

# Dividing Indivisible Items for the Benefit of All: It is Hard to Be Fair Without Social Awareness

Argyrios Deligkas<sup>1</sup>, Eduard Eiben<sup>1</sup>, Tiger-Lily Goldsmith<sup>1</sup>, Dušan Knop<sup>2</sup>, Šimon Schierreich<sup>2,3</sup>

<sup>1</sup>Royal Holloway, University of London

<sup>2</sup>Czech Technical University in Prague

<sup>3</sup>AGH University of Krakow

{argyrios.deligkas,eduard.eiben,tiger-lily.goldsmith}@rhul.ac.uk, {dusan.knop,schiesim}@fit.cvut.cz

## Abstract

In standard fair division models, we assume that all agents are selfish. However, in many scenarios, division of resources has a direct impact on the whole group or even society. Therefore, we study fair allocations of indivisible items that, at the same time, maximize *social impact*. In this model, each agent is associated with two additive functions that define their value and social impact for each item. The goal is to allocate items so that the social impact is maximized while maintaining some fairness criterion. We reveal that the complexity of the problem heavily depends on whether the agents are socially aware, i.e., they take into consideration the social impact functions. For socially unaware agents, we prove that the problem is NP-hard for a variety of fairness notions, and that it is tractable only for very restricted cases, e.g., if, for every agent, the valuation equals social impact *and* it is binary. On the other hand, social awareness allows for fair allocations that maximize social impact, and such allocations can be computed in polynomial time. Interestingly, the problem becomes again intractable as soon as the definition of social awareness is relaxed.

## 1 Introduction

The new equipment for your lab has arrived, and the lab director (LD) has to decide how to allocate it among lab members. Each member has their own preferences, i.e., *subjective valuations*, over the items of equipment. Meanwhile, LD anticipates the *impact* each member can have for the *whole* lab when allocated an item. This can depend on the background and technical skills of each member, and thus it can vary significantly *and* might not be aligned with members' valuations. Having said this, LD knows that the members will *compare* their bundle against their peers. Thus, LD would ideally like to find an allocation that maximizes social impact *and* is perceived as being "*fair*".

Consider the following scenario, though. Two items have to be allocated between reliable-Bill and sloppy-Joe. It is known that the impact of Bill will be 10 from each item, while Joe can contribute only 1 for each. This toy example demonstrates that even the most relaxed fairness criterion – envy-freeness up to one item (EF1) – cannot always be satisfied if we require maximization of social impact. In many

cases in the real world, though, people are *socially aware*. Thus, they are willing to accept a “non-fair” outcome if they realize that this is for the *benefit of all*.

The scenarios above might be seen as *fair division with externalities* (Velez 2016; Aziz et al. 2023b; Deligkas et al. 2024), where agents care not only for the items assigned to them, but also how exactly the remaining items are allocated. Nevertheless, this model is too general, so Flammini, Greco, and Varricchio (2025) have recently introduced a model that captures exactly the setting above: there is a set of goods and every agent is associated with two additive functions that define their value and social impact for each good. The goal is to allocate the goods so that the social impact is maximized while maintaining some fairness criterion. They have studied the “*price of fairness*”, which essentially quantifies how much worse in terms of social welfare a fair outcome is. In addition, they have formally introduced the notion of socially aware agents and which outcomes they perceive as fair. They showed that with socially aware agents, EF1 and social impact maximizing allocations always exist and can be computed in polynomial time. However, the complexity of the underlying computational problems was not examined in detail. Our goal is to extend the fairness solution concepts for this model and settle their complexity.

## Our Contribution

We provide a comprehensive study of the complexity of computing fair allocations that maximize social impact, which we term SIM allocations. We investigate seven different fairness criteria, all of which are strengthenings or weakenings of EF1. For each criterion, we examine the constraints on the valuation and social impact functions that allow for tractability (see Figure 1 for an overview).

Our initial set of technical results considers socially unaware agents, where we study two different dimensions of the problem and for each of them we provide complementary results between tractability and NP-hardness. Interestingly, we show that the problem behaves in the same way for every fairness criterion, albeit we need slight modifications in our proofs to formally establish this. First, we restrict the domain of the valuation and social impact functions. There, we show that the problem is NP-hard even when both functions are binary (Thm. 1), but it becomes tractable if for every agent both functions are binary *and* equal to each other

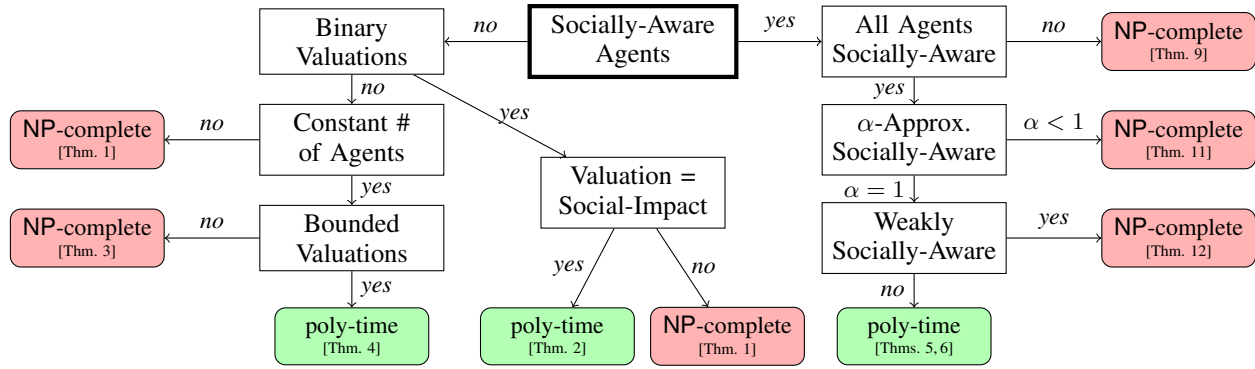


Figure 1: A basic overview of the complexity landscape of deciding whether a SIM and  $\mathcal{F}$ -fair allocation exist.

(Thm. 2). Then, we restrict the number of agents. Unfortunately, the problem is NP-hard even for two agents with binary social impact functions (Thm. 3). On the other hand, we derive a pseudo-polynomial time algorithm that solves the problem for any constant number of agents (Thm. 4). Hence, the problem is tractable when we have a constant number of agents, polynomially bounded valuation functions, and unrestricted social impact functions.

