The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23)

Fairness Concepts for Indivisible Items with Externalities

Haris Aziz[1], Warut Suksompong[2], Zhaohong Sun[3], Toby Walsh[1]

[1]School of Computer Science and Engineering, University of New South Wales, Australia
[2]School of Computing, National University of Singapore, Singapore
[3]AI Lab, CyberAgent, Japan

Abstract

We study a fair allocation problem of indivisible items under additive externalities in which each agent also receives utility from items that are assigned to other agents. This allows us to capture scenarios in which agents benefit from or compete against one another. We extend the well-studied properties of envy-freeness up to one item (EF1) and envy-freeness up to any item (EFX) to this setting, and we propose a new fairness concept called general fair share (GFS), which applies to a more general public decision making model. We undertake a detailed study and present algorithms for finding fair allocations.

1 Introduction

Fair allocation of indivisible items is an active field of research within computer science and economics (Brams and Taylor 1996; Bouveret, Chevaleyre, and Maudet 2016; Thomson 2016). The general problem is to allocate the items among the agents so as to satisfy certain fairness criteria. For example, one important fairness concept is envy-freeness, which stipulates that no agent wants to swap her bundle with another agent's bundle. The field has witnessed several new solution concepts, algorithms, and applications.

In most of the work on fair allocation, agents are assumed to derive value *only* from the set of items allocated to them. In this paper, we consider a significantly more general model in which an agent's value for an allocation may depend both on the agent's own bundle as well as on the bundles of items given to other agents. The latter aspect is referred to in the economics literature as *externalities*. Whereas the theory of fair allocation has progressed tremendously, the topic is relatively less developed when externalities are involved in the valuations of the agents.

Externalities in agent preferences are present in many real-world scenarios. When resources are allocated among agents, an agent may derive positive value from resources given to the agent's friend or family member because the agent has access rights to the resource. Positive externalities can also capture settings where agents are divided into groups and each agent receives the same utility whenever some agent in the group is allocated an item, and no utility when an agent outside the group is allocated the item. Likewise, negative externalities can arise in various resource allocation settings. For example, when dividing assets among conflicting groups, the allocation of a critical asset to another group may hamper one group's functionality. Yet another example of negative externalities is the case of sport drafts, where a team may incur negative value if a valuable player is given to a competing team.

Although externalities have been considered in some prior work on resource allocation problems, the focus was on either allocation of *divisible* resources (Brânzei, Procaccia, and Zhang 2013; Li, Zhang, and Zhang 2015) or concepts based on maximin fair share (Seddighin, Saleh, and Ghodsi 2021). In this paper, we revisit important fairness concepts such as envy-freeness and consider suitable relaxations in the context of indivisible items under externalities.

Since there may not exist an envy-free allocation in general, much of the recent research has focused on relaxations of envy-freeness by removing one or more items from consideration, e.g., *envy-freeness up to one item (EF1)* (Lipton et al. 2004; Budish 2011). Caragiannis et al. (2019) proposed a stronger concept than EF1 called *envy-freeness up to any item (EFX)*. The intuition is that if agent i envies agent j's assignment, then the envy should be eliminated when any item is removed from j's assignment. Aziz et al. (2022a) generalized EFX to the setting of goods and chores (i.e., negative values) when there are no externalities. The difference is that the envy from i towards j can be eliminated by removing agent i's least preferred good from j's bundle, and also by removing agent i's favorite chore from i's own bundle.

For allocation problems under externalities, envy-freeness needs to be carefully extended. When we consider fair allocation of goods without externalities, if agent i envies agent j's assignment, then removing any item from agent j's bundle decreases the envy. However, this no longer holds when externalities exist. For instance, assume that agent i receives value 5 when item a is assigned to agent i and receives value 10 when a is assigned to agent j. In that case, it is unclear that removing item a from j's bundle decreases i's "envy" towards j, since i actually derives more utility when the item is allocated to j than when it is allocated to i herself. This issue becomes more complicated when both positive and negative externalities are allowed.

Another widely studied fairness concept under additive valuations is *proportionality*, which requires each agent to receive at least $1/n$ of the value that she has for the set of all items, where n denotes the number of agents. Conitzer, Freeman, and Shah (2017) proposed a variant of proportionality for a more general *public decision making* problem than allocation of indivisible items under externalities. They showed that their concept is guaranteed to be feasible under positive valuations. However, this guarantee ceases to hold when negative valuations are also allowed.

We summarize our contributions as follows. First, we define the concepts of EF1 and EFX under externalities that still coincide with previous definitions for goods and chores when externalities do not exist. Note that our new concepts work for both positive and negative externalities.

Second, we show how to compute an EFX allocation between two agents in time $O(m \log m)$, where m denotes the number of items, and how to compute an EF1 allocation between two agents in time linear in m.

Third, we show that the set of EFX allocations among three agents could be empty. Under binary values and a "no-chore" assumption, we show that an EF1 allocation always exists among three agents by proposing a new algorithm that computes such an allocation in polynomial time.

Fourth, we propose a new fairness concept called *general fair share (GFS)* based on proportionality. We present a polynomial-time algorithm that computes an allocation satisfying general fair share up to one item (GFS1) for the more general public decision making model where both positive and negative valuations are allowed.

Finally, we present a taxonomy of fairness definitions including both existing and newly proposed concepts.

