

Fair Allocation of Indivisible Goods and Chores

Haris Aziz¹, Ioannis Caragiannis², Ayumi Igarashi³ and Toby Walsh¹

¹UNSW Sydney and Data61 CSIRO, Sydney, Australia

²University of Patras, Patras, Greece

³University of Tokyo, Tokyo, Japan

haris.aziz@unsw.edu.au, caragian@ceid.upatras.gr, igarashi@mist.i.u-tokyo.ac.jp, tw@cse.unsw.edu.au

Abstract

We consider the problem of fairly dividing a set of items. Much of the fair division literature assumes that the items are “goods” i.e., they yield positive utility for the agents. There is also some work where the items are “chores” that yield negative utility for the agents. In this paper, we consider a more general scenario where an agent may have negative or positive utility for each item. This framework captures, e.g., fair task assignment, where agents can have both positive and negative utilities for each task. We show 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 satisfying certain fairness and efficiency properties and further study 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 subjective views regarding how enjoyable each task is. For some people, 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 valued goods so that she does not feel hard done by.

This example can be viewed as an instance of a classic fair division problem. The agents have different preferences over the items and we want to allocate the items to agents as fair as possible. The twist we consider is that whether an agent has positive or negative utility for an item is subjective. Our setting is general enough to encapsulate two well-studied settings: (1) “good allocation” where agents have positive utilities for the items and (2) “chore allocation” in which agents have negative utilities for the items. Our setting also covers (3) “allocation of objective goods and chores” where 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 where 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 chore and good allocation. For this model, we present some case studies that highlight 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. Our central technical contributions are several new efficient algorithms for finding fair allocations.

Contributions.

- We first formalize fairness concepts for the general setting. Some fairness concepts directly extend from the setting of good 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 then show a careful generalization of the decentralized round robin algorithm that finds an EF1 allocation when utilities are additive.
- We also present a different polynomial-time algorithm that always returns an EF1 allocation even when the agents’ utility functions are arbitrary (not necessarily additive).
- Turning our attention to an efficient and fair allocation, we show that for 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 [Brams and Taylor, 1996a; 1996b] that is designed for divisible goods.
- If we weaken EF1 to PROP1, then we show that there exists an allocation that is not only PROP1 but is also contiguous (assuming that items are placed in a line). We further give a polynomial-time algorithm that finds such allocation.

Related work Fair allocation of indivisible items is a central problem in several fields including computer science and economics [Aziz et al., 2015; Brams and Taylor, 1996a; Bouveret, Chevaleyre, and Maudet, 2016; Lipton et al., 2004]. There are several established notions of fairness, including envy-freeness and proportionality.

The idea of envy-freeness up to one good was implicit in the paper by Lipton et al. [2004]. Today, it has become a well-studied fairness concept in its own right [Budish, 2011]. Caragiannis et al. [2016] further popularized it, showing that a natural modification of the Nash welfare maximizing rule satisfies EF1 and PO for the case of goods. Barman, Murthy, and Vaish [2018] recently presented a pseudo-polynomial-time algorithm for computing an allocation that is PO and EF1 for goods. EFX (envy-freeness up to an extreme item) was introduced by Caragiannis et al. [2016].

Aziz [2016] noted that the work on multi-agent chore allocation is less developed than that of goods and that results from one to the other may not necessarily carry over. Barman and Murthy [2017] presented a better approximation algorithm. Caragiannis et al. [2012] studied the efficiency loss in order to achieve several fair allocations in the context of both good and chore divisions. Allocation of a mixture of goods and chores has received recent attention in the context for divisible items [Bogomolnaia et al., 2019; 2017]. Here, we focus on indivisible items.

Full version. The full version of the paper is available on arXiv [Aziz et al., 2018].

2 Our Model and Fairness Concepts

We now define a fair division problem of indivisible items where agents may have both positive and negative utilities. 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 $u_i : O \rightarrow \mathbb{R}$.

We note that under this model, an item can be a good for one agent (i.e., $u_i(o) > 0$) but a chore for another agent (i.e., $u_j(o) < 0$). For $X \subseteq O$, we write $u_i(X) := \sum_{o \in X} u_i(o)$; we assume that the utilities in this paper are additive unless specified otherwise. Each subset $X \subseteq O$ is referred to as a *bundle* of items. An allocation π is a function $\pi : N \rightarrow 2^O$ assigning each agent a different bundle of items, i.e., for every pair of distinct agents $i, j \in N$, $\pi(i) \cap \pi(j) = \emptyset$; it is said to be *complete* if $\bigcup_{i \in N} \pi(i) = O$.

