



Fair allocation of indivisible goods and chores

Haris Aziz¹ · Ioannis Caragiannis² · Ayumi Igarashi³ · Toby Walsh¹

Accepted: 11 September 2021 / Published online: 9 November 2021
© Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We consider the problem of fairly dividing a set of indivisible items. Much of the fair division literature assumes that the items are “goods” that yield positive utility for the agents. There is also some work in which the items are “chores” that yield negative utility for the agents. In this paper, we consider a more general scenario in which an agent may have positive or negative utility for each item. This framework captures, e.g., fair task assignment, where agents can experience both positive and negative utility for each task. We demonstrate that whereas some of the positive axiomatic and computational results extend to this more general setting, others do not. We present several new and efficient algorithms for finding fair allocations in this general setting. We also point out several gaps in the literature regarding the existence of allocations that satisfy certain fairness and efficiency properties and examine the complexity of computing such allocations.

1 Introduction

Consider a group of students who are assigned to a certain set of coursework tasks. Students may have a subjective view regarding how enjoyable each task is. For some, solving a mathematical problem may be fulfilling and rewarding. For others, it may be nothing but torture. A student who gets more cumbersome chores may be compensated by giving her some valuable goods so that she does not feel hard done by.

This example can be viewed as an instance of a classic fair division problem. Given a set of agents that have preferences for a set of indivisible items, we want to allocate the

✉ Ayumi Igarashi
ayumi_igarashi@nii.ac.jp

Haris Aziz
haris.aziz@unsw.edu.au

Ioannis Caragiannis
iannis@cs.au.dk

Toby Walsh
tw@cse.unsw.edu.au

¹ UNSW Sydney and Data61 CSIRO, Sydney, Australia

² Aarhus University, Aarhus, Denmark

³ National Institute of Informatics, Chiyoda City, Japan

items to the agents as fairly as possible. The twist that we consider is that whether an agent has positive or negative utility for an item is subjective. Our setting is sufficiently general to encapsulate two well-studied settings: (1) “the allocation of goods” in which agents have positive utilities for the items and (2) “the allocation of chores” in which agents have negative utilities for the items. The setting we consider includes a third setting (3) “the allocation of objective goods and chores” in which the items can be partitioned into chores (that yield negative utility for all agents) and goods (that yield positive utility for all agents). Setting (3) covers several scenarios in which an agent could be compensated by some goods for doing some chores.

In this paper, we suggest a very simple yet general model of allocation of indivisible items that properly includes the allocation of goods and chores. In our model, we focus on the relaxations of envy-freeness and proportionality, which are two important fairness concepts. Envy-freeness requires that no agent is envious of another agent. Proportionality requires that each agent derives a minimum amount of utility, which depends on the total number of agents and an agent’s utility for the set of all items. We present some case studies that highlight the fact that whereas some existence and computational results can be extended to our general model, in other cases the combination of good and chore allocation poses interesting challenges not faced in subsettings. Some of our results hold even if we assume that the items are placed along a line and that each agent gets a contiguous bundle. These requirements can capture scenarios such as the allocation of rooms in a corridor among research groups in which a research group may want to be assigned to adjacent rooms. Our central technical contributions are several new, efficient algorithms for finding fair allocations. In particular:

- We formalize fairness concepts for the general setting. Some fairness concepts directly extend from the setting of goods allocation to our setting. Other fairness concepts such as “envy-freeness up to one item” (EF1) and “proportionality up to one item” (PROP1) need to be generalized appropriately.
- We show that the round-robin sequential allocation algorithm that returns an EF1 allocation in the case of goods does not work in general. Nevertheless, we present a careful generalization of the decentralized round-robin algorithm that finds an EF1 allocation when utilities are additive.
- Turning our attention to an efficient and fair allocation, we demonstrate that in the case of two agents, there exists a polynomial-time algorithm that finds an EF1 and Pareto-optimal (PO) allocation for our setting. The algorithm can be viewed as an interesting generalization of the Adjusted Winner rule [22, 23] designed for divisible goods.
- Weakening EF1 to PROP1, we show that there exists an allocation that is not only PROP1 but also contiguous (assuming that the items are placed in a line). We further provide a polynomial-time algorithm that finds such an allocation.

1.1 Related work

The fair allocation of indivisible items is a central problem considered by several fields including computer science and economics [19, 22]. A closely intertwined field that uses the same central fairness concepts is that of cake cutting [22, 42], which can be viewed as the allocation of divisible goods.

There are several established notions of fairness, including envy-freeness [34] and proportionality [44]. Allocations of divisible goods that satisfy these fairness requirements

are known to exist in conjunction with Pareto-optimality [47] or contiguity [44–46, 48]; however, for the case of indivisible goods, neither envy-freeness nor proportionality can always be attained. Consider two agents dividing a single valuable good. Envy-freeness up to one good (EF1) and proportionality up to one good (PROP1) are natural relaxations of envy-freeness and proportionality for the case of indivisible goods to allow for the possibility of the existence of fair allocations. In this paper, we formalize EF1 and PROP1 for our general setting of goods and chores. The EF1 concept was implicit in a paper by Lipton et al. [38]; today, it has become a well-studied fairness concept in its own right [25]. Caragiannis et al. [28] further popularized it, showing that a natural modification of the Nash welfare maximizing rule satisfies EF1 and PO for the case of goods. They refer to their rule as the Maximum Nash Welfare (MNW) rule. Barman et al. [11] presented a pseudo-polynomial-time algorithm for computing an allocation that is PO and EF1 for goods. In general, checking whether there exists an envy-free and Pareto-optimal allocation for goods is Σ_2^P -complete [32]; the problem remains NP-hard even when agents have binary additive utilities [17].

A stronger fairness concept, *envy-freeness up to the least valued good* (EFX), was introduced by Caragiannis et al. [28]. In contrast with the weaker requirement of EF1, the existence question concerning EFX remains elusive even for non-negative additive utility: Chaudhury et al. [29] showed that an EFX allocation exists for the number of agents $n \leq 3$. Intriguingly, the existence question is open for $n \geq 4$.

The recently introduced *maximin share* (MMS) notion is weaker than envy-freeness and proportionality and has been heavily studied in the computer science literature. It can be viewed as a relaxation of proportionality for the case of indivisible items. Kurokawa et al. [37] showed that an MMS allocation of goods may not always exist. On the positive side, there exists a polynomial-time algorithm that returns a $2/3$ -approximate MMS allocation [3, 37]. Subsequent papers have presented simpler [10] or even better [36] approximation algorithms for MMS allocations.

Aziz [4] noted that the work on multi-agent chore allocation is less developed than that of goods, and that the results from one may not necessarily carry over to the other. Aziz et al. [7] considered a fair allocation of indivisible chores and showed that there exists a simple polynomial-time algorithm that returns a 2-approximate MMS allocation for chores. Barman and Krishnamurthy [10] presented a better approximation algorithm. Caragiannis et al. [26] studied the efficiency loss to achieve several fair allocations in the context of both good and chore divisions.

The allocation of a mixture of goods and chores has received recent attention in the context of divisible items [15, 16]. In these papers, the authors have focused on the properties of allocations in competitive equilibrium with equal incomes (CEEIs) when the items are divisible. In particular, Bogomolnaia et al. [15] presented that the allocations in CEEIs still possess the properties of envy-freeness and Pareto-optimality, and further obtained a characterization of CEEIs in terms of Nash welfare, which generalizes the classical result of Eisenberg and Gale [33] for non-negative additive utilities. Garg and McGlaughlin [35] and Chaudhury et al. [30] focused on algorithms for computing a competitive equilibrium for goods and chores. While the precise complexity of the problem for additive utilities is still left open, Garg and McGlaughlin [35] demonstrated that it is polynomial-time solvable when the number of items or agents is a constant; Chaudhury et al. [30] showed that the problem of finding CEEIs under additively separable piecewise linear concave (SPLC) utilities is PPAD-hard even when all items are chores.

Subsequent Work. Our study of a general setting for goods and chores and our formalization of general definitions for fairness concepts (that apply well to hybrid settings) has

spurred further work on the topic. In his survey, Moulin [40] discussed the subtle differences between the treatment of goods and chores. Aleksandrov and Walsh [1] considered variations of the concepts and algorithms that we propose. Aziz and Rey [5] considered a stronger concept of group envy-freeness for goods and chores. Aziz et al. [9] focused on our formulation of PROP1, which is a weakening of proportionality, and proposed a polynomial-time algorithm for computing allocations that are PROP1 and PO.