Then we consider the case where the agents are socially aware. Our first set of results for this setting adopts the definition of social awareness from (Flammini, Greco, and Varricchio 2025). Intuitively, we accept allocations that do not satisfy the “standard” fairness constraint, however, any bundle an agent envies (according to the chosen fairness criterion) cannot produce *strictly* more social impact if this agent gets it. We prove that social awareness allows for SIM and fair allocations that can be computed in polynomial time (Thms. 5 and 6). This leads us to the question of whether the definition of social awareness allows for an arbitrary fairness criterion. We prove that this is not the case by showing NP-hardness if we have “*fully envious*” agents (Thm. 7).

Finally, we ask whether we can relax social awareness and still get a fair and SIM allocation in polynomial time. Our first relaxation considers the case where not all agents are socially aware. Unfortunately, even if only one agent is socially unaware existence of SIM and fair allocations is not guaranteed and, in fact, the problem becomes NP-hard even when there are two agents (Thm. 9). On the positive side, we complement this negative result with a pseudo-polynomial time algorithm that works for any constant number of agents where some of them are socially unaware (Thm. 10). Then, we limit social awareness for each agent in two different ways. In the first case, we focus on  $\alpha$ -*approximately* socially aware agents, where  $\alpha \in [0, 1]$ , and an agent accepts an unfair outcome only if every envied bundle produces at least  $\alpha$  times more social impact under the current allocation. Strikingly, we get the same behavior as before: there is not always a solution and the problem is NP-hard for *every* possible  $\alpha$  (Thm. 11). In the second case, we define *weakly socially aware* agents. Intuitively, an agent accepts an unfair allocation if the proportional gain of their valuation from an envied bundle is less than the proportional gain in social impact (Def. 11). Once again, there is no guaranteed solution

and the problem is NP-hard even for two agents (Thm. 12). *The full version is available in (Deligkas et al. 2025b).*

## Related Work

Our paper lies in the area of fair division with indivisible items, where there is a flourish of results in the last 15 years (Lipton et al. 2004; Bouveret and Lang 2008; Budish 2011; Caragiannis et al. 2019); for excellent recent surveys, see (Amanatidis et al. 2023) and (Nguyen and Rothe 2023).

As we have mentioned before, (Flammini, Greco, and Varricchio 2025) is the most relevant paper to ours, since we extend their model and their computational results. Having said this, both papers fall into the larger, more general model of *fair division with externalities* (Velez 2016; Aziz et al. 2023b; Deligkas et al. 2024). In this model, each agent gets value from the *whole* allocation, i.e., they value each bundle depending on the agent who got it. Hence, we can see any instance in our setting as an instance of agents with externalities. On a similar note, Bu et al. (2023) studied the model in which there is an *allocator* with its own view of how the items should be allocated between agents, and the goal is to be fair to both the agents and the allocator. It should be highlighted that users deriving their utility not only from their own benefit but also from the benefit of (or to) others is a well-described phenomenon in social sciences and behavioral economics; see, e.g., (Overvold 1980; Andreoni 1990; Huang et al. 2022; Thomas and Thomas 2025).

A different point of view is to consider our problem as *constrained* fair division (Suksompong 2021). We identify two settings that are closest to ours. The first regards *welfare-maximizing* fair allocations. Observe that when social impact equals the valuation of each agent, fair SIM allocations coincide with welfare-maximizing allocations that are fair, which was proven to be NP-hard (Aziz et al. 2023a) in general. The second is related to *orientations* (Christodoulou et al. 2023; Zhou et al. 2024; Deligkas et al. 2025a; Afshinmehr et al. 2025; Blažej et al. 2025). Here, every agent is associated with a subset of items they are interested in. The goal is to find an allocation where every agent gets items *only* from their subset. In our case, agents are only allowed to get items for which they are social impact maximizers. The crucial difference though is that in our case, agents value the remaining items as well.

## 2 Preliminaries

We assume a set  $M$  of  $m$  items which need to be partitioned between a set  $N$  of  $n$  agents. Each agent  $i \in N$  is associated with a valuation function  $v_i: 2^M \rightarrow \mathbb{N}_0$  which assigns to each subset  $S$  of items, called a bundle, a numerical value expressing how satisfied  $i$  is when  $S$  is allocated to them. Additionally, there is a social impact function  $s_i: 2^M \rightarrow \mathbb{N}_0$  for every agent  $i \in N$  expressing what social impact agent  $i$  generates for the society with each bundle  $S \subseteq M$ . Throughout the paper, we assume that both the valuation functions and the social impact functions are *additive*, meaning that for every agent  $i \in N$  and every  $S \subseteq M$  we have  $v_i(S) = \sum_{g \in S} v_i(\{g\})$  and  $s_i(S) = \sum_{g \in S} s_i(\{g\})$ , respectively, and *normalized*, i.e.,  $v_i(\emptyset) = s_i(\emptyset) = 0$ . An instance of *fair division with social impact* is then a quadruple  $\mathcal{I} = (N, M, (v_i)_{i \in N}, (s_i)_{i \in N})$ .

A *partial allocation* is a tuple  $\mathcal{A} = (A_1, \dots, A_n)$ , where  $A_i \subseteq M$  is a (potentially empty) bundle received by agent  $i \in N$ , for which it holds that  $A_i \cap A_j = \emptyset$  for each pair of distinct  $i, j \in N$ . If additionally  $\bigcup_{i=1}^n A_i = M$ , then we say  $\mathcal{A}$  is a *complete allocation*. In the remainder of the paper, we are interested in complete allocations. Let  $\mathcal{A}$  be a (partial) allocation. We naturally extend the valuations and social impact functions from bundles to allocations by setting  $v_i(\mathcal{A}) = v_i(A_i)$  and  $s_i(\mathcal{A}) = s_i(A_i)$ . Moreover, we denote the overall social impact generated by  $\mathcal{A}$  as  $\text{SI}(\mathcal{A}) = \sum_{i \in N} s_i(A_i)$ .

**Solution Concepts.** The first solution concept crucial for our work is the requirement to maximize social impact. Specifically, we require that in each solution the sum of social impacts is the maximum possible.

**Definition 1.** An allocation  $\mathcal{A}$  is *social impact maximizing (SIM)* if for every allocation  $\mathcal{A}'$  we have  $\text{SI}(\mathcal{A}) \geq \text{SI}(\mathcal{A}')$ .

It is easy to see that SIM allocations always exist and can be found efficiently. It is enough to allocate each item to an agent that maximizes the social impact for it. However, such allocations may be very unfair. Assume an instance with social impact functions that assign one to every combination of an agent and an item. It may happen that the previous procedure allocates all items to a single agent  $i$ , which is arguably very unfair to all remaining agents.