We first observe that the definitions of standard fairness concepts can be naturally extended to this general model. The most classical fairness principle is *envy-freeness*, 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 the other agents. Another appealing notion of fairness is *proportionality* which guarantees each agent an $1/n$ fraction of her utility for the whole items. Formally, an allocation π is *proportional* (PROP) if each agent $i \in N$ receives a bundle $\pi(i)$ of value at least her *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 *An envy-free complete allocation satisfies proportionality.*

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, thereby, focused on approximation of these fairness concepts. A prominent relaxation of envy-freeness, introduced by Budish [2011], is *envy-freeness up to one good* (EF1), which requires that an agent’s envy towards another bundle can be eliminated by removing some good from the envied bundle. We will present a generalized definition for EF1 that have only been considered in the context of good allocation: the envy can diminish by removing either one “good” from the other’s bundle or one “chore” from their own bundle.

Definition 1 (EF1) *An allocation π is envy-free up to one item (EF1) if for all $i, j \in N$, either i does not envy j , or there is an item $o \in \pi(i) \cup \pi(j)$ such that $u_i(\pi(i) \setminus \{o\}) \geq u_i(\pi(j) \setminus \{o\})$.*

Obviously, envy-freeness implies EF1. Conitzer, Freeman, and Shah [2017] introduced a novel relaxation of proportionality, which they called *PROPI*. In the context of good allocation, this fairness relaxation is a weakening of both EF1 and proportionality, requiring that each agent gets her proportional fair share if 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 her bundle.

Definition 2 (PROPI) *An allocation π satisfies proportionality up to one item (PROP1) 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 $o \in \pi(i)$.

It can be easily verified that EF1 implies PROP1. See the online appendix for an illustration of relations between fairness concepts introduced above.

Proposition 2 *An EF1 complete allocation satisfies PROPI.*

Besides 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 .

3 Finding EF1 Allocation

In this section, we focus on EF1, a very permissive fairness concept that admits a polynomial-time algorithm in the case of good allocation. We will present two algorithms which appropriately generalize the known algorithms for finding an EF1 allocation of goods, the *round robin rule* and the *envy-graph algorithm* [Lipton et al., 2004], to our setting.

Double Round Robin Algorithm Consider a *round robin rule* in which agents take turns, and choose their most preferred unallocated item. The round robin rule finds an EF1

allocation if all the items are goods (see e.g., Caragiannis et al. [2016]). 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 already fails to find an EF1 allocation if we have some items that are goods and others that are chores.

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

Proof: Suppose there are two agents and four items with identical utilities described below.

	①	②	③	④
Alice, Bob:	2	-3	-3	-3

Consider the order, in which Alice chooses the only good and then the remaining chores of equal value are allocated accordingly. In that case, Alice gets one a high value good and one chore, whereas Bob gets two chores. So even if one item is removed from the bundles of Alice or Bob, Bob will still remain envious. \square

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 to agents *chores*, i.e., the items for which each agent has non-positive utility, while in the second phase, the reversed round-robin algorithm allocates to agents the remaining *goods*, in the opposite order starting with the agent who is worst off 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; hence the envy of i towards j can be bounded up to one item. We formalize the idea in Algorithm 1; see Figure 1 for an illustration.

Algorithm 1 Double Round Robin Algorithm

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

Output: An allocation π

- 1 Partition O into $O^+ = \{o \in O \mid \exists i \in N \text{ s.t. } u_i(o) > 0\}$, $O^- = \{o \in O \mid \forall i \in N, u_i(o) \leq 0\}$. Suppose $|O^-| = an - k$ for some integer a and $k \in \{0, \dots, n - 1\}$.
- 2 Create k dummy chores for which each agent has utility 0, and add them to O^- . (Hence $|O^-| = an$.)
- 3 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.
- 4 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.
- 5 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 $u_i(\pi(i) \cap X) < u_i(\pi(j) \cap X)$.




Figure 1: Illustration of Double Round Robin Algorithm. The dotted line corresponds to the picking order when allocating chores. The thick line corresponds to the picking order when allocating goods. The solid black circle indicates the agent who starts with the picking. For the dotted (chores) round, agent 1 is the first agent to pick. For the solid (goods) round, agent n is the first agent to pick.

