -approximation to **opt**.*

To this aim, we first provide a sufficient condition for epistemic EF1. Recall that $\mathcal{G}_i^h = \{g_i^{(h-1)n+1}, \dots, g_i^{hn}\}$ and q is the largest integer such that $m \geq qn$.

Lemma 3. *Let \mathcal{A} be an allocation where each agent gets at least q goods, and there exists $\{x_1, \dots, x_q\} \subseteq A_i$ s.t. $x_h \in \mathcal{G}_i^h$, for each $h \in [q]$. Then, \mathcal{A} is epistemic EF1.*

Thanks to Lemma 3 we are ready to show Theorem 6.

Proof of Theorem 6 – Sketch. Assume w.l.o.g $m = q \cdot n$, for some $q \in \mathbb{N}$; if not, we introduce an appropriate number of dummy goods evaluated 0 and having a social impact of 0.

We express our problem as the one of finding a maximum weighted perfect matching on a bipartite graph. On one side, we have m goods, on the other, we have q copies of the agents. Specifically, for each agent i , we have the copies i_1, \dots, i_q . The h -th copy of agent i is connected to all goods in \mathcal{G}_i^h , with edge weights equaling the social impact of the agent for the corresponding good. EF1. By construction, any perfect matching corresponds to an allocation satisfying conditions of Lemma 3. Therefore, a maximum weighted perfect matching in the bipartite graph guarantees epistemic EF1 and hence PROP1.

Approximation. The constructed bipartite graph is n -regular. As a direct consequence of Hall's Theorem, there exist n disjoint perfect matchings that partition the set of edges of the graph. Therefore, the sum of the social values of these n matchings (that equals $\sum_i s_i(\mathcal{G})$) is an upper bound to **opt**. On the other hand, being the computed matching of maximum weight, it is also an n -approximation to the sum of the values of the n matchings, and hence of **opt**. \square

Socially Aware Agents

So far we have shown that, being agents unconscious of their social imprint, it is not possible to guarantee a good social impact while pursuing fairness. Due to the growing attention to social issues, we might imagine scenarios where agents do take into account their and others' social impact when establishing their opinion on the outcome. To introduce some sort of social awareness, a first natural attempt might be to allow envy only towards agents having a worse social impact, that is, agents guaranteeing a better social impact are allowed to have a better bundle. Slightly more formally, we may say that an allocation \mathcal{A} is envy-free (in a social aware sense) if for each agent i one of the following conditions holds:

$$v_i(A_i) \geq v_i(A_j) \quad \text{or} \quad s_i(A_i) < s_j(A_j).$$

On the positive side, according to this definition, the envy an agent may have towards another agent is subjective (depends on her valuation), while who might be envied is objective (depends on the impact on society). Therefore, if an agent envies another this envy is eliminated because the envied agent has objectively a better impact on society. However, this first notion can be arbitrarily inefficient.

Example 1. Assume there are two agents and one item. Both agents value 1 the item; agent 1 has a social impact 1 while agent 2 has social impact ε for receiving the item. In this scenario, $\text{opt} = 1$. Consider the allocation \mathcal{A} where the item is given to agent 2. This provides a social welfare of ε . Moreover, agent 1 does envy agent 2. However, the social impact of 1 is higher than the one of agent 2. If agents 1 would decide to remove the envy towards agent 2 on the sole basis of their current social impact, the allocation \mathcal{A} would be considered fair while, for the sake of the social welfare, it is reasonable that agent 1 envies agent 2.

The main issue with this first attempt at defining a socially aware notion is that there is no comparison in what potentially could guarantee the possibly envious agent. Roughly speaking: It is foolish to believe we cannot envy someone just because she/he has a better impact than us if, in reality, we would make a better contribution if we were in her/his place. This is the idea behind the following definition.

Definition 1 (Socially aware EF). We say an allocation \mathcal{A} is *socially aware envy-free* (sEF) if for each pair of agents i, j one of the following conditions holds:

$$v_i(A_i) \geq v_i(A_j) \quad \text{or} \quad s_i(A_j) < s_j(A_j).$$

In turn, a socially aware agent i envies j (i sa-envies j , in short) if $v_i(A_i) < v_i(A_j)$ and $s_i(A_j) \geq s_j(A_j)$.