Bérczi et al. [12] pointed out that the natural extension of the envy-cycle elimination algorithm of Lipton et al. [38] does not provide an EF1 allocation for a combination of goods and chores. Bhaskar et al. [13] provided further insights into the issue and showed that a modification of the envy-cycle elimination algorithm, which restricts the envy-graph to the edges involving maximum envy, results in an EF1 allocation for doubly-monotonic utilities that are more general than additive utilities.

2 Our model and fairness concepts

We now define a fair division problem of indivisible items in which agents may have either positive or negative utility for the items. For a natural number $s \in \mathbb{N}$, we write $[s] = \{1, 2, \dots, s\}$. An *instance* is a triple $I = (N, O, U)$ where

- $N = [n]$ is a set of *agents*,
- $O = \{o_1, o_2, \dots, o_m\}$ is a set of *indivisible items*, and
- U is an n -tuple of utility functions $u_i : 2^O \rightarrow \mathbb{R}$.

Each subset $X \subseteq O$ is referred to as a *bundle* of items. We may abuse the notation and write $u_i(o) = u_i(\{o\})$. We say that an item o is a *good* (respectively, *chore* and *null item*) for agent i if $u_i(o) > 0$ (respectively, $u_i(o) < 0$ and $u_i(o) = 0$). We note that under this model, an item can be a good for agent i but a chore for another agent j . Such an item o is referred to as a subjective good or chore. Formally, item o is referred to as a *subjective good* (respectively, *subjective chore*) if $u_i(o) > 0$ for some $i \in N$ and $u_j(o) \leq 0$ for some $j \in N$ (respectively, $u_i(o) < 0$ for some $i \in N$ and $u_j(o) \geq 0$ for some $j \in N$). Item o is referred to as an *objective good* (respectively, *objective chore*) if $u_i(o) > 0$ for all $i \in N$ (respectively, $u_i(o) < 0$ for all $i \in N$). An item can be both a subjective good and a subjective chore. An objective good cannot be a null item or a subjective chore. An objective chore cannot be a null item or a subjective good.

We assume that the utilities in this paper are *additive*, namely, $u_i(X) = \sum_{o \in X} u_i(o)$ for each bundle $X \subseteq O$, which implies that $u_i(\emptyset) = 0$ for all $i \in N$. We say that agent i *weakly prefers* (respectively, *strictly prefers*) item o to item o' if $u_i(o) \geq u_i(o')$ (respectively, $u_i(o) > u_i(o')$). An *allocation* π is a function $\pi : N \rightarrow 2^O$ such that $\bigcup_{i \in N} \pi(i) = O$, and $\pi(i) \cap \pi(j) = \emptyset$ for every pair of distinct agents $i, j \in N$.

We first observe that the definitions of some fairness concepts can be naturally extended to this general model. The most classical fairness principle is *envy-freeness*, a condition requiring that agents do not envy each other. Specifically, given an allocation π , we say that i *envies* j if $u_i(\pi(i)) < u_i(\pi(j))$. An allocation π is *envy-free* (EF) if no agent envies any other agent. Another appealing notion of fairness is *proportionality*, which guarantees each agent an $1/n$ fraction of his or her utility for the whole set of items. Formally, an allocation π is *proportional* (PROP) if each agent $i \in N$ receives a bundle $\pi(i)$ of utility that is at least

a *proportional fair share* $u_i(O)/n$. The following implication, which is well-known for the case of goods, holds in our setting as well.

Proposition 1 *For additive utilities, an envy-free allocation satisfies proportionality.*

Proof Suppose that an allocation π is an envy-free allocation. Consider any agent $i \in N$. Then, by envy-freeness, $u_i(\pi(i)) \geq u_i(\pi(j))$ for all $j \in N$. Thus, by summing up all of the inequalities, $n \cdot u_i(\pi(i)) \geq \sum_{j \in N} u_i(\pi(j)) = u_i(O)$. Therefore each $i \in N$ receives a bundle of utility at least $u_i(O)/n$, so π satisfies proportionality. \square

A simple example of one good with two agents already suggests the impossibility in achieving envy-freeness and proportionality. The recent literature on indivisible allocation has, therefore, focused on approximations of these fairness concepts. A prominent relaxation of envy-freeness, introduced by Budish [25], is *envy-freeness up to one good* (EF1), which requires that an agent's envy toward another bundle can be eliminated by removing some good from the envied bundle. We will present a generalized definition of EF1 for our setting: the envy can disappear by removing either one "good" from the other's bundle or one "chore" from their own bundle.

Definition 1 (EF1) Given allocation π , we say that i envies j by more than one item if i envies j and $u_i(\pi(i) \setminus \{o\}) < u_i(\pi(j) \setminus \{o\})$ for each item $o \in \pi(i) \cup \pi(j)$. An allocation π is *envy-free up to one item* (EF1) if for all $i, j \in N$, i does not envy j by more than one item.

Obviously, envy-freeness implies EF1. Conitzer et al. [31] introduced a novel relaxation of proportionality that is referred to as *PROPI*. In the context of goods allocation, this fairness relaxation is a weakening of both EF1 and proportionality, requiring that each agent gets his or her proportional fair share if he or she obtains one additional good from the others' bundles. Now, we will extend this definition to our setting: under our definition, each agent receives her proportional fair share by obtaining an additional good or removing some chore from his or her bundle.

Definition 2 (PROPI) Allocation π satisfies *proportionality up to one item* (PROPI) if for each agent $i \in N$,

- $u_i(\pi(i)) \geq u_i(O)/n$; or
- $u_i(\pi(i)) + u_i(o) \geq u_i(O)/n$ for some $o \in O \setminus \pi(i)$; or
- $u_i(\pi(i)) - u_i(o) \geq u_i(O)/n$ for some $o \in \pi(i)$.


We can verify that EF1 implies PROPI.

Proposition 2 *For additive utilities, an EF1 allocation satisfies PROPI.*

Proof The claim clearly holds when $|N| \leq 1$ or $O = \emptyset$; thus suppose $|N| \geq 2$ and $O \neq \emptyset$. Consider any allocation π that satisfies EF1, and any agent $i \in N$.

First, consider the case when $\pi(i) = O$. If $u_i(O) \geq 0$, then it is clear that $u_i(\pi(i)) = u_i(O) \geq u_i(O)/n$. If $u_i(O) < 0$, consider any $j \in N \setminus \{i\} \neq \emptyset$. Since π is EF1, there is an item $o \in \pi(i)$ such that $u_i(\pi(i)) - u_i(o) \geq u_i(\pi(j)) = u_i(\emptyset) = 0$. Thus, $u_i(\pi(i)) - u_i(o) \geq 0 > u_i(O)/n$ for some $o \in \pi(i)$.

Fig. 1 Relationship between fairness concepts



Next, consider the case when $\pi(i) = \emptyset$. If $u_i(O) \leq 0$, then it is clear that $u_i(\pi(i)) = u_i(\emptyset) = 0 \geq u_i(O)/n$. Thus, suppose $u_i(O) > 0$. Then, according to EF1, for every $j \in N \setminus \{i\}$, i does not envy j (i.e., $u_i(\pi(i)) \geq u_i(\pi(j))$), or there is an item $o \in \pi(j)$ such that $u_i(\pi(i)) \geq u_i(\pi(j) \setminus \{o\})$, which means that $u_i(\pi(i)) = u_i(\emptyset) = 0 \geq u_i(\pi(j)) - u_i(o)$. Let $o^* \in O$ be such that $o^* \in \operatorname{argmax}_{o \in O} u_i(o)$. Note that $u_i(o^*) > 0$ since $u_i(O) > 0$. Then, $u_i(\pi(i)) + u_i(o^*) \geq u_i(\pi(j))$ for every $j \in N$, which implies that $u_i(\pi(i)) + u_i(o^*) \geq u_i(O)/n$.

Finally, consider the case when $O \setminus \pi(i) \neq \emptyset$ and $\pi(i) \neq \emptyset$. Let $x = \max_{o \in O \setminus \pi(i)} u_i(o)$ and $y = \min_{o \in \pi(i)} u_i(o)$. Since π satisfies EF1, for any agent $j \in N \setminus \{i\}$,

- i does not envy j ; or
- there exists an item $o \in \pi(j)$ such that $u_i(\pi(i)) \geq u_i(\pi(j)) - u_i(o)$; or
- there exists an item $o \in \pi(i)$ such that $u_i(\pi(i)) - u_i(o) \geq u_i(\pi(j))$,

which implies the following:

- $u_i(\pi(i)) \geq u_i(\pi(j))$; or
- $u_i(\pi(i)) + x \geq u_i(\pi(j))$; or
- $u_i(\pi(i)) - y \geq u_i(\pi(j))$.