Theorem 1 *The double round robin algorithm (Algorithm 1) returns an EF1 complete allocation in $O(\max\{m^2, mn\})$ time.*

Proof: We note that the algorithm ensures that all agents receive the same number of chores, by introducing k virtual 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, these agents do not envy each other.

First, consider i 's envy for j . We first observe that for the k -th item in O^- allocated to i is weakly preferred by i than the k -th item in O^- allocated to j . Hence, agent i does not envy j with respect to O^- . As for the goods allocation, agent i may envy agent j with respect to O^+ , which implies that j picks at least one more item from O^+ . But if the first item o^* picked by j from O^+ is removed from j 's bundle, then the envy will diminish, i.e., i does not envy j with respect to $O^+ \setminus \{o^*\}$. The reason is that for each item in $O^+ \setminus \{o^*\}$ picked by j there is a corresponding item picked by i before j 's turn that is at least as preferred by i . Thus $u_i(\pi(i)) \geq u_i(\pi(j) \setminus \{o^*\})$.

Second, consider j 's envy for i . Similarly, agent j does not envy agent i with respect to O^+ because he takes the first pick among i and j ; that is, for every item in $o \in O^+ \cap \pi(i)$ such that $u_j(o) > 0$, agent j picks a corresponding item before i that she weakly prefers. As for the items in O^- , let o^* be the last item from O^- chosen by j . Then, for each item $o \in O^- \setminus \{o^*\}$ picked by i , there is a corresponding item picked by j before i that j weakly prefers to o , which implies j does not envy i with respect to $O^- \setminus \{o^*\}$. Thus $u_j(\pi(j) \setminus \{o^*\}) \geq u_j(\pi(i))$.

In either case, agents do not envy each other up to one item. We conclude that π is EF1 and so does the final allocation π^* as removing dummy items does not affect the utilities of each agent. It remains to analyze the running time of Algorithm 1. Line 1 requires $O(mn)$ time as each item needs to be examined by all agents. Lines 3 and 4 require $O(m^2)$ time as there are at most m iterations, and for each iteration, each agent considers at most m candidates. Thus, the total running time can be bounded by $O(\max\{m^2, mn\})$, which completes the proof. \square

Generalized Envy Graph Algorithm Algorithm 1 is designed for additive utilities. We construct another algorithm

(Algorithm 2) that finds an EF1 allocation for arbitrary utility functions $u_i : 2^O \rightarrow \mathbb{R}$ (not necessarily additive). The algorithm is based on a generalization of an algorithm presented by Lipton et al. [2004] for finding an EF1 allocation for goods. For an allocation π , the *envy-graph* $G(\pi)$ is a directed graph where the vertices is given by the set of agents N , and there is an arc from i to j if and only if i envies j . For each directed cycle $C = \{i_1, i_2, \dots, i_k\}$ of the envy graph $G(\pi)$ where i_j envies i_{j+1} for each $j \in [k]$ and $i_{k+1} = i_1$, we may implement an exchange over the cycle, and define the resulting allocation π_C as follows:

$$\pi_C(i) = \begin{cases} \pi(i) & \text{if } i \notin C, \\ \pi(i_{j+1}) & \text{if } i = i_j \in C. \end{cases}$$

A *source* of a directed graph is a vertex with no incoming arcs, whereas a *sink* is a vertex with no outgoing arcs. Given an allocation π , the *marginal utility* of an item $o \in O \setminus \pi(i)$ to an agent $i \in N$ is defined as $\delta_i(\pi, o) := u_i(\pi(i) \cup \{o\}) - u_i(\pi(i))$.

Theorem 2 *Suppose that we are given an oracle access to the utility functions $u_i : 2^O \rightarrow \mathbb{R}$ of all agents $i \in N$. Then the generalized envy-graph algorithm (Algorithm 2) finds an EF1 complete allocation in $O(mn^3)$.*

Proof: We will prove by induction that each time a new item is allocated and a while loop of Algorithm 2 is executed, the envy-graph $G(\pi)$ is acyclic and the allocation π is EF1. The base case clearly holds since the initial allocation corresponds to the null allocation. Suppose that $k-1$ items have been allocated and we want to allocate the k -th item o . If $N^+ \neq \emptyset$, then let G^+ be the envy graph induced by N^+ . Since by the induction hypothesis, the envy-graph is acyclic, which means that G^+ is acyclic as well. Hence, at least one agent i^* has no incoming arc towards i^* in G^+ and the algorithm can give i^* item o . Since no agent envied agent i^* before, even if some agent now envies agent i^* , the allocation π satisfies EF1. If $N^+ = \emptyset$, then we know that all the agents have negative marginal utility for item o . Since by the induction hypothesis, the envy graph $G(\pi)$ is acyclic, and thus at least one agent i^* has no out-degree, which follows that we can give i^* item o . Since i^* envied no one before, even if i^* envies other agents now, the allocation π satisfies EF1.