2 Related Work

Fair allocation of indivisible items is an active topic of research in computer science and economics (Brams and Taylor 1996; Bouveret, Chevaleyre, and Maudet 2016; Thomson 2016). For some recent overviews, we refer to the surveys of Amanatidis et al. (2022) and Aziz et al. (2022b).

For allocation problems under externalities, fairness concepts need to be carefully revisited and extended. Velez (2016) proposed a natural adaptation of envy-freeness which requires that no agent prefers the allocation obtained by swapping her bundle with another agent. Brânzei, Procaccia, and Zhang (2013) considered both the envy-freeness concept of Velez (2016) and proportionality in the context of cake-cutting. Li, Zhang, and Zhang (2015) studied truthful mechanisms in the setting of cake-cutting under externalities. Since cake-cutting involves the allocation of a divisible resource, one can obtain existence results without relaxations even when externalities are present.

In our paper, we focus on allocation of indivisible items. Seddighin, Saleh, and Ghodsi (2021) presented an algorithm for computing an allocation that satisfies a relaxation of a concept called maximin share fairness, which can in turn be viewed as a relaxation of proportionality. Note that both Seddighin, Saleh, and Ghodsi (2021) and Brânzei, Procaccia, and Zhang (2013) restricted their attention to settings with positive externalities, whereas we allow both positive and negative externalities. Li, Zhang, and Zhang (2015) made the restrictive assumption that agents derive externalities that are percentages of other agents' values. Mishra, Padala, and Gujar (2022) studied a special form of externalities in which an agent receives the same externality from an item regardless of which other agent receives the item.

A related line of work concerns house or residential allocation with externalities, where an agent's value for an allocation is influenced by other agents assigned to her neighborhood (Chauhan, Lenzner, and Molitor 2018; Massand and Simon 2019; Elkind et al. 2020; Agarwal et al. 2021; Bullinger, Suksompong, and Voudouris 2021; Gross-Humbert et al. 2021).

3 Model

We consider a setting where a set of indivisible items $A = \{a_1, \ldots, a_m\}$ are to be allocated among a set of agents $N = \{1, \ldots, n\}$ under additive externalities.

An allocation is denoted by $\pi = (\pi_1, \ldots, \pi_n)$ where each $\pi_i \subseteq A$ is the bundle assigned to agent i such that for any distinct $i, j \in N$, we have $\pi_i \cap \pi_j = \emptyset$. If $\bigcup_{i \in N} \pi_i = A$, then we call π a *complete* allocation of A. Unless specified otherwise, we only consider complete allocations. Let Π denote the set of all allocations. For any item $a \in A$, let $\pi(a)$ denote the agent who receives item a in allocation π.

Every agent $i \in N$ is associated with a valuation function $V_i : \Pi \to \mathbb{R}$, which assigns a real value to every allocation $\pi \in \Pi$. We assume that agents have additive valuations and externalities. Under the additive preference domain, we have $V_i(\pi) = \sum_{a \in A} V_i(\pi(a), a)$, where we abuse notation and let $V_i(j, a)$ represent the value that agent i receives when item a is assigned to agent j. Note that in problems without externalities, an agent receives the same value from an allocation as long as the agent receives the same bundle.

4 EF1 and EFX under Externalities

In this section, we consider how to generalize the definitions of EF1 and EFX to the setting of externalities. Note that we need to carefully design both definitions to ensure that they coincide with the previous definitions without externalities.

Velez (2016) proposed a natural adaptation of envy-freeness which requires that no agent prefers the allocation obtained by swapping her bundle with another agent's bundle. Since this notion has become the standard of envy-freeness in the setting of externalities, we simply refer to it as envy-freeness. In this work, we follow this idea of *swapping bundles* to define EF1 and EFX. Let $\pi^{i \leftrightarrow j}$ represent a new allocation in which only agents i and j swap their bundles in π while other agents' bundles remain the same.

Definition 4.1 (Envy-Freeness (Velez 2016)). *An allocation π is envy-free (EF) if there do not exist agents $i, j \in N$ such that $V_i(\pi^{i \leftrightarrow j}) > V_i(\pi)$.*

Recall that envy-freeness cannot be guaranteed in the indivisible domain even if there are two agents and one item, and the agents have no externalities. In view of this challenge, a natural recourse is to explore "up to one item relaxations" of fairness concepts. The intuition is that when an agent is envious, she would like to swap her bundle with another agent. The "up to k" relaxation ensures that such a swap is not desirable if at most k items are removed from consideration. We next formalize an "up to k items relaxation" of EF under externalities.

Definition 4.2 (Envy-Freeness up to k Items)**.** *An allocation π is envy-free up to k items (EFk) if for every pair of agents $i, j \in N$, there exists a set of items $C \subseteq A$ and an allocation λ such that the following conditions hold:*

1. *$|C| \leq k$;*
2. *$\lambda_\ell = \pi_\ell \setminus C$ for all $\ell \in N$;*
3. *$V_i(\lambda) \geq V_i(\lambda^{i\leftrightarrow j})$.*

In words, Definition 4.2 states that an allocation π is EFk if for each pair of agents i and j, there exists a set of items C of size at most k such that for the new allocation λ obtained by removing items in C from each agent ℓ's bundle π_ℓ in π, agent i would not like to swap her bundle λ_i with agent j's bundle λ_j. Note that if $k = 1$ and there are no externalities, then Definition 4.2 coincides with the EF1 concept as formalized by Budish (2011) for goods and by Aziz et al. (2022a) for goods and chores.