Notice that, if $s_i(g) = 0$ for each $g \in \mathcal{G}$ and $i \in \mathcal{N}$ – which means, agents have no social impact – sEF \equiv EF; this implies sEF may not exist as well. We, therefore, relax this notion in the “up to one good” sense.

Definition 2 (sEF up to one good). An allocation \mathcal{A} is sEF *up to one good* (sEF1) if for each pair of agents i, j , such that $A_j \neq \emptyset$, one of the following conditions holds:

- there exists $g \in A_j$ such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$, or
- $s_i(A_j) < s_j(A_j)$.

We notice we could have defined the “up to one good” relaxation saying i envies j if, for each $g \in A_j$, $s_i(A_j \setminus \{g\}) \geq s_j(A_j)$ and $v_i(A_i) < v_i(A_j \setminus \{g\})$; however, this is weaker than the sEF1 we just defined.

Surprisingly, we are able to show the following result.

Theorem 7. *There exists a poly-time algorithm computing an sEF1 allocation which is optimal for MAXUT.*

Sketch. The theorem is a consequence of the following observation: Given an optimal allocation OPT, if i sa-envies j , then $s_i(g) = s_j(g)$, for each $g \in \text{OPT}_j$.

Our algorithm sequentially assigns the goods in \mathcal{G} optimally while maintaining EF1. Assume at a certain step the current partial allocation \mathcal{A} is EF1 and g has to be assigned to an agent $i^* \in \arg \max_i s_i(g)$, to guarantee optimality.

We need to make sure that the new allocation is EF1. If not, there exists j who would not be sEF1 toward i^* , and hence not sEF, in the \mathcal{A}' obtained from \mathcal{A} by allocating g to i^* . \mathcal{A}' is optimal, for the currently allocated goods, and our previous observation imply 1) $s_j(g) = s_{i^*}(g)$ and 2) $s_j(A_{i^*}) = s_{i^*}(A_{i^*})$, as the social impact of j is the same of i for all goods in $A'_{i^*} = A_{i^*} \cup \{g\}$. 1) means that the only agents that might envy i^* after allocating g are the ones in $\arg \max_i s_i(g)$, so, we should choose the i^* who is not envied by any other agent in $\arg \max_i s_i(g)$. However, it is possible that such an agent does not exist – because there are no sources among $\arg \max_i s_i(g)$ in the sa-envy-graph. 2) establishes that if there exists a cycle of sa-envy in $\arg \max_i s_i(g)$ applying CYCLEELIMINATION does not affect the social impact. Hence, we can apply CYCLEELIMINATION until we find a $i^* \in \arg \max_i s_i(g)$ who is not envied. This shows there always exists a way to sequentially assign goods optimally while maintaining EF1. \square

Discussion

We have investigated the fair division model when the underlying allocation problem has an impact on society.

Our work also paves the way for a large number of future challenges. One might study the scenario where goods are chores for the agents; in such case, not all our results naturally extend. In fact, the envy-cycle elimination no longer guarantees EF1, and a different approach is needed (Bhaskar, Sricharan, and Vaish 2021). However, this cannot be immediately translated into an approximation algorithm for MAXUT. Other appealing directions include the study of social welfare functions other than the utilitarian.

We have also introduced social awareness. One may extend this definition to encode different levels of awareness by introducing a parameter α establishing how altruistic agents are. If $\alpha = 0$ we would have EF1, which means, agents completely ignore their impact, while for $\alpha = 1$, the notion would coincide with sEF1. We have seen that an approximation to opt better than linear is not possible when requiring EF1, while an outcome that is optimal and sEF1 always exists. We wonder if there exists an explicit connection between the parameter α and the attainable approximation.

Acknowledgments

This work was partially supported by: the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI; the PNRR MIUR project VITALITY (ECS00000041), Spoke 2 - Advanced Space Technologies and Research Alliance (ASTRA); the European Union - Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3, CUP J33C22002880001, partnership on “Telecommunications of the Future”(PE00000001 - program “RESTART”), project MoVeOver/SCHEDULE (“Smart interseCtions witH connEctED and aUtonomous vehicLEs”, CUP J33C22002880001); and Gruppo Nazionale Calcolo Scientifico-Istituto Nazionale di Alta Matematica (GNCS-INdAM).