To overcome this issue, we combine SIM requirement with various notions of fairness. We focus on axioms based on pairwise comparison of agents' bundles. The prototypical concept in this line of research is *envy-freeness* (Foley 1967), which requires that each agent prefers their own bundle before the bundles of other agents and is defined as follows.

**Definition 2.** An allocation  $\mathcal{A}$  is called *envy-free (EF)* if for any pair of agents  $i, j \in N$  we have  $v_i(A_i) \geq v_i(A_j)$ .

It is well known that envy-free allocations are not guaranteed to exist, even without social impact functions. Simply assume an instance with two agents and one item strictly positively valued by both agents; the agent who does not receive this item is always envious. Moreover, deciding the existence of EF allocations is a notoriously hard computational problem. Therefore, in the rest of the paper, we focus on several relaxations of envy-freeness for which, without

social impact function, existence is guaranteed and there are efficient algorithms finding such desirable outcomes.

We start with a relaxation called *envy-freeness up to one item* (Lipton et al. 2004; Budish 2011) which requires that if there is envy from agent  $i$  to agent  $j$ , then this envy can be eliminated by the removal of one item from  $j$ 's bundle.

**Definition 3.** An allocation  $\mathcal{A}$  is called *envy-free up to one item (EF1)* if for every pair of agents  $i, j \in N$  there exists an item  $g \in M$  such that  $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ .

Note that in the definition of EF1, the item removed from  $j$ 's bundle can be *specific* for each agent  $i \in N$ . A stronger variant of this notion requires that the removal of one *universal* item from  $A_j$  prevents potential envy from all other agents. This notion is due to Conitzer et al. (2019).

**Definition 4.** An allocation  $\mathcal{A}$  is called *strongly envy-free up to one item (sEF1)* if for every  $j \in N$  there exists an item  $g \in M$  such that  $v_i(A_i) \geq v_i(A_j \setminus \{g\})$  for every  $i \in N$ .

As noted by Wu, Zhang, and Zhou (2025), in traditional EF1 (or sEF1), it is assumed that all agents have equal obligations. However, this assumption can be very unrealistic in many scenarios. E.g., recall our scientific laboratory example from the beginning of the paper. It is reasonable to expect that the division of the equipment should also consider factors such as seniority or merit. Therefore, we also study *weighted envy-freeness* (Chakraborty et al. 2021; Aziz et al. 2024; Suksompong 2025). In this model, each agent  $i \in N$  is additionally associated with its *weight*  $w_i$ .

**Definition 5.** An allocation  $\mathcal{A}$  is called *weighted envy-free up to one item (wEF1)* if for every pair of agents  $i, j \in N$  there exists an item  $g \in M$  such that

$$\frac{v_i(A_i)}{w_i} \geq \frac{v_i(A_j \setminus \{g\})}{w_j}.$$

Clearly, if all the weights are the same, wEF1 is equivalent to EF1. Therefore, wEF1 is a generalization of EF1. Similarly to sEF1, we can analogously define *strong weighted envy-freeness up to one item (swEF1)*.

A different point of criticism towards EF1 is that it allows the envy between  $i$  and  $j$  to vanish simply by the removal of the best good from the  $j$ 's bundle, regardless of how valued this item is for the agent  $i$ . For example, assume that  $i$  and  $j$  are allocated one item each, an  $i$  values  $A_j$  as 1000 and  $A_i$  as 1. Although formally there is no EF1-envy from  $i$  towards  $j$ , agent  $i$  may still consider the allocation very unfair, since there is a huge difference in the valuations of both bundles. Motivated by this, Barman et al. (2018) introduced *envy-freeness up to one less preferred good*.

**Definition 6.** An allocation  $\mathcal{A}$  is called *envy-free up to one less preferred item (EFL)* if for every pair of agents  $i, j \in N$  at least one of the following two conditions hold:

1.  $A_j$  contains at most one item which is positively valued by  $i$ , or
2. there exists  $g \in A_j$  such that  $v_i(A_i) \geq v_i(A_j \setminus \{g\})$  and  $v_i(A_i) \geq v_i(\{g\})$ .

Finally, in light of our negative results, we also study a relaxation of EF1 that, instead of removing one item to eliminate envy, uses a transfer of one item. In the literature, this

notion is usually called weak envy-freeness up to one item. However, to avoid possible confusion with weighted EF1, we call it *envy-freeness up to one transfer*.

**Definition 7.** An allocation  $\mathcal{A}$  is called *envy-free up to one transfer (tEF1)* if for every pair of agents  $i, j \in N$  such that  $v_i(A_i) < v_i(A_j)$  there exists an item  $g \in A_j$  so that  $v_i(A_i \cup \{g\}) \geq v_j(A_j \setminus \{g\})$ .

It is easy to see that whenever an allocation is EF1, it is tEF1, as if the item  $g$  that removes the envy from  $i$  to  $j$  is additionally moved to  $A_i$ , it cannot decrease the value of  $A_i$  for  $i$ . In the opposite direction, tEF1 does not imply EF1. The example here is an allocation where  $A_j$  contains two items, both valued 1 by  $i$ , and  $A_i = \emptyset$ . If we transfer one item from  $A_j$  to  $A_i$ , the value of agent  $i$  for both bundles is 1, so this allocation is tEF1. However, the removal of a single item does not eliminate envy, so it is not EF1.

**Remark 1.** It is not hard to see that, given an allocation  $\mathcal{A}$ , one can check whether this allocation is SIM and  $\mathcal{F}$ -fair, where  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL\}$ , in polynomial time. Hence, deciding whether a SIM and  $\mathcal{F}$ -fair allocation exists is trivially in NP and we will not stress this explicitly in our hardness proofs.

### 3 Socially Unaware Agents

In this section, we explore the computational complexity of deciding the existence of SIM and fair allocations. In general, hardness for EF1 follows from the result of Aziz et al. (2023a), who showed that deciding whether an EF1 allocation maximizing social-welfare<sup>1</sup> exists is NP-complete. This corresponds to the case with  $s_i = v_i$  for every  $i \in N$ . We extend their result by studying also additional notions of fairness and giving a much more detailed picture of its complexity, even in as restricted domains as binary valuations, or with a constant number of agents.

First, we observe that, without loss of generality, we can assume that the social impact function is binary.

**Lemma 1.** For any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL, EF\}$ , an instance  $\mathcal{I} = (N, M, (v_i)_{i \in N}, (s_i)_{i \in N})$  admits a SIM and  $\mathcal{F}$ -fair allocation if and only if an instance  $\mathcal{J} = (N, M, (v_i)_{i \in N}, (s_i^*)_{i \in N})$ , where  $s_i^*(g) = 1$  if  $i \in \arg \max_{i \in N} s_i(g)$  and 0 otherwise, admits a SIM and  $\mathcal{F}$ -fair allocation.