Thus, if i gets bonus utility $b^* := \max\{x, -y, 0\}$ by getting some good or eliminating some chore, his or her updated utility is such that $u_i(\pi(i)) + b^* \geq u_i(\pi(j))$ for any agent $j \in N \setminus \{i\}$. This would imply the following:

$$n \cdot (u_i(\pi(i)) + b^*) \geq \sum_{j \in N} u_i(\pi(j)) = u_i(O),$$

which implies that $u_i(\pi(i)) + b^* \geq u_i(O)/n$. Therefore, PROP1 is satisfied. □

Figure 1 illustrates the relationship between the fairness concepts introduced previously.

In addition to fairness, we will also consider an efficiency criterion. The most commonly used efficiency concept is *Pareto-optimality*. Given an allocation π , another allocation π' is a *Pareto-improvement* of π if $u_i(\pi'(i)) \geq u_i(\pi(i))$ for all $i \in N$ and $u_j(\pi'(j)) > u_j(\pi(j))$ for some $j \in N$. We say that an allocation π is *Pareto-optimal* (PO) if there is no allocation that is a Pareto-improvement of π .

Next, we give an example of our problem setting.

Example 1 Consider an instance with four agents $N = [4]$ and item set $O = \{o_1, o_2, \dots, o_9\}$. The utility functions of the agents are represented in Table 1. Consider the following allocation π in which

Table 1 An instance of four agents with subjective goods and chores

	o_1	o_2	o_3	o_4	o_5	o_6	o_7	o_8	o_9
Agent 1:	1	-1	2	1	-2	-4	-6	-1	-1
Agent 2:	4	-3	6	2	-2	-2	-2	-1	-1
Agent 3:	0	11	8	11	0	0	0	10	0
Agent 4:	0	11	8	11	0	0	0	0	10

$$\pi(1) = \{o_2, o_4\},$$

$$\pi(2) = \{o_1, o_3, o_5, o_6, o_7\},$$

$$\pi(3) = \{o_8\} \text{ and}$$

$$\pi(4) = \{o_9\}.$$

The resultant utility of the agents are as follows.

$$u_1(\pi(1)) = u_1(o_2) + u_1(o_4) = 0,$$

$$u_2(\pi(2)) = u_2(o_1) + u_2(o_3) + u_2(o_5) + u_2(o_6) + u_2(o_7) = 4,$$

$$u_3(\pi(3)) = u_3(o_8) = 10, \text{ and}$$

$$u_4(\pi(4)) = u_4(o_9) = 10.$$

Allocation π satisfies proportionality since $u_i(\pi(i)) \geq u_i(O)/n$ for all $i \in \{1, 2, 3, 4\}$. However, π does not satisfy envy-freeness as agents 3 and 4 are envious of agent 1: $u_3(\pi(3)) < u_3(\pi(1))$ and $u_4(\pi(4)) < u_4(\pi(1))$. Allocation π even violates the weaker property of EF1 as both agents 3 and 4 envy agent 1 by more than one item.

Allocation π is not Pareto-optimal since agents 1 and 2 get items from $\{o_5, o_6, o_7\}$, for which they have negative utility. These items can be given to either agent 3 or 4, who has zero utility for them.

3 Finding an EF1 allocation

In this section, we focus on EF1, a very permissive fairness concept that admits a polynomial-time algorithm in the case of goods allocation. For instance, consider the *round-robin rule* in which agents take turns and choose the most preferred unallocated item. The round-robin rule finds an EF1 allocation if all of the items are goods (see e.g., [28]). The round-robin rule falls under a general class of allocation rules in which there is a picking sequence of agents and agents pick the most preferred available item on their turn [18]. By a very similar argument, it can be shown that the algorithm also finds an EF1 allocation if all the items are chores. However, we will show that the round-robin rule fails to find an EF1 allocation if we have some items that are objective goods and others that are objective chores.

Proposition 3 *The round-robin rule does not satisfy EF1.*

Proof Suppose that there are two agents and four items with identical utilities described in Table 2.

Consider an ordering in which Alice first chooses the only good, and then the remaining chores of equal value are allocated accordingly. In that case, Alice gets the positively

Table 2 An instance in which the round-robin rule fails to satisfy EF1

	①	②	③	④
Alice, Bob:	2	- 3	- 3	- 3




Fig. 2 Illustration of the double round-robin algorithm. The dotted line corresponds to the picking order when allocating the items for which each agent has a non-positive utility. The thick line corresponds to the picking order when allocating the items for which some agent has a positive utility. The solid black circle indicates the agent who starts the picking. For the dotted round, agent 1 is the first agent to pick. For the solid round, agent n is the first agent to pick

valued good and one chore, whereas Bob gets two chores. So, even if one item is removed from the bundles of Alice and Bob, Bob will still remain envious. □

Nevertheless, a careful adaptation of the round-robin method to our setting, which we call the *double round-robin algorithm*, constructs an EF1 allocation. In essence, the algorithm will apply the round-robin method twice: clockwise and anticlockwise. In the first phase, the round-robin algorithm allocates the items for which each agent has a non-positive utility, while in the second phase, the reversed round-robin algorithm allocates to agents the remaining items for which some agent has a positive utility, in the opposite order starting with the agent who chooses last in the first phase. Intuitively each agent i may envy agent j who comes earlier than her at the end of one phase, but i does not envy j with respect to the items allocated in the other round; therefore, the envy of i toward j can be bounded up to one item. We present a formal description of the algorithm in Algorithm 1; see Fig. 2 for an illustration. In the algorithm description, when we use picking sequence $(1, 2, \dots, n)^*$, we mean that the picking sequence $1, 2, \dots, n$ repeats.

Algorithm 1 Double Round-Robin Algorithm

Input: An instance $I = (N, O, U)$.

Output: An allocation π .

- 1: Initialize $\pi(i) = \emptyset$ for each agent $i \in N$.
 - 2: Partition O into $O^+ = \{o \in O \mid \exists i \in N \text{ s.t. } u_i(o) > 0\}$, $O^- = \{o \in O \mid \forall i \in N, u_i(o) \leq 0\}$ and suppose $|O^-| = an - k$ for some positive integer a and $k \in \{0, \dots, n - 1\}$.
 - 3: Create k dummy null items for which each agent has utility 0, and add them to O^- (hence, $|O^-| = an$).
 - 4: Let the agents come in a round-robin sequence $(1, 2, \dots, n)^*$ and pick their most preferred item in O^- until all items in O^- are allocated.
 - 5: Let the agents come in a round-robin sequence $(n, n - 1, \dots, 1)^*$ and pick their most preferred item in O^+ until all items in O^+ are allocated. If an agent has no available item which gives her strictly positive utility, she does not get a real item but pretends to *pick* a dummy one for which she has utility 0.
 - 6: Remove the dummy items from the current allocation π and return the resulting allocation π^* .
-

In the following, for an allocation π and a bundle X , we say that i envies j with respect to X if $u_i(\pi(i) \cap X) < u_i(\pi(j) \cap X)$.

Theorem 1 For additive utility, the double round-robin algorithm returns an EF1 allocation in $O(\max\{m \log m, mn\})$ time.

Proof We note that the algorithm ensures that all agents receive the same number of chores, by introducing k dummy chores. Now let π be the output of Algorithm 1. To see that π satisfies EF1, consider any pair of two agents i and j where $i < j$. We will show that by removing one item from either i 's bundle or j 's bundle, these agents will not envy each other. We denote c_t^i and c_t^j as the t -th items allocated to agent i and agent j for $t = 1, 2, \dots, a$ in Line 4, respectively. We denote g_t^i and g_t^j as the t -th items allocated to agent i and agent j for $t = 1, 2, \dots, b$ in Line 5, respectively, where b denotes the number of rounds in which each agent chooses an item (including a dummy item) in Line 5.