We next generalize EFX to the case of externalities. A first attempt is to define the generalization so that if agent i envies agent j, then removing any item from either of their bundles should eliminate the envy. However, this fails to capture the original idea of Caragiannis et al. (2019) when externalities exist, as shown in Example 4.3.

Example 4.3. *Consider two agents $N = \{1, 2\}$ and three items $A = \{a, b, c\}$. The values of items and externalities are described in Table 1. For allocation $\pi = \{(1, ab), (2, c)\}$ in which agent 1 receives items a, b and agent 2 receives item c, agent 2 has envy towards agent 1:*

$$V_2(\pi) - V_2(\pi^{1\leftrightarrow 2}) = (1 + 2 + 2) - (4 + 1 + 3) = -3.$$

If we remove item b from agent 1's bundle, then agent 2 envies agent 1 even more. That is, for allocation $\tilde{\pi} = \{(1, a), (2, c)\}$,

$$V_2(\tilde{\pi}) - V_2(\tilde{\pi}^{1\leftrightarrow 2}) = (1 + 2) - (4 + 3) = -4.$$

	a	b	c
1	$3, 1$	$1, 2$	$2, 1$
2	$1, 4$	$2, 1$	$3, 2$

Table 1: For each row $i \in \{1, 2\}$ and each entry (x, y) in row i, x and y denote the value that agent i receives when the corresponding item is assigned to agent 1 and 2, respectively.

When we consider only goods (i.e., indivisible items with positive values) without externalities, if some agent i envies another agent j, then removing any item from agent j's bundle decreases the envy. However, this is no longer true when externalities exist. As shown in Example 4.3, removing item b from agent 1's bundle does not reduce the envy from agent 2 towards agent 1. Instead, it increases this envy.

We thus propose a more suitable generalization of EFX in Definition 4.4. Intuitively, if agent i envies agent j, then for any item a such that removing a from the bundle of agent i or j reduces the envy, i should no longer envy j after removing a. This idea coincides with the definition of EFX for goods and chores when there are no externalities (Aziz et al. 2022a). Recall that in the definition by Aziz et al. (2022a), agent i's envy towards agent j can be eliminated by removing i's least preferred good from j's bundle as well as by removing i's favorite chore (i.e., one yielding the least disutility) from i's own bundle.

Definition 4.4 (Envy-Freeness up to Any Item)**.** *An allocation π is envy-free up to any item (EFX) if for all agents $i, j \in N$, if i envies j, then for any item $a \in A$ and allocation λ with the properties*

1. *$\lambda_\ell = \pi_\ell \setminus \{a\}$ for all $\ell \in N$ and*
2. *$V_i(\lambda) - V_i(\lambda^{i\leftrightarrow j}) > V_i(\pi) - V_i(\pi^{i\leftrightarrow j})$,*

the following holds:

$$V_i(\lambda) \geq V_i(\lambda^{i\leftrightarrow j}).$$

We next explain Definition 4.4 in detail. $V_i(\pi) - V_i(\pi^{i\leftrightarrow j})$ represents the envy from agent i towards agent j with respect to allocation π. Because agent i envies agent j, we have $V_i(\pi) < V_i(\pi^{i\leftrightarrow j})$, which implies $V_i(\pi) - V_i(\pi^{i\leftrightarrow j}) < 0$. Allocation λ is obtained by removing some item a from the bundle π_ℓ containing a. (Note that if $a \notin \pi_i \cup \pi_j$, then the envy of i towards j does not change upon removing a, so we may assume that $a \in \pi_i \cup \pi_j$.) Similarly, $V_i(\lambda) - V_i(\lambda^{i\leftrightarrow j})$ represents the envy from i towards j with respect to the new allocation λ. Thus $V_i(\lambda) - V_i(\lambda^{i\leftrightarrow j}) > V_i(\pi) - V_i(\pi^{i\leftrightarrow j})$ means that removing item a reduces the envy of i towards j. Finally, $V_i(\lambda) \geq V_i(\lambda^{i\leftrightarrow j})$ requires that i does not envy j with respect to the new allocation λ.

Note that our new definition of EFX in Definition 4.4 still implies EF1 in Definition 4.2. To see this, consider an EFX allocation π. For any pair of agents i and j, if i envies j in allocation π, then because valuations and externalities are additive, there must exist an item $a \in A$ such that removing a from either bundle π_i or π_j helps decrease the envy of i towards j. Since π is EFX, removing a must eliminate i's envy towards j. Thus the allocation π is also EF1. We next show an example of an EF1 allocation that is not EFX.