With the previous lemma in hand, we show our main hardness result of this section, which settles that even if both the valuations and the social impact functions are binary, the problem is NP-complete. Our strategy in the reduction is to introduce some initial unfairness between agents so that the rest of the items have to be allocated in an envy-free way, which is known to be computationally hard (Aziz et al. 2015; Hosseini et al. 2020).

**Theorem 1.** For any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL\}$ , it is NP-complete to decide whether a SIM and  $\mathcal{F}$ -fair allocation exist, even if both the valuations and the social impact functions are binary.

<sup>1</sup>Social welfare of an allocation is simply the sum of utilities of all agents in this allocation.

Note that in the previous results, the social impact functions and the valuations are not the same. If we additionally assume that the valuations and social impacts are the same, we obtain a positive result modifying the Yankee-swap algorithm of Viswanathan and Zick (2025) and the weighted round-robin of Chakraborty et al. (2021), respectively.

**Theorem 2.** If the valuations are binary and  $v_i = s_i$  for every  $i \in N$ , a SIM and  $\mathcal{F}$ -fair allocation is guaranteed to exist and can be found in polynomial time for any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL\}$ .

Next, we turn our attention to instances with a constant number of agents. We show that with two agents already, deciding the existence of SIM and fair allocation is computationally intractable.

**Theorem 3.** For any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL\}$ , it is NP-complete to decide whether a SIM and  $\mathcal{F}$ -fair allocation exist, even if the social impact functions are binary and  $n = 2$ .

However, the previous hardness result is highly dependent on the fact that the agents' valuations for the items are exponential in the input size. That is, the result shows only weak NP-hardness. In the following, we complement the hardness and show that if the number of agents is a fixed constant and the valuations are polynomially bounded in the input size, then there is an efficient algorithm. Notably, the algorithm is not based on an explicit DP formulation, as one might expect, but rather models our problem as graph reachability. Thanks to this, the algorithm can be implemented very efficiently by generating the graph on-demand and finding the solution using, e.g., a variant of the A\* algorithm (Hart, Nilsson, and Raphael 1968; Dechter and Pearl 1985; Korf 1997; Zhou and Zeng 2015).

**Theorem 4.** For every fixed constant number of agents  $n$ , there exists a pseudo-polynomial-time algorithm deciding whether a SIM and  $\mathcal{F}$ -fair allocation exist for any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL, EF\}$ .

### 4 Socially Aware Agents

Motivated by the experimental work of Herreiner and Puppe (2009) and Hosseini et al. (2025), which explore what real-life users perceive to be fair outcomes, Hosseini (2024) argues that “developing fair algorithms may require axioms that are able to capture solutions that take the social context into account beyond perceived envy.” In this section, we build up on this idea and follow the approach introduced by Flammini, Greco, and Varricchio (2025) and assume *socially aware* agents. Intuitively, an agent is socially aware if they are willing to accept a formally unfair outcome in case it leads to a greater social impact.

**Definition 8.** We say that an allocation  $\mathcal{A}$  is *envy-free up to one item with socially aware agents (SA-EF1)*, if for each pair of agents  $i, j \in N$ , one of the following conditions hold:

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

Other fairness notions, such as SA-wEF1 and SA-EFL, are then defined analogously to SA-EF1, just with the first

condition modified accordingly. Based on the above definition, we will say that an agent  $i$  SA-envies  $j$ , if and only if  $v_i(A_i) < v_i(A_j)$  and  $s_i(A_j) \geq s_j(A_j)$ .

First, we observe that the second condition in the definition of social awareness is meaningful in the sense that it captures the intuition of socially aware agents but is not strong enough to allow for some very unfair solutions.

**Observation 1.** *If we replace  $<$  with  $\leq$  in the second part of the definition of SA-EF (and its relaxations), any SIM allocation is a solution.*

Flammini, Greco, and Varricchio (2025) proved that there is a modification of the famous envy-cycle elimination algorithm (Lipton et al. 2004) that always finds an SA-EF1 allocation. We extend their result in two ways. In our first result, we show that SA-swEF1 allocations are also guaranteed to exist. Interestingly, our approach is based on a picking sequence similar to the round-robin algorithm. More precisely, the algorithm is a modification of the weighted picking sequence protocol proposed by Chakraborty et al. (2021), which reduces to a round-robin approach in the case of identical agent weights. However, we only allow agents to pick goods for which they maximize social impact, and we force agents to skip their turn whenever they do not maximize the social impact of any unallocated good. The proof then basically follows from Chakraborty et al. (2021) and a simple observation that agent  $i$  can SA-envy agent  $j$  only if  $j$  chose only goods for which  $i$  maximizes social impact. Therefore, from the point of view of  $i$ , we can completely ignore all agents that get an item for which  $i$  does not maximize social impact. On the other hand, if we consider an auxiliary instance in which agents only value items for which they maximize impact and get a dummy item that everyone values 0 whenever they skip a turn we get an allocation that is swEF1 according to Chakraborty et al. (2021) and in which: 1) All agents get same bundle plus some dummy items. 2) Every agent values every bundle for which the agent maximizes the social impact and all non-dummy items in it exactly the same.

**Theorem 5.** *SIM and SA-swEF1 allocation is guaranteed to exist and can be found in polynomial time.*

The above approach naturally extends to all fairness notions we consider except for SA-EFL. It is not a coincidence, as it turns out that *any* picking sequence protocol that computes picking sequence disregarding the valuation functions is doomed from the start for SA-EFL. Observe that the first agent in the sequence that is allowed to pick more than once can choose an item that everyone, including the agent, considers more valuable than the sum of all remaining items (not assigned until the first pick of the agent); to give a more concrete example, consider the following.