First, consider i 's envy for j . We first observe that the t -th item c_t^i in O^- allocated to i is weakly preferred by i to the t -th item c_t^j in O^- allocated to j . Therefore, agent i does not envy j with respect to O^- . Namely,

$$u_i(\pi(i) \cap O^-) = \sum_{t=1}^a u_i(c_t^i) \geq \sum_{t=1}^a u_i(c_t^j) = u_i(\pi(j) \cap O^-). \tag{1}$$

As for the allocation of the items in O^+ , agent i may envy agent j with respect to O^+ . But if the first item, g_1^j , picked by j from O^+ is removed from j 's bundle, then the envy will disappear, that is, i will not envy j with respect to $O^+ \setminus \{g_1^j\}$. Namely,

$$u_i(\pi(i) \cap O^+) = \sum_{t=1}^b u_i(g_t^i) \geq \sum_{t=2}^b u_i(g_t^j) = u_i((\pi(j) \cap O^+) \setminus \{g_1^j\}). \tag{2}$$

This is because for each item g_t^j picked by j where $t = 2, 3, \dots, b$, there is a corresponding item, g_{t-1}^i , picked by i before j 's turn that is weakly as preferred by i to g_t^j . Combining (1) and (2) yields $u_i(\pi(i)) \geq u_i(\pi(j) \setminus \{g_1^j\})$.

Second, consider j 's envy for i . Similar to the preceding scenario, agent j does not envy agent i with respect to O^+ because agent j takes the first pick among i and j ; that is, for every item g_i^i chosen by i , agent j picks an item g_j^j before i that he or she weakly prefers to g_i^i . Thus,

$$u_j(\pi(j) \cap O^+) = \sum_{t=1}^b u_j(g_t^j) \geq \sum_{t=1}^b u_i(g_t^i) = u_j(\pi(i) \cap O^+). \tag{3}$$

As for the items in O^- , for each item c_t^i picked by i where $t = 2, 3, \dots, a$, there is an item c_{t-1}^j picked by j before i that j weakly prefers to c_t^i , which implies that j does not envy i with respect to $O^- \setminus \{c_a^j\}$. Thus

$$u_j((\pi(j) \cap O^-) \setminus \{c_a^j\}) = \sum_{t=1}^{a-1} u_j(c_t^j) \geq \sum_{t=2}^a u_j(c_t^i) \geq \sum_{t=1}^a u_j(c_t^i) = u_j(\pi(i)). \tag{4}$$

Note that the last inequality holds since $u_j(c_1^i) \leq 0$. Combining (3) and (4) yields $u_j(\pi(j) \setminus \{c_a^j\}) \geq u_j(\pi(i))$.

In either case, agents do not envy each other by more than one item. We conclude that π is EF1 and so is the final allocation π^* as removing dummy items does not affect the utility of each agent.

It remains to analyze the running time of Algorithm 1. Line 2 requires $O(mn)$ time, as each item needs to be examined by every agent. Lines 4 and 5 require $O(m \log m)$ time, as there are at most m iterations, and for each iteration, each agent has to choose the most preferred item out of at most m items, which can be done by sorting all of the items according to the preference of each agent at the beginning. Thus, the total running time can be bounded by $O(\max\{m \log m, mn\})$, which completes the proof. \square

4 Finding a PO and EF1 allocation

We now move on to the next question: whether fairness is achievable along with efficiency. In the context of goods allocation in which agents have non-negative additive utility, Caragiannis et al. [28] proved that an outcome that maximizes the *Nash welfare* (i.e., the product of the utilities) satisfies EF1 and Pareto-optimality. The question regarding whether a PO and EF1 allocation exists for chores is unresolved. Starting from an EF1 allocation and finding Pareto-improvements, one encounters two challenges. First, Pareto-improvements may not necessarily preserve EF1; second, finding Pareto-improvements is NP-hard [6, 32]. Even if we ignore the second challenge, the question regarding the existence of a PO and EF1 allocation for chores is open. Next, we show that the problem of finding a PO and EF1 allocation is completely resolved for the restricted but important case of two agents. Note that the problem of distributing a resource between two agents, which arises in a number of applications, such as a divorce settlement and land division, has been regarded as fundamental in the fair division literature. Indeed, there are several prominent works that study the two-agent case, ranging from classical [21, 24] to more recent ones [2, 41].

The main theorem in this section is stated as follows.

Theorem 2 *For two agents with additive utility, a PO and EF1 allocation always exists and can be computed in $O(m^2)$ time.*

Table 3 An instance of two agents with objective goods and chores

	①	②	③	④	⑤	⑥	⑦
Alice (winner):	1	-1	2	1	-2	-4	-6
Bob (loser):	4	-3	6	2	-2	-2	-2
$ u_\ell(o) / u_w(o) $	4	3	3	2	1	1/2	1/3

Our algorithm, which we present formally in Algorithm 2, can be viewed as a discrete version of the well-known Adjusted Winner (AW) rule [22, 23]. Just like the AW rule, our algorithm finds a PO and EF1 allocation. In contrast to AW, which is designed for goods, our algorithm can handle both goods and chores.

Without loss of generality, assume that there is no item for which each agent has utility 0. The algorithm begins by giving each subjective chore to the agent who considers it as a good or a null item; similarly, it gives each subjective good to the agent who considers it as a good. So, in the following, we assume that we have objective items only, that is, for each item $o \in O$, either o is a good ($u_i(o) > 0$ for each $i \in N$); or o is a chore ($u_i(o) < 0$ for each $i \in N$). Now we call one of the two agents the *winner* (denoted by w) and another the *loser* (denoted by ℓ).

1. Initially, all goods are allocated to the winner and all chores to the loser.
2. We sort the items in terms of $|u_\ell(o)|/|u_w(o)|$ (monotone non-increasing order) from left to right, and consider reallocation of the items according to the ordering (from the left-most to the right-most item).
3. When considering a good, we move it from the winner to the loser. When considering a chore, we move it from the loser to the winner. We stop when the loser does not envy the winner by more than one item.

The example below illustrates our discrete adaptation of AW.

Example 2 (Example of the generalized AW) Consider two agents, Alice and Bob, and five items with the additive utility represented in Table 3 in which the gray circles correspond to goods and the white circles correspond to chores. The generalized AW initially allocates the goods to Alice and the chores to Bob. Then, it transfers the first good from Alice to Bob and moves the second chore from Bob to Alice. After moving the third good from Alice to Bob, Bob stops being envious (by more than one item). Therefore, the final allocation gives items 2 and 4 to Alice and the rest to Bob.

We will first prove that at any point in the algorithm, the allocation π is Pareto-optimal, and so is the final allocation.

Algorithm 2 Generalized Adjusted Winner Algorithm

Input: An instance $I = (N, O, U)$ where $N = \{w, \ell\}$.

Output: An allocation π .

- 1: Initialize $\pi(i) = \emptyset$ for each agent $i \in N$.
 - 2: Let $O_w^* = \{o \in O \mid u_w(o) \geq 0 \text{ and } u_\ell(o) \leq 0\}$ and $O_\ell^* = \{o \in O \mid u_\ell(o) \geq 0 \text{ and } u_w(o) < 0\}$.
 - 3: Let $O^+ = \{o \in O \mid u_i(o) > 0 \ \forall i \in N\}$ and $O^- = \{o \in O \mid u_i(o) < 0 \ \forall i \in N\}$.
 - 4: For each item $o \in O^+ \cup O_w^*$, allocate o to agent w . For each item $o \in O^- \cup O_\ell^*$, allocate o to agent ℓ .
 - 5: Sort the items in $O^+ \cup O^- = \{o_1, o_2, \dots, o_r\}$ where $|u_\ell(o_1)|/|u_w(o_1)| \geq |u_\ell(o_2)|/|u_w(o_2)| \geq \dots \geq |u_\ell(o_r)|/|u_w(o_r)|$.
 - 6: Set $t = 1$.
 - 7: **while** agent ℓ envies agent w by more than one item **do**
 - 8: **if** $o_t \in O^+$ **then**
 - 9: Set $\pi(w) = \pi(w) \setminus \{o_t\}$ and $\pi(\ell) = \pi(\ell) \cup \{o_t\}$.
 - 10: **else if** $o_t \in O^-$ **then**
 - 11: Set $\pi(w) = \pi(w) \cup \{o_t\}$ and $\pi(\ell) = \pi(\ell) \setminus \{o_t\}$.
 - 12: **end if**
 - 13: Update $t = t + 1$.
 - 14: **end while**
-

Lemma 1 *During the execution of Algorithm 2, the allocation π is Pareto-optimal at any point after Line 4.*

Proof It can be easily verified that the allocation π just after Line 4 is Pareto-optimal. Thus, consider some time step after the algorithm enters the **while**-loop of Line 8. Assume the contradiction that π' is a Pareto-improvement of π . We assume without loss of generality that all items in O_w^* remain assigned to w under π' because transferring an item in O_w^* from w to ℓ improves neither the utility of w nor that of ℓ . Likewise, we assume that all items in O_ℓ^* remain assigned to ℓ under π' .