Example 4.5. *Consider the instance in Example 4.3. Allocation $\pi' = \{(1, bc), (2, a)\}$ is EF1 but not EFX for agent 1. To see this, first note that agent 1 envies agent 2:*

$$V_1(\pi') - V_1(\pi'^{1\leftrightarrow 2}) = (1 + 2 + 1) - (3 + 2 + 1) = -2.$$

If we remove item a from agent 2's bundle π'_2, then agent 1 does not envy agent 2. That is, for allocation $\hat{\pi} = \{(1, bc), (2, \emptyset)\}$,

$$V_1(\hat{\pi}) - V_1(\hat{\pi}^{1\leftrightarrow 2}) = (1 + 2) - (2 + 1) = 0.$$

On the other hand, if we remove item b from agent 1's bundle π'_1, which helps decrease agent 1's envy, then agent 1 still envies agent 2. That is, for allocation $\overline{\pi} = \{(1, c), (2, a)\}$,

$$V_1(\overline{\pi}) - V_1(\overline{\pi}^{1\leftrightarrow 2}) = (1 + 2) - (3 + 1) = -1.$$

5 Two Agents

In this section, we prove the existence of EFX allocations between two agents by mapping onto a simplified problem where agents have "symmetric valuations" and an EFX allocation can be constructed in polynomial time. In Lemma 5.1, we show how to construct an EFX allocation between two agents when valuations are symmetric. Based on this result, we then prove the existence of EFX allocations between two agents in Theorem 5.2. We further show that an EF1 allocation between two agents can be computed in linear time.

We say that agents' valuations are *symmetric* if for each pair $i, j \in N$ and each item $a \in A$, the following holds: $V_i(i, a) = V_j(j, a)$ and $V_i(j, a) = V_j(i, a)$.

Lemma 5.1. *An EFX allocation always exists for two agents with symmetric valuations and can be computed in time $O(m \log m)$.*

Proof. Consider two agents with symmetric valuations. For an item a, let $\Delta_{12}(a) = V_1(1, a) - V_1(2, a)$ be the difference between the value $V_1(1, a)$ that agent 1 receives when a is assigned to agent 1 and the value $V_1(2, a)$ that she receives when a is assigned to agent 2. Since valuations are symmetric, $V_1(1, a) - V_1(2, a) = V_2(2, a) - V_2(1, a)$.

Create an allocation $\tilde{\pi}$ as follows. We iteratively allocate each item in decreasing order of $|\Delta_{12}(\cdot)|$. At each step, there are two possible bundles for the item, leading to two different allocations. We choose one that the agent with the smaller current total value weakly prefers from these two allocations. That is, if $\Delta_{12}(a) \geq 0$, then the item is assigned to the agent with the smaller current total value; otherwise the item is assigned to the other agent. Break ties arbitrarily.

We next prove that allocation $\tilde{\pi}$ is EFX for both agents. Suppose we allocate all items in the order $a_1, a_2, \ldots, a_m$. For the base case, assigning a_1 to either agent is EFX. For the induction, assume that a partial allocation of items a_1, $\ldots$, a_k is EFX. Since the agents have symmetric valuations, at most one agent can be envious. Without loss of generality, assume agent 1 has at most the same value as agent 2 and the algorithm allocates a_{k+1} according to agent 1's preference. Then agent 1's envy towards agent 2 weakly decreases and the allocation is still EFX for agent 1. If agent 2 becomes envious, then removing item a_{k+1} will eliminate the envy. For any item a_j allocated to agent 1 with $\Delta_{12}(a_j) > 0$, we have $\Delta_{12}(a_j) \geq \Delta_{12}(a_{k+1})$ and removing any such item will eliminate the envy from agent 2 as well; a similar argument holds for any item a_j allocated to agent 2 with $\Delta_{12}(a_j) < 0$. Thus the allocation remains EFX for both agents.

We can sort all items in decreasing order of $|\Delta_{12}(\cdot)|$ in time $O(m \log m)$ and thus we can compute an EFX allocation between two agents with symmetric valuations in polynomial time. This completes the proof of Lemma 5.1. □

Based on Lemma 5.1, we prove the existence of EFX allocation between two agents in Theorem 5.2.

Theorem 5.2. *There always exists an EFX allocation between two agents which can be computed in time $O(m \log m)$.*

Proof. First create a dummy agent $1'$ of 1. Both agent 1 and agent $1'$ treat each other as agent 2 and they have a symmetric valuation function such that for any item $a \in A$, we have $V_1(1, a) = V_{1'}(1', a)$ and $V_1(1', a) = V_{1'}(1, a) = V_1(2, a)$. That is, if item a is assigned to $1'$, then agent $1'$ receives the value $V_1(1, a)$ and agent 1 receives the value $V_1(2, a)$ as if item a is assigned to 2 from the perspective of agent 1.

Compute an EFX allocation $\tilde{\pi}$ between agents 1 and $1'$ via the algorithm in the proof of Lemma 5.1. Allocation $\tilde{\pi}$ divides all items A into two bundles; let agent 2 first choose the bundle she prefers and leave the remaining bundle to agent 1. Since agent 2 chooses first, she does not envy agent 1. We showed that $\tilde{\pi}$ is EFX between 1 and $1'$ in Lemma 5.1, so it is EFX no matter which bundle agent 1 receives. This completes the proof of Theorem 5.2. □

We remark that constructing an EF1 allocation is easier and can be done in linear time, because we do not need to sort all items based on $|\Delta(\cdot)|$. The detailed proof is provided in the full version of our paper (Aziz et al. 2022c).

Corollary 5.3. *There always exists an EF1 allocation between two agents which can be computed in time $O(m)$.*

6 Three Agents

In this section, we consider EF1 and EFX allocations among three agents. Note that several real-world problems involve a limited number of agents (e.g., divorce settlement and inheritance division). We first show that, in contrast to the positive results of EFX allocations between two agents with externalities (Section 5) and among three agents *without* externalities (Chaudhury, Garg, and Mehlhorn 2020), there may not exist an EFX allocation for three agents with externalities. We then prove that an EF1 allocation always exists among three agents under binary values and a "no-chore" assumption by proposing a polynomial-time algorithm for this case.