**Example 1.** *Let us assume an instance with two agents, four items, and both agents have identical social impact and identical valuations for these items. Let these identical valuations be 100, 1, 1, 1. Consider the round-robin picking sequence, in which each agent picks two items. So, the final allocation is  $(\{100, 1\}, \{1, 1\})$ . Clearly, agent 2 with the bundle  $\{1, 1\}$  envies agent 1 with the bundle  $\{100, 1\}$ ,*

*and only removal of 100 would remove this envy. However,  $100 > 2 = 1 + 1$ , so this allocation is not EFL. The only way to get an EFL allocation for this instance is to give the large item to one agent and the remaining items to the other.*

However, we obtain a positive result for EFL anyway. The algorithm to obtain SA-EFL is a combination of the ideas for SA-EF1 from Flammini, Greco, and Varricchio (2025) and for EFL from Barman et al. (2018). The algorithm is basically identical to the algorithm for finding EFL allocation from Barman et al. (2018); however, agents are only allowed to pick goods for which they maximize social impact, and instead of using the standard envy-graph, whose vertices are agents and directed edges represent envy from one agent to another, we consider the so-called SA-envy graph, where directed edge represent SA-envy. The correctness follows again from the observation that throughout the computation, each agent maximizes social impact of their bundle. So, once a bundle  $A_j$  receives an item for which an agent  $i$  does not maximize the social impact, agent  $i$  will never envy an agent that has (a superset of) the bundle  $A_j$ . Hence, considering only the SA-envy graph is sufficient. The proof that the EFL property is satisfied if two agents SA-envy one another is then analogous to the proof of Barman et al. (2018).

**Theorem 6.** *SIM and SA-EFL allocation is guaranteed to exist and can be found in polynomial time.*

Previous results indicate that whenever we combine a fairness notion with social awareness, we obtain an existence guarantee accompanied by polynomial-time algorithms. One may wonder whether these very positive results are not caused simply by the social awareness of our agents. Maybe it is the case that whenever we have socially aware agents, we can actually achieve an arbitrary fairness notion. To formally study this possibility, we assume fully envious but socially aware agents. An agent  $i \in N$  is not satisfied with an allocation  $\mathcal{A}$  whenever another agent  $j$  exists with a non-empty bundle.

**Definition 9.** *We say that an allocation  $\mathcal{A}$  is SA- $\emptyset$ , if for each pair of agents  $i, j$  either  $A_j = \emptyset$  or  $s_i(A_j) < s_j(A_j)$ .*

However, as the next theorem proves, the problem of deciding the existence of SIM and SA- $\emptyset$  allocations is computationally hard. That is, socially aware agents are not enough to guarantee an arbitrary fairness notion.

**Theorem 7.** *It is NP-complete to decide whether there is a SIM and SA- $\emptyset$  allocation, even if the valuations and social impacts are binary and  $v_i = v_j$  for every  $i, j \in N$ .*

We conclude with one positive algorithmic result for SA- $\emptyset$ . Specifically, we show that whenever the number of agent-types or item-types is bounded, there is an efficient algorithm deciding the existence of SIM and SA- $\emptyset$  allocations. This result, especially for the latter parameterization, is not without interest, as in the area of fair division of indivisible items, parameterization by the number of different item-types is a notoriously hard open problem; see, e.g., (Eiben et al. 2023; Nguyen and Rothe 2023; Bredereck et al. 2023).

**Theorem 8.** *When parameterized by the number of item-types or the number of agent-types, an FPT algorithm deciding the existence of SIM and SA- $\emptyset$  allocation exist.*

## 5 Limited Social Awareness

In the previous section, we studied the impact of social awareness on the existence of fair and social impact maximizing allocations. It turned out that with socially aware agents, such a desirable allocation always exists and can be found efficiently. However, the notion of social awareness, as defined, is somewhat idealistic, and in practical scenarios, it may not be the case that agents accept very bad bundles just because of the social impact they generate with these chores. Or, some agents can be even very selfish and may be interested only in fairness and not in the social impact at all.

### Some Agents Socially Unaware

The first limitation of social awareness we study is a setting with some agents socially aware and some socially unaware. This is arguably a very natural setting. Unfortunately, as we show in the following example, even if we allow as few as one agent to be socially unaware, we lose all existence guarantees.

**Example 2.** Let  $a_1$  be a socially unaware agent,  $a_2$  be socially aware agent, and  $g_1, g_2$  be two identical elements such that for  $j \in [2]$  we have  $v_{a_1}(g_j) = v_{a_2}(g_j) = 10$ , and  $s_{a_1}(g_j) = 0$  and  $s_{a_2}(g_j) = 1$ . Due to the social impact function, both items have to be allocated to agent  $a_2$ . However,  $a_1$  is now very envious of agent  $a_2$ , and this envy cannot be eradicated by removing any item from  $A_1$ .

Note that we can extend the previous example to show the non-existence of the SA-EF $r$  solution for arbitrary  $r \in \mathbb{N}$ , simply by adding at least  $r + 1$  copies of our items. Hence, even if we have exactly one socially unaware agent, we cannot provide any guarantee on (relaxations of) EF.

In the next result, we draw the situation even more negatively. Specifically, we show that even one socially unaware agent  $a_1$  makes the decision of the existence of a social impact maximizing allocation, which is EF1 for  $a_1$  and SA-EF1 for all other agents, computationally hard.

**Theorem 9.** It is NP-complete to decide if a SIM allocation which is EF1 for agent  $a_1$  and SA-EF1 for every  $a_i \in N \setminus \{a_1\}$  exist, even if there are only two agents.

In the previous hardness result, we exploited the fact that some items are forced to be allocated to the not socially aware agent. In the following, we show that in the case of two agents, this property is crucial to draw an instance computationally hard because if there is no such item, we can decide the instance in polynomial time.

**Proposition 1.** If there are two agents and no item where  $a_2$  is the unique agent maximizing social impact, an allocation that is simultaneously SIM, EF1 for agent  $a_1$ , and SA-EF1 for agent  $a_2$  always exists and can be found in poly-time.

We conclude this subsection with a pseudo-polynomial algorithm similar to the one of Theorem 4, complementing the hardness from Theorem 9.

**Theorem 10.** For every fixed number of agents  $n$  and any  $c \leq n$ , there exists a pseudo-polynomial time algorithm deciding whether a SIM allocation, which is simultaneously  $\mathcal{F}$ -fair for  $a_i, i \in [c]$ , and SA- $\mathcal{F}$ -fair

for any  $a_i, i \in [c + 1, \dots, n]$ , exists for any  $\mathcal{F} \in \{tEF1, EF1, sEF1, wEF1, swEF1, EFL\}$ .

### Approximate Awareness

A clear outcome of the previous subsection is that having all agents socially aware is a necessary condition for the problem to be tractable. However, having fully socially aware agents seems highly unrealistic. Suppose, for example, an instance with two agents  $a_1$  and  $a_2$ , both socially aware, and one item  $g$  valued as 1000 by both agents and with  $s_{a_1}(g) = 1$  and  $s_{a_2}(g) = 0$ . By the social impact functions, we have to allocate the single item to  $a_1$ . Observe that  $a_2$  values its bundle at 0 and  $a_1$ 's bundle at 1000. By Definition 8, this tremendous envy is outweighed by the social impact, which is, however, larger only by one. Motivated by this obvious overestimation of the social awareness, in the rest of this section we study several relaxations of SA.