In the following, we call each item $o \in O^+$ a *good* and item $o \in O^-$ a *chore*. For each $i, j \in \{w, \ell\}$ with $i \neq j$, let

- G_{ii} be the set of goods in $\pi(i) \cap \pi'(i)$;
- C_{ii} be the set of chores in $\pi(i) \cap \pi'(i)$;
- G_{ij} be the set of goods in $\pi(i) \cap \pi'(j)$; and
- C_{ij} be the set of chores in $\pi(i) \cap \pi'(j)$.

Consider first the case when in π , the winner has a utility that is at least as high as in π' , while the loser is strictly better off. Taking into account the fact that the bundles of goods G_{ww} and $G_{\ell\ell}$ and the bundles of chores C_{ww} and $C_{\ell\ell}$ are allocated to the same agent in both allocations, this means the following:

$$u_w(G_{\ell w}) + u_w(C_{\ell w}) - u_w(G_{w\ell}) - u_w(C_{w\ell}) \geq 0; \text{ and} \tag{5}$$

$$u_\ell(G_{w\ell}) + u_\ell(C_{w\ell}) - u_\ell(G_{\ell w}) - u_\ell(C_{\ell w}) > 0 \tag{6}$$

The crucial observation now is that the algorithm considered all items in $G_{\ell w}$ and $C_{w\ell}$ before the items in $G_{w\ell}$ and $C_{\ell w}$ in the ordering. Indeed, recall that all of the goods are initially assigned to the winner, $G_{\ell w} \subseteq \pi(\ell)$, and $G_{w\ell} \subseteq \pi(w)$. Thus the goods in $G_{\ell w}$ are those transferred from the winner w to the loser ℓ in the **while**-loop of Line 8, while the goods in $G_{w\ell}$ are those that stay in the winner's bundle. Similarly, recall that all of the chores are initially assigned to the loser, $C_{w\ell} \subseteq \pi(w)$, and $C_{\ell w} \subseteq \pi(\ell)$. Thus, the chores in $C_{w\ell}$ are those transferred from the loser ℓ to the winner w , while the chores in $C_{\ell w}$ are those that stay in the loser's bundle. Now, let α be such that

$$\max_{o \in G_{w\ell} \cup C_{\ell w}} |u_\ell(o)|/|u_w(o)| \leq \alpha \leq \min_{o \in G_{\ell w} \cup C_{w\ell}} |u_\ell(o)|/|u_w(o)|.$$

This definition implies the inequalities,

$$\begin{aligned} u_\ell(G_{w\ell}) &\leq \alpha u_w(G_{w\ell}); u_\ell(G_{\ell w}) \geq \alpha u_w(G_{\ell w}); \\ -u_\ell(C_{w\ell}) &\geq -\alpha u_w(C_{w\ell}); -u_\ell(C_{\ell w}) \leq -\alpha u_w(C_{\ell w}), \end{aligned}$$

which, together with inequality (6), yield

$$\begin{aligned} 0 &< u_\ell(G_{w\ell}) + u_\ell(C_{w\ell}) - u_\ell(G_{\ell w}) - u_\ell(C_{\ell w}) \\ &\leq -\alpha(u_w(G_{\ell w}) + u_w(C_{\ell w}) - u_w(G_{w\ell}) - u_w(C_{w\ell})) \leq 0, \end{aligned}$$

a contradiction. The last inequality follows by (5) and by the fact that α is non-negative. A similar argument applies when in π , the loser has a utility that is at least as high as in π' , while the winner is strictly better off. \square

We are now ready to prove Theorem 2.

Proof (of Theorem 2) We will prove that the final output π of Algorithm 2 satisfies EF1. Together with Lemma 1, this proves the desired claim.

Now, observe that at the final allocation π , one agent at most envies the other: if the loser still envies the winner and the winner also envies the loser, then exchanging the bundles would result in a Pareto-improvement, contradicting Lemma 1. Thus, if the loser envies the winner at π , the winner does not envy the loser, which implies that π is EF1 according to the termination condition.

Consider when at π , the loser does not envy the winner but the winner envies the loser. Let π' be the previous allocation just before the final transfer in the **while**-loop of Line 8. Let $W = \pi'(w) \cap \pi(w)$ and $L = \pi'(\ell) \cap \pi(\ell)$. Namely, W (respectively, L) is the set of items in the winner's bundle (respectively, the loser's bundle) excluding the transferred item at π' and π . By construction, the loser envies the winner by more than one item at π' , which implies $u_\ell(L) < u_\ell(W)$. Suppose towards a contradiction that the winner envies the loser by more than one item at π , which implies $u_w(W) < u_w(L)$.

- If g is the last good that has been moved from the winner to the loser, then allocating W to ℓ and $L \cup \{g\}$ to w would be a Pareto-improvement of π' , a contradiction.
- If c is the last chore that has been moved from the loser to the winner, then allocating $W \cup \{c\}$ to ℓ and L to w would be a Pareto-improvement of π' , a contradiction.

Therefore, the winner does not envy the loser by more than one item at π , and we conclude that π is EF1.

It remains to analyze the running time of the algorithm. First, the items can be sorted in $O(m \log m)$ time. The adjustment process takes $O(m^2)$ time. Each iteration checks if the allocation is EF1 from the perspective of the loser, which requires at most m comparisons of utilities, and there are at most m iterations. Thus, the number of operations is bounded by $O(m^2)$.

A natural question is whether PO and EF1 allocations exist for three or more agents; we leave this as an interesting open question. We remark that Pareto-optimality by itself is easy to achieve in $O(nm)$ time. It suffices to give each item to the agent who values it most.

5 Finding a connected PROP1 allocation

We saw that there always exists an EF1 allocation for subjective goods and chores. If we weaken EF1 to PROP1, one can achieve one additional requirement aside from fairness, *connectivity*. In this section, we will consider a situation when items are placed on a path, and each agent desires a connected bundle of the path. Finding a connected set of items is important in many scenarios. For example, the items can be a set of rooms in a corridor and the agents could be research groups wherein each research group wants to obtain adjacent rooms (see e.g., [14, 20]).

We will show that a connected PROP1 allocation exists, and that it can be found efficiently. In what follows, we assume that the path is given by a sequence of items (o_1, o_2, \dots, o_m) . Formally, we say that an allocation π is *connected* if for each agent $i \in N$, $\pi(i)$ is connected in the path (o_1, o_2, \dots, o_m) . We will consider a slightly more stringent notion of PROP1: a connected allocation π is PROP1_{outer} if for each agent $i \in N$,

- agent i receives a bundle of utility at least her proportional fair share, i.e., $u_i(\pi(i)) \geq u_i(O)/n$, or
- $u_i(\pi(i)) + u_i(o) \geq u_i(O)/n$ for some item $o \in O \setminus \pi(i)$ such that $\pi(i) \cup \{o\}$ is connected; or
- $u_i(\pi(i)) - u_i(o) \geq u_i(O)/n$ for some $o \in \pi(i)$ such that $\pi(i) \setminus \{o\}$ is connected.

We first prove the result for a case of the cake-cutting setting [22, 42] that is of independent interest. In the following, a *mixed cake* is the interval $[0, m]$. Each agent $i \in N$ has a value density function \hat{u}_i , which maps a subinterval of the cake to a real value, where i has uniform utility $u_i(o_j)$ for the interval $[j - 1, j]$ for each $j \in [m]$. The *proportional fair share* of agent i for a mixed cake $[0, m]$ is given by $\hat{u}_i([0, m])/n$. A *contiguous allocation* of a mixed cake assigns each agent a disjoint sub-interval of the cake in which the union of the intervals equals the entire cake $[0, m]$; it satisfies *proportionality* if each agent i gets an interval of utility that is at least his or her proportional fair share.

Theorem 3 *For additive utilities, a contiguous proportional allocation of a mixed cake exists and can be computed in polynomial time.*

Proof Let N^+ be the set of agents with strictly positive total utility for O .