Theorem 6.1. *The set of EFX allocations could be empty when there are three agents.*

Proof. We prove Theorem 6.1 through the following counterexample. Consider three agents $N = \{1, 2, 3\}$ and seven items $A = \{a_1, a_2, a_3, a_4, a_5, a_6, g\}$. The values of items and externalities are described in Table 2.

	a_k	g
1	$21, 16, 16$	$17, 16, 16$
2	$16, 21, 16$	$16, 24, 0$
3	$16, 16, 21$	$16, 0, 24$

Table 2: Values for items and externalities. For each row $i \in \{1, 2, 3\}$ and each entry (x, y, z) in row i, x, y, and z denote the value that agent i receives when the corresponding item is assigned to agent 1, 2, and 3, respectively.

First consider the case where item g is assigned to agent 1.

- If at most one item from $\{a_1, \ldots, a_6\}$ is assigned to agent 1, then either agent 2 or 3 receives at least three items from this set. Suppose agent 1 receives $\{a_4, g\}$ (or just $\{g\}$) and agent 2 receives $\{a_1, a_2, a_3\}$ (or more). Then

agent 1 has envy towards agent 2, and the envy remains when one item a_1 is removed from agent 2's bundle.
- If at least two items from $\{a_1, \ldots, a_6\}$ are assigned to agent 1, then either agent 2 or 3 receives at most two items from this set. Suppose agent 1 receives $\{a_1, a_2, g\}$ (or more) and agent 2 receives $\{a_3, a_4\}$ (or less). Then agent 2 has envy towards agent 1, and the envy remains when one item a_1 is removed from agent 1's bundle.

The case where item g is assigned to agent 2 or 3 can be handled by a similar type of analysis; we provide the details in the full version of our paper (Aziz et al. 2022c). □

We next consider EF1 allocations and restrict our attention to the case where the three agents have binary valuations, i.e., each $V_i(j, a)$ is either 0 or 1. Note that binary valuations without externalities have previously been considered in fair division both for the desirable normative properties that they allow and for their ease of elicitation (Aleksandrov et al. 2015; Darmann and Schauer 2015; Bouveret and Lemaître 2016; Barman, Krishnamurthy, and Vaish 2018; Freeman et al. 2019; Halpern et al. 2020; Suksompong and Teh 2022).

Even under the setting of binary valuations, it still appears to be challenging to check if a given instance admits an EF1 allocation and to determine whether there always exists an EF1 allocation for any instance. To make the question tractable, we additionally impose one more assumption in Definition 6.2: for each item, all agents prefer to own the item rather than to have it assigned to others. Note that the example in Theorem 6.1 satisfies this assumption.

Definition 6.2 (No-Chore Assumption). *For any item a and any pair of agents i, j, we have $V_i(i, a) \geq V_i(j, a)$.*

We prove that an EF1 allocation always exists in this setting and can be computed in polynomial time. Our method may be useful for further exploration of more general settings, e.g., whether there exists an EF1 allocation among multiple agents under more general preference domains.

Theorem 6.3. *For three agents under No-Chore Assumption and binary valuations, there always exists an EF1 allocation which can be computed in polynomial time.*

Due to space limitation, we give a high-level description of our algorithm here and present a detailed proof of Theorem 6.3 in the full version of our paper (Aziz et al. 2022c). The key idea is that given any instance, we iteratively apply some reduction rules that assign one item, one pair, or three items to some agents in an EF or EF1 manner. We show that any instance can be reduced to a certain number of cases where each case consists of a small number of items (no more than 12). We then wrote a program to verify that for each case there always exists an EF1 allocation by exhaustive search.

For the sake of illustration, we next describe two simple reduction rules. Given an item $a \in A$, we can create a 3-by-3 matrix to represent each agent i's valuation function $V_i(\cdot, a)$, where the ith row corresponds to the values that agent i receives when item a is assigned to each agent.

- Assign an item a to some agent i if it does not generate envy from any agent towards i. For instance, if we have an item with the following matrix, then we can assign it to agent 1 without generating envy from any other agent.

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

- Suppose that item a will not generate envy from agent i if it is assigned to any other agent. Then we can leave this item aside until we cannot apply any other reduction rules and then consider assigning item a to the other two agents in an EF1 manner. For instance, if we have an item with the following matrix, then assigning it to either agent 2 or agent 3 does not generate envy from agent 1.

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Note that we have only taken an initial step towards a complete understanding of EF1 allocations under externalities. We conjecture that an EF1 allocation always exists for three agents under binary valuations even without the No-Chore Assumption. For larger numbers of agents n, we may need to relax EF1 to EFk where k is a function of n.

7 GFS and Public Decision Making

In this section, we propose a new fairness concept based on proportionality that we call *general fair share (GFS)*. This concept works even for the "public decision making" setting (Conitzer, Freeman, and Shah 2017), which generalizes fair division of indivisible items under externalities. We show that there always exists an allocation satisfying general fair share up to one item (GFS1) even when the valuations can be positive or negative, and such an allocation can be computed in polynomial time via a variant of round robin. We also discuss how GFS1 is superior to an existing proportionality concept in public decision making.