The first approach we explore is with agents who are not fully socially aware but require their social impact to be significantly larger compared to the social impact of every other agent with the same bundle in order to accept an unfair allocation. We define the notion as follows.

**Definition 10** ( $\alpha$ -SA-EF1). We say that an allocation  $\mathcal{A}$  is  $\alpha$ -envy-free up to one item with socially aware agents ( $\alpha$ -SA-EF1) for  $0 \leq \alpha \leq 1$ , if for each pair of agents  $i, j \in N$ , one of the following conditions hold: 1/ there is no EF1-envy from  $i$  towards  $j$ , or 2/  $s_i(A_j) < \alpha \cdot s_j(A_j)$ .

It is easy to see that 1-SA-EF1 is equivalent to SA-EF1 and that 0-SA-EF1 is equivalent to the standard EF1. Therefore, 1-SA-EF1 allocations are guaranteed to exist and can be found in polynomial time, and deciding the existence of 0-SA-EF1 is computationally hard. Any value of  $\alpha$  between these two extreme values then naturally limits social awareness of our agents. In our first result, we show that for any  $\alpha$ -relaxation,  $\alpha < 1$ , of social awareness, unfortunately, the SIM and  $\alpha$ -SA-EF1 allocations are not guaranteed to exist.

**Proposition 2.** For any  $\alpha \in \mathbb{R}, 0 \leq \alpha < 1$ , SIM and  $\alpha$ -SA-EF1 is not guaranteed to exist.

*Proof.* Let  $\alpha$  be as stated. We construct an instance with two agents  $a_1$  and  $a_2$  and two items  $g_1$  and  $g_2$ . The valuation functions are identical and assign the same value of 1 to both items. The social impact functions are  $s_{a_1}(g_j) = 1$  and  $s_{a_2}(g_j) = \alpha$ . Due to the social impact function, both items have to be allocated to agent  $a_1$ . Moreover,  $v_{a_2}(A_1) = 2$  and  $v_{a_2}(A_2) = 0$ , meaning that agent  $a_2$  is EF1-envious towards  $a_1$ . The social impacts are  $s_{a_1}(A_1) = 2$  and  $s_{a_2}(A_1) = 2\alpha$ . That is, the second condition from the definition of  $\alpha$ -SA-EF1 is not met  $\square$

Moreover, the following result shows that relaxing social awareness not only eliminates any existence guarantee, but also renders the associated decision problem computationally intractable. The reduction is similar to the one used to prove Theorem 9.

**Theorem 11.** For any  $\alpha \in \mathbb{R}, 0 \leq \alpha < 1$ , it is NP-complete to decide whether there is a SIM and  $\alpha$ -SA-EF1 allocation, even if there are only two agents.

Both Proposition 2 and Theorem 11 hold also in the setting where only agent  $a_1$  is not fully socially aware, which paints our results even stronger.

### Weak Social Awareness

Although  $\alpha$ -SA seems to be a natural relaxation of social awareness, it fails to address one of our most important criticisms. Specifically, the level of social awareness is completely unrelated to the amount of envy between two agents (unless we set  $\alpha = v_i(A_i)/v_j(A_j)$ , which is specific for every pair of  $i, j \in N$ ). In order to also take into account this perspective, in the remainder of this section, we explore a novel notion of *weak social awareness* (WSA). Intuitively, the notion formalizes the idea that if an agent  $i$  is envious of the agent  $j$  and the bundle  $A_j$  is twice as good as the bundle  $A_i$  according to the agent  $i$ , the social awareness condition should override envy if the social impact generated by  $i$  from  $A_i$  is at least twice the social impact of  $j$  with  $A_j$ . Before we formally define weak social awareness, we provide an example that illustrates how powerful standard social awareness is.

**Example 3.** *Let us have two agents 1 and 2, three items  $g_1, g_2,$  and  $g_3,$  and let the valuations and social impacts be as follows.*

	$g_1$	$g_2$	$g_3$
1	$s = 1, v = 1$	$s = 1, v = 5$	$s = 0, v = 5$
2	$s = 0, v = 5$	$s = 1, v = 5$	$s = 1, v = 1$

*Let the allocation  $\mathcal{A}$  be with  $A_1 = \{g_1\}$  and  $A_2 = \{g_3, g_2\}$ . Observe that  $v_1(A_1) = 1$  and  $v_1(A_2) = 10$ . That is, the bundle  $A_2$  is ten times more valuable than  $A_1$  according to agent 1, and the envy cannot be eliminated by the removal of a single item. However, such an allocation is SA-EF1, since agent 2 generates strictly more social impact with  $A_2$ . However, the social impact of agent 2 is only twice the social impact agent 1 would generate with bundle  $A_2$ .*

**Definition 11.** *We say that an allocation  $\mathcal{A}$  is envy-free up to one item with weakly socially aware agents (WSA-EF1), if for every pair of agents  $i, j \in N$ , one of the following conditions hold: 1/ there is no EF1-envy from  $i$  towards  $j$ , or 2/  $v_i(A_j) \cdot s_i(A_j) \leq v_i(A_i) \cdot s_j(A_j)$ .*

The WSA variants for other fairness notions are then defined analogously.

**Example 4.** *Recall the instance from Example 3. Clearly, the presented allocation  $\mathcal{A}$  is not EF1 from the perspective of agent 1. At the same time, we have  $v_1(A_2) \cdot s_1(A_1) = 10 \cdot 1 = 10$  and  $v_1(A_1) \cdot s_2(A_2) = 1 \cdot 2 = 2$ . Therefore, neither condition 2 from the definition of WSA-EF1 is satisfied, which means that  $\mathcal{A}$  is not a WSA-EF1 allocation.*

It is easy to see that the instance from Example 3 admits no SIM and WSA-EF1 allocation, as only other possible allocation is with  $A_1 = \{g_1\}$  and  $A_2 = \{g_2, g_3\}$ , which is also not WSA-EF1 by symmetric arguments. Moreover, also with weakly SA agents, the problem remains intractable.

**Theorem 12.** *For any  $\mathcal{F} \in \{tEF1, sEF1, wEF1, swEF1, EF1, EFL\}$ , it is NP-complete to decide whether SIM and WSA- $\mathcal{F}$ -fair allocation exist, even if the social impact function is binary and  $n = 2$ .*

## 6 Discussion

We have presented a comprehensive study of computing allocations that maximize social impact while maintaining some fairness criterion. Our results completely resolve the complexity of the problem for goods when the fairness criterion is based on EF1 or EFL.