References

- Amanatidis, G.; Aziz, H.; Birmapas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322: 103965.
- Aziz, H.; Bouveret, S.; Caragiannis, I.; Giagkousi, I.; and Lang, J. 2018. Knowledge, Fairness, and Social Constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 4638–4645.
- Aziz, H.; Freeman, R.; Shah, N.; and Vaish, R. 2023. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*.
- Aziz, H.; Ganguly, A.; and Micha, E. 2023. Best of both worlds fairness under entitlements. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 941–948.
- Aziz, H.; Huang, X.; Mattei, N.; and Segal-Halevi, E. 2019. The Constrained Round Robin Algorithm for Fair and Efficient Allocation. *CoRR*, abs/1908.00161.
- Aziz, H.; Moulin, H.; and Sandomirskiy, F. 2020. A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5): 573–578.
- Barman, S.; Bhaskar, U.; and Shah, N. 2020. Optimal bounds on the price of fairness for indivisible goods. In *International Conference on Web and Internet Economics*, 356–369. Springer.
- Barman, S.; Khan, A.; Shyam, S.; and Sreenivas, K. V. N. 2023. Guaranteeing Envy-Freeness under Generalized Assignment Constraints. In *Proceedings of the 24th ACM Conference on Economics and Computation, EC 2023*, 242–269. ACM.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, 7–13. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM.
- Bei, X.; Lu, X.; Manurangsi, P.; and Suksompong, W. 2021. The price of fairness for indivisible goods. *Theory of Computing Systems*, 65: 1069–1093.
- Bhaskar, U.; Sricharan, A. R.; and Vaish, R. 2021. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021, August 16-18, 2021, University of Washington, Seattle, Washington, USA (Virtual Conference)*, volume 207 of *LIPICs*, 1:1–1:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Bilò, V.; Caragiannis, I.; Flammini, M.; Igarashi, A.; Monaco, G.; Peters, D.; Vinci, C.; and Zwicker, W. S. 2022. Almost envy-free allocations with connected bundles. *Games Econ. Behav.*, 131: 197–221.
- Bredereck, R.; Kaczmarczyk, A.; and Niedermeier, R. 2022. Envy-free allocations respecting social networks. *Artif. Intell.*, 305: 103664.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2022. On the complexity of maximizing social welfare within fair allocations of indivisible goods. *arXiv preprint arXiv:2205.14296*.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2023. Fair Division with Allocator’s Preference. In *Web and Internet Economics - 19th International Conference, WINE 2023*, volume 14413 of *Lecture Notes in Computer Science*, 77–94. Springer.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.
- Caragiannis, I.; Garg, J.; Rathi, N.; Sharma, E.; and Varricchio, G. 2023. New Fairness Concepts for Allocating Indivisible Items. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023*, 2554–2562. ijcai.org.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3): 1–32.
- Chakraborty, M.; Igarashi, A.; Suksompong, W.; and Zick, Y. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation (TEAC)*, 9(3): 1–39.
- Foley, D. K. 1966. *Resource allocation and the public sector*. Yale University.
- Hofer, M.; Schmalhofer, M.; and Varricchio, G. 2024. Best of both worlds: Agents with entitlements. *Journal of Artificial Intelligence Research*, 80: 559–591.
- Li, Z.; Liu, S.; Lu, X.; Tao, B.; and Tao, Y. 2024. A Complete Landscape for the Price of Envy-Freeness. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024*, 1183–1191.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131. ACM.
- Nicosia, G.; Pacifici, A.; and Pferschy, U. 2017. Price of fairness for allocating a bounded resource. *European Journal of Operational Research*, 257(3): 933–943.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.
- Suksompong, W. 2021. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2): 46–61.
- Suksompong, W. 2025. Weighted Fair Division of Indivisible Items: A Review. *Information Processing Letters*, 187: 106519.