7.1 Public Decision Making

An instance I^P of public decision making consists of a set of agents N and a set of issues A. Each issue $a \in A$ is associated with a set of choices a^T, exactly one of which needs to be selected. For each choice $a^t \in a^T$ of issue a, each agent i derives a value $V_i(a^t)$, where we reuse the notation V_i in a slightly different way than in fair allocation under externalities. An allocation π of instance I^P is a set of choices for all issues; let $\pi(a)$ denote the choice made for issue a. The value that agent i receives from allocation π is $V_i(\pi) = \sum_{a \in A} V_i(\pi(a))$.

A fair allocation problem with additive externalities can be reduced to an equivalent public decision making problem as follows: Each item a is viewed as an issue and associated with exactly n choices, where each choice corresponds to an agent to whom the item could be given. For public decision making, the number of choices is flexible, whereas for fair allocation with externalities, the number of choices is equal to the number of agents n.

Conitzer, Freeman, and Shah (2017) proposed the following concept for the public decision making problem, which requires that each agent i should receive at least $1/n$ of the maximum value she can get from all of the issues.[1] For each issue a, let $V_i^{max}(a) = \max_{a^t \in a^T} V_i(a^t)$

Definition 7.1 (PROP-Max)**.** *Given an allocation π, the Proportional-Max share of agent i (PROP-Max$_i$) is defined as*

$$\textit{PROP-Max}_i = \frac{1}{n} \sum_{a \in A} V_i^{max}(a).$$

An allocation π satisfies Proportionality-Max (PROP-Max) if $V_i(\pi) \geq$ PROP-Max$_i$ holds for all $i \in N$.

Conitzer et al. also introduced an "up to one" relaxation of PROP-Max.

Definition 7.2 (PROP-Max up to One Issue)**.** *An allocation π satisfies Proportionality-Max up to one issue (PROP-Max-1) if for all $i \in N$, there exists $a \in A$ such that*

$$V_i(\pi) - V_i(\pi(a)) + V_i^{max}(a) \geq \textit{PROP-Max}_i.$$

In other words, an allocation π satisfies PROP-Max-1 if for each agent i, there exists an issue a such that changing the assignment of a from $\pi(a)$ to agent i's best assignment yielding $V_i^{max}(a)$ ensures that the value that agent i receives is at least her PROP-Max$_i$.

Conitzer et al. showed that when all valuations are positive, there always exists a PROP-Max-1 allocation. However, our next proposition shows that a PROP-Max-1 may not exist if negative valuations are allowed.

Proposition 7.3. *There may not exist a PROP-Max-1 allocation when negative valuations are allowed, even if there are only two agents.*

Proof. We show this negative result for the more restricted setting of fair allocation with externalities.

Consider $N = \{1, 2\}$ and $A = \{a_1, a_2, a_3\}$. Suppose that for distinct $i, j \in N$ and each item $a \in A$, we have $V_i(i, a) = 0$ and $V_i(j, a) = -100$. One agent (say, 1) must receive at least two items, and the value of agent 2 is -200. However, agent 2's maximum value is 0, but it is not possible to attain this by reassigning one item. □

7.2 GFS and GFS1 Concepts

The incompatibility of PROP-Max-1 with negative valuations in Proposition 7.3 motivates us to propose a new fairness concept called *general fair share (GFS)*. For agent i and issue a, let $V_i^{min}(a) = \min_{a^t \in a^T} V_i(a^t)$.

Definition 7.4 (General Fair Share)**.** *The general fair share of agent i (GFS$_i$) is defined as*

$$\begin{aligned} GFS_i &= \frac{1}{n} \sum_{a \in A} V_i^{max}(a) + \frac{n-1}{n} \sum_{a \in A} V_i^{min}(a) \\ &= \sum_{a \in A} V_i^{min}(a) + \frac{1}{n} \sum_{a \in A} (V_i^{max}(a) - V_i^{min}(a)). \end{aligned}$$

[1] In their paper, this concept is simply called proportionality, but we refer to it as PROP-Max to distinguish it from another variant of proportionality that we will discuss later.

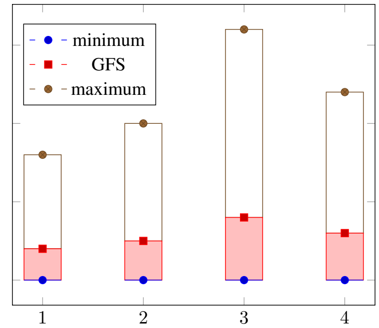

Figure 1: Illustration of GFS with four agents. Bottom lines and top lines denote the minimum and the maximum values each agent can receive among all possible allocations. Middle lines denote the GFS for each agent, and the colored region in each bar equals one-fourth of the difference between the maximum value and the minimum value of that agent.

An allocation π satisfies general fair share (GFS) if $V_i(\pi) \geq$ GFS$_i$ for all $i \in N$.