It is evident from our results that to achieve SIM, social awareness is essential for tractability. An immediate question is whether there exist other natural definitions of social awareness that allow for SIM and fair allocations.

**Relaxing SIM.** One obvious direction for future research is to relax the SIM constraint, as someone can argue that SIM is too restrictive and *this* is the reason for non-existence of fair allocations for socially unaware agents; Flammini, Greco, and Varricchio (2025) show that we cannot get better than  $\mathcal{O}(n)$ -approximation for SIM under EF1 fairness. What if we replace SIM with a Pareto optimal allocation with respect to social impact? Unfortunately, there is no positive news for this: the proof of Theorem 3 shows that the problem remains NP-complete.

**Other Fairness Notions.** A complementary approach to the question above is to further relax the fairness criterion. Does there always exist an allocation that is SIM and *approximately* fair, for some established fairness notion? A different idea is to try to strengthen the fairness guarantee for socially aware agents. Can we prove the existence of SIM and EFX allocations with socially aware agents for the settings that admit EFX allocations? What about SIM allocations satisfying epistemic EFX, that are guaranteed to exist and can be efficiently found (Aziz et al. 2018; Caragiannis et al. 2023; Akrami and Rathi 2025)?

**Division of Chores.** Another exciting avenue, which we believe deserves further study, is the setting where we have *chores* to allocate (Aziz 2016; Aziz et al. 2017, 2022). This is a very natural setting with many practical applications; think of the admin tasks in a CS department. Our initial results provided in the appendix show that, again, social unawareness does not allow for SIM and EF1 allocations. Besides this, even with social awareness, it is unclear whether we can always find SIM-EF1 allocations. In fact, the round-robin approach fails when we have chores.

Similarly, it is unclear whether algorithms based on envy-cycle elimination can be extended to SA-EF1 for chores. The main issue is that in the setting of chores, it is possible that agent  $i$  does not SA-envy bundle  $A_j$ , when agent  $j$  has the bundle, as agent  $j$  has a smaller impact on the bundle of agent  $i$ ; however, if  $A_j$  is moved to some other agent  $k$ , who has same social impact for the  $i$ -th's bundle as  $i$ , the agent  $i$  can start SA-envying the same bundle. This is a problem, as obtaining SIM allocation requires eliminating the cycle only in the SA-envy graph. However, this means that it is possible that the agent to whom we allocate a chore already envies some bundle, just not SA-envies it on the current agent. Moving this bundle then can introduce envy that is not SA-EF1. Hence, the discussion above is a strong indication that if SA-EF1 allocations do indeed exist for chores, they require a different algorithmic technique.

## Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 101002854). This work was co-funded by the European Union under the project Robotics and advanced industrial production (reg. no. CZ.02.01.01/00/22\_008/0004590). AD acknowledges the support of the EPSRC grant EP/X039862/1. SS acknowledges the additional support of the Grant Agency of the Czech Technical University in Prague, grant No. SGS23/205/OHK3/3T/18.