We combine moving-knife algorithms for goods and chores as follows.¹ First, if there is an agent who has positive proportional fair share ($N^+ \neq \emptyset$), we apply the moving-knife algorithm only to the agents in N^+ . Our algorithm moves a knife from *left* to *right*, and agents shout whenever the left part of the cake has a utility of exactly equal to the proportional fair share. The first agent who shouts is allocated the left bundle, and the algorithm recurs on the remaining instance. Second, if no agent has a positive proportional fair share, our algorithm moves a knife from *right* to *left*, and agents shout whenever the left part of the cake has utility exactly proportional fair share. Again, the first agent who shouts is allocated the left bundle, and the algorithm recurs on the remaining instance. Below, for an allocation π and agent set $N' \subseteq N$, we define $\pi|_{N'} : N' \rightarrow 2^O$ to be the restriction of π to N' , i.e., $\pi|_{N'}(i) = \pi(i)$ for each agent $i \in N'$.

Algorithm 3 Generalized Moving-knife Algorithm \mathcal{A}

Input: A sub-interval $[\ell, r]$, agent set N' , utility functions \hat{u}_i for each $i \in N'$.

Output: An allocation $\hat{\pi}$ of a mixed cake $[\ell, r]$ to N' .

```

1: Initialize  $\hat{\pi}(i) = \emptyset$  for each  $i \in N'$ .
2: Set  $N^+ = \{i \in N' \mid \hat{u}_i([\ell, r]) > 0\}$ .
3: if  $N^+ \neq \emptyset$  then
4:   if  $|N^+| = 1$  then
5:     Allocate  $[\ell, r]$  to the unique agent in  $N^+$ .
6:   else
7:     Let  $x_i$  be the minimum point where  $\hat{u}_i([\ell, x_i]) = \hat{u}_i([\ell, r])/|N^+|$  for  $i \in N^+$ .
8:     Find agent  $j$  with minimum  $x_j$  among all agents in  $N^+$ .
9:     return  $\hat{\pi}$  where  $\hat{\pi}(j) = [\ell, x_j]$  and  $\hat{\pi}|_{N' \setminus \{j\}} = \mathcal{A}([x_j, r], N' \setminus \{j\}, (\hat{u}_i)_{i \in N' \setminus \{j\}})$ 
10:  end if
11: else
12:   Let  $x_i$  be the maximum point where  $\hat{u}_i([\ell, x_i]) = \hat{u}_i([\ell, r])/n$  for  $i \in N'$ .
13:   Find agent  $j$  with maximum  $x_j$  among all agents in  $N'$ .
14:   return  $\hat{\pi}$  where  $\hat{\pi}(j) = [\ell, x_j]$  and  $\hat{\pi}|_{N' \setminus \{j\}} = \mathcal{A}([x_j, r], N' \setminus \{j\}, (\hat{u}_i)_{i \in N' \setminus \{j\}})$ 
15: end if

```

Algorithm 3 formalizes the idea. To prove its correctness, we will prove by induction on the number of agents $|N'|$ that the allocation of a mixed cake $[\ell, r]$ produced by \mathcal{A} satisfies the following conditions:

- if $N^+ \neq \emptyset$, then each agent in N^+ receives an interval of utility that is at least his or her proportional fair share $\hat{u}_i([\ell, r])/|N^+|$ and each agent not in N^+ receives an empty piece; and
- if $N^+ = \emptyset$, then each agent receives an interval of utility at least her proportional fair share $\hat{u}_i([\ell, r])/|N'|$.

The claim is clearly true when $|N'| = 1$. Suppose that \mathcal{A} returns a proportional allocation of a mixed cake with desired properties when $|N'| = k - 1$; we will prove it for $|N'| = k$.

Suppose that some agent has a positive proportional fair share, i.e., $N^+ \neq \emptyset$. Note that each agent i not in N^+ has a non-positive proportional fair share and gets nothing; thus, it suffices to show that the agents in N^+ receive bundles of utility that is at least his or her

¹ A moving-knife procedure for computing a proportional allocation for a cake is known as the Dubins-Spanier moving-knife procedure [22].

proportional fair share. If $|N^+| = 1$, the claim is trivial; thus, we assume otherwise. Clearly, agent j receives an interval of utility that is at least his or her proportional fair share. Further, all other agents in N^+ have utility at most their proportional fair shares for the left piece $[\ell, x_j]$. Indeed, if there is an agent $i' \in N^+$ whose utility for the left piece $[\ell, x_j]$ is greater than his or her proportional fair share $\hat{u}_{i'}([\ell, r])/k$, then i' would have shouted when the knife reaches before x_j by the continuity of $\hat{u}_{i'}$, i.e., $x_{i'} < x_j$, contradicting the minimality of x_j . Thus, the remaining agents in N^+ have at least $(k-1) \cdot \hat{u}_i([\ell, r])/k$ utility for the rest of the cake $[x_j, r]$; therefore, by the induction hypothesis each agent in N^+ gets an interval of utility that is at least his or her proportional fair share, and each of the remaining agents gets an empty piece.

Suppose that no agent has a positive proportional fair share. Again, if there is an agent i' whose utility for the left piece $[\ell, x_j]$ is greater than his or her proportional fair share $\hat{u}_{i'}([\ell, r])/k$, then i' would have shouted when the knife reaches before x_j by the continuity of $\hat{u}_{i'}$, i.e., $x_{i'} > x_j$, contradicting the maximality of x_j . Thus, all of the remaining agents have utility of at least $(k-1) \cdot \hat{u}_i([\ell, r])/k$ for the rest of the cake $[x_j, r]$, and therefore, according to the induction hypothesis, each agent gets an interval of utility that is at least his or her proportional fair share $\hat{u}_i([\ell, r])/k$. It can be easily verified that Algorithm 3 runs in polynomial time. \square

The theorem stated above also applies to a general cake-cutting model in which information about the agent's utility function over an interval can be inferred by a series of queries. We note that in contrast with proportionality, the existence of a contiguous envy-free allocation of a mixed cake remains elusive: it is known to exist only when the number n of agents is four or a prime number [39, 43]. Next, we show how a fractional proportional allocation (an allocation that achieves proportionality but treats the items as divisible) can be used to achieve a contiguous PROP1 division of indivisible items.

Theorem 4 *For additive utilities, a connected PROP1_{outer} allocation of a path always exists and can be computed in polynomial time.*

Proof Given a path (o_1, o_2, \dots, o_m) , consider a mixed cake $[0, m]$ and each agent with a value density function \hat{u}_i , where i has uniform utility $u_i(o_j)$ for the interval $[j-1, j]$ for each $j \in [m]$. We know that this instance admits a contiguous and proportional allocation $\hat{\pi}$ from Theorem 3. Suppose without loss of generality that under such an allocation $\hat{\pi}$, agents $1, 2, \dots, n$ receive the 1st, 2nd, ..., and n -th bundles from left to right. That is, each agent $i = 1, 2, \dots, n$ receives the sub-interval $[x_{i-1}, x_i]$ of the mixed cake, where $0 = x_0 \leq x_1 \leq \dots \leq x_{n-1} \leq x_n = m$. Without loss of generality, we also assume that no agent gets the empty bundle under this fractional allocation, i.e., $x_{i-1} < x_i$ for each $i = 1, 2, \dots, n$.

From left to right, we show how to allocate each item o_j for $j = 1, 2, \dots, m$ to construct an integral allocation π . If item o_j is fully contained in some agent's bundle, namely, $x_{i-1} \leq j-1 \leq j \leq x_i$ for some $i \in N$, then we assign each item o_j to agent i . If not (i.e., the item o_j is on the boundary), we allocate it according to the left-most/right-most agents' preference. Formally, suppose that $j-1 \leq x_\ell \leq x_{\ell+1} \leq \dots \leq x_r \leq j$ such that $x_\ell = \min\{x_i \mid x_i \geq j-1\}$ and $x_r = \max\{x_i \mid x_i \leq j\}$. Then we do the following:

1. If two agents ℓ and r disagree on the sign of o_j , i.e., $\min\{u_\ell(o_j), u_r(o_j)\} < 0 < \max\{u_\ell(o_j), u_r(o_j)\}$, then we give the item o_j to the agent $i \in \{\ell, r\}$ who has a positive utility for it.
2. If two agents ℓ and r agree on the sign of o_j , i.e., $\min\{u_\ell(o_j), u_r(o_j)\} \geq 0$ or $\max\{u_\ell(o_j), u_r(o_j)\} < 0$, then we allocate the item o_j in such a way that:
 - the left-agent ℓ takes o_j if both agents have non-negative utility (i.e., $\min\{u_\ell(o_j), u_r(o_j)\} \geq 0$);
 - the right-agent r takes o_j if both agents have negative utility (i.e., $\max\{u_\ell(o_j), u_r(o_j)\} < 0$).