We next illustrate the intuition of GFS. Consider a GFS allocation π. For any agent i, the improvement that π offers upon agent i's worst allocation is at least $\frac{1}{n} \sum_{a \in A} (V_i^{max}(a) - V_i^{min}(a))$, i.e., $1/n$ of agent i's largest possible improvement—the improvement required by GFS is shown in the colored regions of Figure 1. That is, if we subtract $\sum_{a \in A} V_i^{min}(a)$ from $V_i(\pi)$, then we have

$$\begin{aligned} V_i(\pi) - \sum_{a \in A} V_i^{min}(a) &\geq \mathrm{GFS}_i - \sum_{a \in A} V_i^{min}(a) \\ &= \frac{1}{n} \sum_{a \in A} (V_i^{max}(a) - V_i^{min}(a)). \end{aligned}$$

Similarly to PROP-Max, GFS is too strong to guarantee corresponding allocations, so we relax it in the same manner as PROP-Max-1. We refer to this concept as *general fair share up to one item (GFS1)*.

Definition 7.5 (General Fair Share up to One Item)**.** *An allocation π satisfies general fair share up to one item (GFS1) if for all $i \in N$, there exists $a \in A$ such that*

$$V_i(\pi) - V_i(\pi(a)) + V_i^{max}(a) \geq GFS_i.$$

In other words, an allocation π satisfies GFS1 if for each agent i, there exists an item a such that changing the assignment of a from $\pi(a)$ to agent i's best assignment yielding $V_i^{max}(a)$ ensures that the value that agent i receives is at least her general fair share GFS$_i$.

With positive valuations, our GFS/GFS1 concepts are stronger than PROP-Max/PROP-Max-1.

Proposition 7.6. *For public decision making with positive valuations, GFS implies PROP-Max, and GFS1 implies PROP-Max-1.*

Proof. Given a GFS allocation π, for any agent i we have

$$V_i(\pi) - \frac{1}{n} \sum_{a \in A} V_i^{max}(a) \geq \frac{n-1}{n} V_i^{min}(a) \geq 0,$$

where we use the assumption of positive valuations for the latter inequality. Thus, GFS implies PROP-Max. The proof that GFS1 implies PROP-Max-1 is similar. $\square$

We next show that a GFS1 allocation always exists. Combined with Propositions 7.3 and 7.6, this means that GFS1 is a more suitable concept in public decision making than PROP-Max-1, both when valuations are only positive and when negative valuations are allowed.

7.3 Max-Min Round Robin

In this subsection, we present a polynomial-time algorithm "Max-Min Round Robin" that computes a GFS1 allocation for public decision making.

We give here a brief description of Max-Min Round Robin. For agent $i \in N$ and issue $a \in A$, let $\beta_i(a) = V_i^{max}(a) - V_i^{min}(a)$ denote the difference between the maximum and minimum value that agent i can receive from issue a. The Max-Min Round Robin algorithm works as follows: First, fix a round robin sequence of agents. Then, for each agent i's turn, let i determine the choice of some issue a in her favor such that $\beta_i(a)$ is the largest among all remaining issues. Repeat this procedure until there are no issues left. The details are described in Algorithm 1.

Algorithm 1: Max-Min Round Robin

Require: an instance of public decision making
Ensure: a GFS1 allocation

Let $\beta_i(a) = V_i^{max}(a) - V_i^{min}(a), \forall i \in N, a \in A$.
Fix a round robin sequence of agents, say $1, 2, \ldots, n$.
$\pi \leftarrow \emptyset, j \leftarrow 1$
while $A \neq \emptyset$ **do**
  For agent j's turn, find an issue a and a choice a^t such that
  - $\forall a' \in A, \beta_j(a) \geq \beta_j(a')$
  - $V_j^{max}(a) = V_j(a^t)$

  $\pi \leftarrow \pi \cup \{a^t\}$ {Choose a^t for issue a}
  $A \leftarrow A \setminus \{a\}$ {Remove a from A}
  $j \leftarrow (j \bmod n) + 1$ {Move on to the next agent}
end while
return an allocation π

Theorem 7.7. *Max-Min Round Robin returns a GFS1 allocation for public decision making in polynomial time.*

We reiterate that Theorem 7.7 works for both positive and negative valuations; its proof is provided in the full version of our paper (Aziz et al. 2022c).

8 Taxonomy of Fairness Concepts

In this section, we present a taxonomy of fairness concepts for fair allocation with externalities including existing and newly proposed ones (Figure 2). Formal definitions and proofs are provided in the full version of our paper (Aziz et al. 2022c).

Besides PROP-Max, another extension of proportionality is PROP-Ave, proposed by Seddighin, Saleh, and Ghodsi

$$\begin{tikzcd}
 & \text{EF} \arrow[dl] \arrow[dr] & & n\text{-P-PROP} \arrow[dl] \arrow[dr] & & \\
\text{EFX} \arrow[d] & & \text{2-P-PROP} & & \text{PROP-Ave} \arrow[r] \arrow[d] & \text{EMMS} \\
\text{EF1} \arrow[d] & & & & \text{GFS} \arrow[r] \arrow[d, "\text{positive}"] & \text{GFS1} \\
\text{EF}k & & & & \text{PROP-Max} &
\end{tikzcd}$$

Figure 2: Relationships among fairness concepts for fair allocation with externalities. An arrow from A to B denotes that A implies B. EF, EMMS, PROP-Ave, and PROP-Max under externalities were proposed in previous work.

(2021).[2] Maximin Share (MMS) is a relaxation of proportionality for fair division of indivisible items, introduced by Budish (2011). Seddighin, Saleh, and Ghodsi (2021) proposed Extended Maximin Share (EMMS) which generalizes MMS to the case of externalities.