## References

- Afshinmehr, M.; Danaei, A.; Kazemi, M.; Mehlhorn, K.; and Rathi, N. 2025. EFX Allocations and Orientations on Bipartite Multi-Graphs: A Complete Picture. In Das, S.; Nowé, A.; and Vorobeychik, Y., eds., *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '25*, 32–40. Richland, SC: IFAAMAS.
- Akrami, H.; and Rathi, N. 2025. Epistemic EFX Allocations Exist for Monotone Valuations. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI '25*, 13520–13528. Washington, DC, USA: AAAI Press.
- Amanatidis, G.; Aziz, H.; Birmpas, 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.
- Andreoni, J. 1990. Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal*, 100(401): 464–477.
- Aziz, H. 2016. Computational Social Choice: Some Current and New Directions. In Kambhampati, S., ed., *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI '16*, 4054–4057. IJCAI/AAAI Press.
- Aziz, H.; Bouveret, S.; Caragiannis, I.; Giagkousi, I.; and Lang, J. 2018. Knowledge, Fairness, and Social Constraints. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI '18*, 4638–4645. Palo Alto, CA, USA: AAAI Press.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2022. Fair Allocation of Indivisible Goods and Chores. *Autonomous Agents and Multiagent Systems*, 36(1): 3.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015. Fair Assignment of Indivisible Objects Under Ordinal Preferences. *Artificial Intelligence*, 227: 71–92.
- Aziz, H.; Huang, X.; Mattei, N.; and Segal-Halevi, E. 2023a. Computing Welfare-Maximizing Fair Allocations of Indivisible Goods. *European Journal of Operational Research*, 307(2): 773–784.
- Aziz, H.; Li, B.; Moulin, H.; Wu, X.; and Zhu, X. 2024. Almost Proportional Allocations of Indivisible Chores: Computation, Approximation and Efficiency. *Artificial Intelligence*, 331: 104118.
- Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017. Algorithms for Max-Min Share Fair Allocation of Indivisible Chores. In Singh, S.; and Markovitch, S., eds., *Proceedings of the 31st AAAI Conference on Artificial Intelligence, AAAI '17*, 335–341. Palo Alto, CA, USA: AAAI Press.
- Aziz, H.; Suksompong, W.; Sun, Z.; and Walsh, T. 2023b. Fairness Concepts for Indivisible Items With Externalities. In Williams, B.; Chen, Y.; and Neville, J., eds., *Proceedings of the 37th AAAI Conference on Artificial Intelligence, AAAI '23*, 5472–5480. Washington, DC, USA: AAAI Press.
- Barman, S.; Biswas, A.; Murthy, S. K. K.; and Narahari, Y. 2018. Groupwise Maximin Fair Allocation of Indivisible Goods. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the 32nd AAAI Conference on Artificial Intelligence, AAAI '18*, 917–924. Palo Alto, CA, USA: AAAI Press.
- Błażej, V.; Gupta, S.; Sridharan, R.; and Strulo, P. 2025. Tractable Graph Structures in EFX Orientation. In Lavi, R.; and Zhang, J., eds., *Proceedings of the 18th International Symposium on Algorithmic Game Theory, SAGT '25*, volume 15953 of *Lecture Notes in Computer Science*, 175–190. Cham: Springer.
- Bouveret, S.; and Lang, J. 2008. Efficiency and Envy-Freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. *Journal of Artificial Intelligence Research*, 32: 525–564.
- Bredereck, R.; Kaczmarczyk, A.; Knop, D.; and Niedermeier, R. 2023. High-Multiplicity Fair Allocation Using Parametric Integer Linear Programming. In Gal, K.; Nowé, A.; Nalepa, G. J.; Fairstein, R.; and Radulescu, R., eds., *Proceedings of the 26th European Conference on Artificial Intelligence, ECAI '23*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, 303–310. IOS Press.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2023. Fair Division with Allocator’s Preference. In Garg, J.; Klimm, M.; and Kong, Y., eds., *Proceedings of the 19th International Conference on Web and Internet Economics, WINE '23*, volume 14413 of *Lecture Notes in Computer Science*, 77–94. Cham: 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 Elkind, E., ed., *Proceedings of the 32nd International Joint Conference on Artificial Intelligence, IJCAI '23*, 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*, 9(3): 18.
- Christodoulou, G.; Fiat, A.; Koutsoupias, E.; and Sgouritsa, A. 2023. Fair Allocation in Graphs. In Leyton-Brown, K.; Hartline, J. D.; and Samuelson, L., eds., *Proceedings of the 24th ACM Conference on Economics and Computation, EC '23*, 473–488. New York, NY, USA: ACM.
- Conitzer, V.; Freeman, R.; Shah, N.; and Wortman Vaughan, J. 2019. Group Fairness for the Allocation of Indivisible Goods. In van Hentenryck, P.; and Zhou, Z.-H., eds., *Proceedings of the 33rd AAAI Conference on Artificial Intelligence, AAAI '19*, 1853–1860. Palo Alto, CA, USA: AAAI Press.
- Dechter, R.; and Pearl, J. 1985. Generalized Best-First Search Strategies and the Optimality of A\*. *Journal of the ACM*, 32(3): 505–536.
- Deligkas, A.; Eiben, E.; Goldsmith, T.; and Korchemna, V. 2025a. EF1 and EFX Orientations. In Kwok, J., ed., *Proceedings of the 34rd International Joint Conference on Artificial Intelligence, IJCAI '25*, 56–63. ijcai.org.
- Deligkas, A.; Eiben, E.; Goldsmith, T.-L.; Knop, D.; and Schierreich, Š. 2025b. Dividing Indivisible Items for the Benefit of All: It is Hard to Be Fair Without Social Awareness. arXiv:2511.08160.
- Deligkas, A.; Eiben, E.; Korchemna, V.; and Schierreich, Š. 2024. The Complexity of Fair Division of Indivisible Items With Externalities. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Proceedings of the 38th AAAI Conference on Artificial Intelligence, AAAI '24*, 9653–9661. Washington, DC, USA: AAAI Press.
- Eiben, E.; Ganian, R.; Hamm, T.; and Ordyniak, S. 2023. Parameterized Complexity of Envy-Free Resource Allocation in Social Networks. *Artificial Intelligence*, 315: 103826.
- Flammini, M.; Greco, G.; and Varricchio, G. 2025. Fair Division With Social Impact. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI '25*, 13856–13863. Washington, DC, USA: AAAI Press.
- Foley, D. K. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays*, 7: 45–98.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Herreiner, D. K.; and Puppe, C. D. 2009. Envy Freeness in Experimental Fair Division Problems. *Theory and Decision*, 67: 65–100.
- Hosseini, H. 2024. The Fairness Fair: Bringing Human Perception Into Collective Decision-Making. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Proceedings of the 38th AAAI Conference on Artificial Intelligence, AAAI '24*, 22624–22631. Washington, DC, USA: AAAI Press.
- Hosseini, H.; Kavner, J.; Sikdar, S.; Vaish, R.; and Xia, L. 2025. Epistemic vs. Counterfactual Fairness in Allocation of Resources. In *Proceedings of the 5th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, EAAMO '25*, 93–106. New York, NY, USA: ACM.
- Hosseini, H.; Sikdar, S.; Vaish, R.; Wang, H.; and Xia, L. 2020. Fair Division Through Information Withholding. In Conitzer, V.; and Sha, F., eds., *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI '20*, 2014–2021. Palo Alto, CA, USA: AAAI Press.
- Huang, T.; y. Leung, A. K.; Eom, K.; and Tam, K.-P. 2022. Important to Me and My Society: How Culture Influences the Roles of Personal Values and Perceived Group Values in Environmental Engagements via Collectivistic Orientation. *Journal of Environmental Psychology*, 80: 101774.
- Korf, R. E. 1997. Finding Optimal Solutions to Rubik's Cube Using Pattern Databases. In Kuipers, B.; and Webber, B. L., eds., *Proceedings of the 14th National Conference on Artificial Intelligence, AAAI '97*, 700–705. AAAI Press.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In Breese, J. S.; Feigenbaum, J.; and Seltzer, M. I., eds., *Proceedings of the 5th ACM Conference on Electronic Commerce, EC '04*, 125–131. New York, NY, USA: ACM.
- Nguyen, T. T.; and Rothe, J. 2023. Complexity Results and Exact Algorithms for Fair Division of Indivisible Items: A Survey. In Elkind, E., ed., *Proceedings of the 32nd International Joint Conference on Artificial Intelligence, IJCAI '23*, 6732–6740. ijcai.org.
- Overvold, M. C. 1980. Self-Interest and the Concept of Self-Sacrifice. *Canadian Journal of Philosophy*, 10(1): 105–118.
- 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.
- Thomas, D. W.; and Thomas, M. D. 2025. Behavioral Symmetry With Humanomics: Public Choice and Moral Community. *Public Choice*, 202(3): 419–431.
- Velez, R. A. 2016. Fairness and Externalities. *Theoretical Economics*, 11(1): 381–410.
- Viswanathan, V.; and Zick, Y. 2025. A General Framework for Fair Allocation under Matroid Rank Valuations. *ACM Transactions on Economics and Computation*, 13(3): 14.
- Wu, X.; Zhang, C.; and Zhou, S. 2025. Weighted EF1 Allocations for Indivisible Chores. *Artificial Intelligence*, 347: 104386.
- Zhou, Y.; Wei, T.; Li, M.; and Li, B. 2024. A Complete Landscape of EFX Allocations on Graphs: Goods, Chores and Mixed Manna. In Larson, K., ed., *Proceedings of the 33rd International Joint Conference on Artificial Intelligence, IJCAI '24*, 3049–3056. ijcai.org.
- Zhou, Y.; and Zeng, J. 2015. Massively Parallel A\* Search on a GPU. In Bonet, B.; and Koenig, S., eds., *Proceedings of the 29th AAAI Conference on Artificial Intelligence, AAAI '15*, 1248–1255. Palo Alto, CA, USA: AAAI Press.