Note that if under the proportional fractional division, agent i gets a fraction of one item only and there are two other agents $i - 1$ and $i + 1$ on the left and right who get a fraction of the same item, agent i gets nothing under our final allocation.

The resulting integral allocation π is PROPI_{outer} . To see this, take any agent i . Clearly, when one of the knife positions x_{i-1} and x_i is integral, the bundle satisfies PROP_{outer} . Further, if $[x_{i-1}, x_i] \subseteq [j - 1, j]$ for some $j \in [m]$, agent i gets utility $1/n$ by receiving either the item o_j or the empty bundle. Thus, assume otherwise, that is, $x_{i-1}, x_i \notin \{0, 1, \dots, m\}$ and $|x_i - x_{i-1}| > 1$. We will show that such an agent gets utility $1/n$ by either receiving the item next to its bundle or deleting the left-most item of her bundle. Let o_r and o_ℓ be the left and right boundary items where $x_{i-1} \in (r - 1, r)$ and $x_i \in (\ell - 1, \ell)$. Note that we have $\{o_{\ell+1}, o_{\ell+2}, \dots, o_{r-1}\} \subseteq \pi(i)$. Consider the following four cases.

- Both o_ℓ and o_r are goods or null items for i , i.e., $\min\{u_i(o_\ell), u_i(o_r)\} \geq 0$. In this case, agent i receives at least o_r . Thus, if $o_\ell \in \pi(i)$, agent i obtains utility $1/n$. If not, agent i gets utility $1/n$ by receiving the item o_ℓ .
- Both o_ℓ and o_r are chores for i , i.e., $\max\{u_i(o_\ell), u_i(o_r)\} < 0$. In this case, agent i does not receive o_r . Thus, if $o_\ell \notin \pi(i)$, agent i obtains utility $1/n$. If not, agent i gets utility $1/n$ by removing the item o_ℓ .
- The item o_ℓ is a good or a null item but o_r is a chore for i , i.e., $u_i(o_\ell) \geq 0$ and $u_i(o_r) < 0$. In this case, agent i does not receive o_r . Thus, if $o_\ell \in \pi(i)$, agent i obtains utility $1/n$. If not, agent i gets utility $1/n$ by receiving the item o_ℓ .
- The item o_ℓ is a chore but o_r is a good or a null item for i , i.e., $u_i(o_\ell) < 0$ and $u_i(o_r) \geq 0$. In this case, agent i receives at least o_r . Thus, if $o_\ell \notin \pi(i)$, agent i obtains utility $1/n$. If not, agent i gets utility $1/n$ by removing the item o_ℓ .

We conclude that π is a connected PROPI_{outer} allocation. By Theorem 3, it is immediate to see that one can compute a connected PROPI_{outer} allocation in polynomial time. \square

6 Discussion

In this paper, we have formally analyzed fair allocation when the indivisible items are a combination of subjective goods and chores. Our work paves the way for a detailed examination of the allocation of goods/chores, and opens up an interesting line of research with many problems left open for future exploration. We conclude with several directions for future research.

EF1 allocations for general utility functions. In Section 3, we have shown that for general additive utilities, an EF1 allocation exists and can be computed efficiently by

Table 4 An instance of two agents and four objective chores

	o_1	o_2	o_3	o_4
Agent 1:	- 1	- 100	- 100	- 100
Agent 2:	- 1	- 2	- 2	- 2

using the double round-robin procedure. The most intriguing open question as a result of our study may be the existence of an EF1 allocation under arbitrary non-monotonic utilities; the most general result obtained thus far is that of Bhaskar et al. [13], who showed that an EF1 allocation exists for the so-called doubly monotonic utilities.

PO and EF1 allocations. Another open question is the existence of PO and EF1 allocations in our setting. While our work establishes the existence of such allocations for two agents with additive utilities, the problem for an arbitrary number of agents is open, even for the restricted setting in which the items are chores. Here, we discuss some unsuccessful approaches while trying to prove that an EF1 and PO outcome exists for chores.

- Recall that the MNW solution of Caragiannis et al. [27] gives an PO and EF1 outcome for goods. But naively maximizing the product of the utilities appears futile because the objective is positive or negative depending on the parity of the number of agents.
- Another potential approach for chores is to minimize the product of the *disutilities*; if the product is zero, we can identify a largest set of agents for which the product is positive. We can then apply the solution to this set of agents. Such an approach does not give an EF1 guarantee. To see this, consider the case of two agents and four identical chores of utility -1 . In that case, the outcome that allocates at least one item to each agent and minimizes the product of the disutilities is one in which one agent gets one chore and the other gets three chores. The outcome does not satisfy EF1 despite the existence of a balanced outcome allocating an equal number of chores to each agent, which is arguably fair by all reasonable measures.
- An approach that works well for fairness in the case of *divisible* chores is the rule that among all PO allocations, *maximizes* the product of the disutilities of those agents who derive a non-zero disutility [15, 16]. However, there exists a simple example with two agents such that maximizing the product of the disutilities subject to PO does not provide an EF1 guarantee. Consider the example in Table 4 that was shared by Hervé Moulin, Ariel Procaccia, and Nisarg Shah. Consider allocation π such that $\pi_1 = \{o_2, o_3, o_4\}$ and $\pi_2 = \{o_1\}$. Allocation π maximizes the product of the disutilities and is PO. However, it is not EF1.

Round-robin share (RRS) and PO allocations. We also note that while the relationship between Pareto-optimality and most fairness notions is still unclear, Conitzer et al. [31] proposed a fairness concept called *round-robin share* (RRS) that can be achieved along with Pareto-optimality. In our context, RRS can be formalized as follows. Given an instance $I = (N, O, U)$, consider the round-robin sequence in which all agents have the same utility as agent i . In that case, the minimum utility achieved by any of the agents is $RRS_i(I)$. An allocation satisfies RRS if each agent i derives utility of at least $RRS_i(I)$. It would then be very natural to ask what the computational complexity of finding an allocation satisfying both properties is.

EFX allocations. There are further fairness concepts that could be examined from both existence and complexity perspectives, most notably *envy-freeness up to the least valued item* (EFX) [28]. In our setting, one can define an allocation π to be EFX if for any pair of agents i and j , agent i does not envy agent j , or the following two conditions hold:

1. $\forall o \in \pi(i)$ s.t. $u_i(\pi(i) \setminus \{o\}) > u_i(\pi(i)) : u_i(\pi(i) \setminus \{o\}) \geq u_i(\pi(j))$; and
2. $\forall o \in \pi(j)$ s.t. $u_i(\pi(j) \setminus \{o\}) < u_i(\pi(j)) : u_i(\pi(i)) \geq u_i(\pi(j) \setminus \{o\})$.

That is, i 's envy towards j can be eliminated by either removing i 's least valuable good from j 's bundle or removing i 's favorite chore from i 's bundle. Caragiannis et al. [28] mentioned the following “enigmatic” problem: does an EFX allocation exist for goods? It would be interesting to investigate the same question for subjective or objective goods/chores under additive utility.

Connectivity constraints. Finally, recent papers of Bouveret et al. [20] and Bilò et al. [14] demonstrated that a connected allocation satisfying several fairness notions, such as MMS and EF1, is guaranteed to exist for some restricted domains. These existence results rely crucially on the fact that the agents have monotonic utility, and it remains open whether similar results can be obtained in the fair division of indivisible goods and chores.

Acknowledgements We would like to thank Bérczi et al. [12] and Bhaskar et al. [13] for pointing out an error in the IJCAI version [8]. We thank the IJCAI 2019 reviewers and Warut Suksumpong for helpful feedback. Ayumi Igarashi is supported by the KAKENHI Grant-in-Aid for JSPS Fellows number 18J00997 and JST, PRESTO number JPMJPR20C1. Toby Walsh is funded by the European Research Council under the Horizon 2020 Programme via AMPLify 670077.