For fair division without externalities, EF implies proportionality. However, we show that EF implies neither PROP-Max nor PROP-Ave when externalities exist. We propose a new notion k-Partial-Proportionality (k-P-PROP) that connects both EF and PROP-Ave. The intuition is that for any subset of agents $N' \subseteq N$ with $|N'| \leq k$, each agent $i \in N'$ should receive at least $1/|N'|$ of the total value she can receive from all items assigned to the group N'.

Note that Aziz et al. (2018) considered a general framework called $\mathcal{H}$-HG-PROP for defining fairness concepts when allocating indivisible items in the presence of a social graph. If there are no externalities and $\mathcal{H}$ contains all subsets of size at most k of the agents as hyperedges, then k-P-PROP is equivalent to $\mathcal{H}$-HG-PROP.

9 Conclusion

In this paper, we proposed several fairness concepts for fair division of indivisible items under externalities including EF1, EFX and GFS. We presented efficient algorithms for finding the corresponding fair allocations. An important open question that remains from our work is whether there always exists an EF1 allocation among three or more agents. Note that a positive answer to this question would generalize the corresponding result of Aziz et al. (2022a) for goods and chores without externalities. On the other hand, if the answer is negative, it would be reasonable to ask for the optimal relaxation EFk that can be attained. Finally, it will be interesting to consider more general valuation functions that are not necessarily additive. While the existence of EF1 and EFX allocations is known beyond additive valuations in certain settings (Lipton et al. 2004; Plaut and Roughgarden 2020), it remains to be seen whether these guarantees can be extended to incorporate externalities.

[2]These authors called the notion *average-share*.

Acknowledgments

This work was partially supported by ARC Laureate Project FL200100204 on "Trustworthy AI", by the Singapore Ministry of Education under grant number MOE-T2EP20221-0001, and by an NUS Start-up Grant. We thank the anonymous reviewers of AAAI 2023 for their comments.

References

Agarwal, A.; Elkind, E.; Gan, J.; Igarashi, A.; Suksompong, W.; and Voudouris, A. A. 2021. Schelling games on graphs. *Artificial Intelligence*, 301: 103576.

Aleksandrov, M.; Aziz, H.; Gaspers, S.; and Walsh, T. 2015. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2540–2546.

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; and Voudouris, A. A. 2022. Fair division of indivisible goods: a survey. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 5385–5393.

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.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2022a. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1): 3:1–3:21.

Aziz, H.; Li, B.; Moulin, H.; and Wu, X. 2022b. Algorithmic fair allocation of indivisible items: a survey and new questions. *ACM SIGecom Exchanges*, 20(1): 24–40.

Aziz, H.; Suksompong, W.; Sun, Z.; and Walsh, T. 2022c. Fairness concepts for indivisible items with externalities. *CoRR*, abs/2110.09066.

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)*, 7–13.

Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair allocation of indivisible goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 12, 284–310. Cambridge University Press.

Bouveret, S.; and Lemaître, M. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multi-Agent Systems*, 30(2): 259–290.

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

Brânzei, S.; Procaccia, A. D.; and Zhang, J. 2013. Externalities in cake cutting. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 55–61.

Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.

Bullinger, M.; Suksompong, W.; and Voudouris, A. A. 2021. Welfare guarantees in Schelling segregation. *Journal of Artificial Intelligence Research*, 71: 143–174.

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): 12:1–12:32.

Chaudhury, B. R.; Garg, J.; and Mehlhorn, K. 2020. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, 1–19.

Chauhan, A.; Lenzner, P.; and Molitor, L. 2018. Schelling segregation with strategic agents. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*, 137–149.

Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (ACM-EC)*, 629–646.

Darmann, A.; and Schauer, J. 2015. Maximizing Nash product social welfare in allocating indivisible goods. *European Journal of Operational Research*, 247(2): 548–559.

Elkind, E.; Patel, N.; Tsang, A.; and Zick, Y. 2020. Keeping your friends close: Land allocation with friends. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 318–324.

Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.

Gross-Humbert, N.; Benabbou, N.; Beynier, A.; and Maudet, N. 2021. Sequential and swap mechanisms for public housing allocation with quotas and neighbourhood-based utilities. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1521–1523.

Halpern, D.; Procaccia, A. D.; Psomas, A.; and Shah, N. 2020. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 370–383.

Li, M.; Zhang, J.; and Zhang, Q. 2015. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much! In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 589–595.

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 Economics and Computation (ACM-EC)*, 125–131.

Massand, S.; and Simon, S. 2019. Graphical one-sided markets. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 492–498.

Mishra, S.; Padala, M.; and Gujar, S. 2022. Fair allocation with special externalities. In *Proceedings of the 19th Pacific Rim International Conference on Artificial Intelligence (PRICAI)*, 3–16.

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

Seddighin, M.; Saleh, H.; and Ghodsi, M. 2021. Maximin share guarantee for goods with positive externalities. *Social Choice and Welfare*, 56(2): 291–324.

Suksompong, W.; and Teh, N. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences*, 117: 101–108.

Thomson, W. 2016. Introduction to the theory of fair allocation. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 11, 261–283. Cambridge University Press.

Velez, R. A. 2016. Fairness and externalities. *Theoretical Economics*, 11(1): 381–410.