References

1. Aleksandrov, M., & Walsh, T. (2020). Two algorithms for additive and fair division of mixed manna. In *KI 2020: Advances in artificial intelligence—43rd German conference on AI, Proceedings*, Bamberg, Germany, September 21–25, 2020 (pp. 3–17).
2. Amanatidis, G., Birmpas, G., Christodoulou, G., & Markakis, E. (2017). Truthful allocation mechanisms without payments: characterization and implications on fairness. In *Proceedings of the 18th ACM conference on economics and computation (EC)*.
3. Amanatidis, G., Markakis, E., Nikzad, A., & Saberi, A. (2017). Approximation algorithms for computing maximin share allocations. *ACM Transactions on Algorithms*, 13(4), 52:1–52:28.
4. Aziz, H. (2016). Computational social choice: Some current and new directions. In *Proceedings of the 25th international joint conference on artificial intelligence (IJCAI)* (pp. 4054–4057).
5. Aziz, H., & Rey, S. (2020). Almost group envy-free allocation of indivisible goods and chores. In *Proceedings of the 29th international joint conference on artificial intelligence (IJCAI)* (pp. 39–45).
6. Aziz, H., Biro, P., Lang, J., Lesca, J., & Monnot, J. (2016). Optimal reallocation under additive and ordinal preferences. In *Proceedings of the 15th international conference on autonomous agents and multiagent systems (AAMAS)* (pp. 402–410).
7. Aziz, H., Rauchecker, G., Schryen, G., & Walsh, T. (2017). Algorithms for max-min share fair allocation of indivisible chores. In *Proceedings of the 31st AAAI conference on artificial intelligence (AAAI)* (pp. 335–341).
8. Aziz, H., Caragiannis, I., Igarashi, A., & Walsh, T. (2019). Fair allocation of indivisible goods and chores. In *Proceedings of the 28th international joint conference on artificial intelligence (IJCAI)* (pp. 53–59).
9. Aziz, H., Moulin, H., & Sandomirskiy, F. (2020). A polynomial-time algorithm for computing a Pareto optimal and almost proportional allocation. *Operations Research Letters*, 48(5), 573–578.
10. Barman, S., & Krishnamurthy, S. K. (2020). Approximation algorithms for maximin fair division. *ACM Transactions on Economics and Computation*, 8(1), 5:1–5:28.

11. Barman, S., Kumar Krishna Murthy, S., & Vaish, R. (2018). Finding fair and efficient allocations. In *Proceedings of the 19th ACM conference on economics and computation (EC)* (pp. 557–574).
12. Bérczi, K., Bérczi-Kovács, E. R., Boros, E., Gedefa, F. T., Kamiyama, N., Kavitha, T., Kobayashi, Y., & Makino, K. (2020). Envy-free relaxations for goods, chores, and mixed items. CoRR [arXiv: abs/2006.04428](https://arxiv.org/abs/2006.04428).
13. Bhaskar, U., Sricharan, A. R., & Vaish, R. (2020). On approximate envy-freeness for indivisible chores and mixed resources. *CoRR*[arXiv:abs/2012.06788](https://arxiv.org/abs/2012.06788).
14. Bilò, V., Caragiannis, I., Flammini, M., Igarashi, A., Monaco, G., Peters, D., Vinci, C., & Zwicker, W. S. (2019). Almost envy-free allocations with connected bundles. In *Proceedings of the 10th innovations in theoretical computer science conference (ITCS)* (pp. 14:1–14:21).
15. Bogomolnaia, A., Moulin, H., Sandomirskiy, F., & Yanovskaya, E. (2017). Competitive division of a mixed manna. *Econometrica*, 85(6), 1847–1871.
16. Bogomolnaia, A., Moulin, H., Sandomirskiy, F., & Yanovskaya, E. (2019). Dividing goods and bads under additive utilities. *Social Choice and Welfare*, 52, 395–417.
17. Bouveret, S., & 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.
18. Bouveret, S., & Lang, J. (2011). A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd international joint conference on artificial intelligence (IJCAI)* (pp. 73–78). AAAI Press.
19. Bouveret, S., Chevaleyre, Y., & Maudet, N. (2016). Chapter 12: Fair allocation of indivisible goods. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, & A. D. Procaccia (Eds.), *Handbook of computational social choice* (pp. 284–310). Cambridge University Press.
20. Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., & Peters, D. (2017). Fair division of a graph. In *Proceedings of the 26th international joint conference on artificial intelligence (IJCAI)* (pp. 135–141).
21. Brams, S. J., & Fishburn, P. C. (2000). Fair division of indivisible items between two people with identical preferences: Envy-freeness, Pareto-optimality, and equity. *Social Choice and Welfare*, 17(2), 247–267.
22. Brams, S. J., & Taylor, A. D. (1996). *Fair division: From cake-cutting to dispute resolution*. Cambridge University Press.
23. Brams, S. J., & Taylor, A. D. (1996). A procedure for divorce settlements. *Issue Mediation Quarterly*, 13(3), 191–205.
24. Brams, S. J., & Taylor, A. D. (1996). A procedure for divorce settlements. *Mediation Quarterly*, 13(3), 191–205.
25. Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6), 1061–1103.
26. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., & Kyropoulou, M. (2012). The efficiency of fair division. *Theory of Computing Systems*, 50(4), 589–610.
27. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., & Wang, J. (2016). The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM conference on economics and computation (EC)* (pp. 305–322). ACM Press.
28. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., & Wang, J. (2019). The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3), 12:1–12:32.
29. Chaudhury, B. R., Garg, J., & Mehlhorn, K. (2020). EFX exists for three agents. In *Proceedings of the 21st ACM conference on economics and computation (EC)* (pp. 1–19).
30. Chaudhury, B. R., Garg, J., McGlaughlin, P., & Mehta, R. (2021). Competitive allocation of a mixed manna. In *Proceedings of the 32nd annual ACM-SIAM symposium on discrete algorithms (SODA)*.
31. Conitzer, V., Freeman, R., & Shah, N. (2017). Fair public decision making. In *Proceedings of the 18th ACM conference on economics and computation (EC)* (pp. 629–646).
32. de Keijzer, B., Bouveret, S., Klos, T., & Zhang, Y. (2009). On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proceedings of the 1st international conference on algorithmic decision theory* (pp. 98–110).
33. Eisenberg, E., & Gale, D. (1959). Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1), 165–168.
34. Foley, D. K. (1967). Resource allocation and the public sector. *Yale Economic Essays*, 7, 45–98.
35. Garg, J., & McGlaughlin, P. (2020). Computing competitive equilibria with mixed manna. In *Proceedings of the 19th international conference on autonomous agents and multiagent systems (AAMAS)* (pp. 420–428).

36. Ghodsi, M., HajiAghayi, M., Seddighin, M., Seddighin, S., & Yami, H. (2018). Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)* (pp. 539–556).
37. Kurokawa, D., Procaccia, A. D., & Wang, J. (2018). Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 65(2), 8:1-8:27.
38. Lipton, R. J., Markakis, E., Mossel, E., & Saberi, A. (2004). On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM conference on electronic commerce (EC)* (pp. 125–131).
39. Meunier, F., & Zerbib, S. (2019). Envy-free cake division without assuming the players prefer non-empty pieces. *Israel Journal of Mathematics*, 234, 907–925.
40. Moulin, H. (2019). Fair division in the internet age. *Annual Review of Economics*, 11, 1–37.
41. Plaut, B., & Roughgarden, T. (2018). Almost envy-freeness with general valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2018* (pp. 2584–2603).
42. Robertson, J. M., & Webb, W. A. (1998). *Cake cutting algorithms: Be fair if you can*. A. K. Peters.
43. Segal-Halevi, E. (2018). Fairly dividing a cake after some parts were burnt in the oven. In *Proceedings of the 17th international conference on autonomous agents and multiagent systems (AAMAS)* (pp. 1276–1284).
44. Steinhaus, H. (1949). The problem of fair division. *Econometrica*, 17, 315–319.
45. Stromquist, W. (1980). How to cut a cake fairly. *The American Mathematical Monthly*, 87(8), 640–644.
46. Su, F. E. (1999). Rental harmony: Sperner’s lemma in fair division. *The American Mathematical Monthly*, 106(10), 930–942.
47. Varian, H. R. (1974). Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1), 63–91.
48. Woodall, D. R. (1980). Dividing a cake fairly. *Journal of Mathematical Analysis and Applications*, 78(1), 233–247.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.