We now focus on the while loop in the algorithm whereby envy cycles are removed by exchanging allocations along the cycle. After exchanging bundles over cycle C , we observe that the agents in the cycle improve their utility and the envy towards each agent in C is still bounded up to one good as the set of bundles does not change. Hence, the resulting allocation π_C still satisfies EF1. Further, by removing one cycle, we note that each agent in the cycle has her out-degree decreased by one. Furthermore no agent outside the cycle has her outdegree changed. Hence, after a linear number of cycles removed, the graph has no more envy-cycle and hence is acyclic. Similarly to [Lipton et al., 2004], the running time of the algorithm can be bounded by $O(mn^3)$. \square

4 Finding EF1 and PO allocation

We move on the the next question as to whether fairness is achievable together with efficiency. In the context of goods

Algorithm 2 Generalized Envy Graph Algorithm

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

Output: An allocation π

```

1 Initialize allocation  $\pi(i) = \emptyset$  for all  $i \in N$ 
2 for  $o \in O$  do
3   Let  $N^+$  be the set of agents  $i$  who have a non-
   negative marginal utility  $\delta_i(\pi, o) \geq 0$ .
4   if  $N^+ \neq \emptyset$  then
5     Choose a source  $i^* \in N^+$  in the graph  $G(\pi)$  in-
     duced by  $N^+$ .
6   else
7     Choose a sink  $i^* \in N$  in  $G(\pi)$ .
8   Update  $\pi(i^*) \leftarrow \pi(i^*) \cup \{o\}$ .
9   while  $G(\pi)$  contains a directed cycle  $C$  do
10    Update  $\pi \leftarrow \pi_C$ .
```

allocation where agents have non-negative additive utilities, Caragiannis et al. [2016] proved that an outcome that maximizes the *Nash welfare* (i.e., the product of utilities) satisfies EF1 and Pareto-optimality simultaneously. The question regarding whether a Pareto-optimal and EF1 allocation exists for chores is unresolved. If one starts from an EF1 allocation and finds Pareto improvements, one runs into two challenges: first, Pareto improvements may not necessarily preserve EF1; second, finding Pareto improvements is NP-hard [Aziz et al., 2016; de Keijzer et al., 2009]. Even if we ignore the second challenge, the question regarding the existence of a Pareto-optimal and EF1 allocation for chores is open.

Generalized Adjusted Winner Next we show that the problem of finding an EF1 and Pareto-optimal allocation is completely resolved for the restricted but important case of two agents. Our algorithm can be viewed as a discrete version of the well-known Adjusted Winner (AW) rule [Brams and Taylor, 1996a; 1996b]. Just like AW, our algorithm is Pareto-optimal and EF1. In contrast to AW that is designed for goods, our algorithm can handle both goods and chores.

Theorem 3 *For two agents, a Pareto-optimal and EF1 complete allocation always exists and can be computed in polynomial time.*

Proof: The algorithm begins by giving each subjective item to the agent who considers it as a good; that is, for each item $o \in O$ allocate o to agent i if $u_i(o) \geq 0$ and $u_j(o) < 0$ where $j \in N \setminus \{i\}$. So, in the following, let us assume that there is no item for which each agent has utility 0. Also we assume that we have objective items only, i.e., 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 *winner* (denoted by w) and another *loser* (denoted by ℓ).

- (i) Initially, all goods are allocated to the winner and all chores to the loser.
- (ii) We sort the items in terms of $|u_\ell(o)|/|u_w(o)|$ (monotone non-increasing order), and consider reallocation of the items according to the ordering (from the left-most to the right-most item).

(iii) 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 we find an EF1 allocation from the point of view of the loser. Note that the loser is envious up to one item of the winner.

We will first prove that at any point of the algorithm, the allocation π is Pareto-optimal, and so is the final allocation π^* . Assume towards a contradiction that the allocation π is Pareto-dominated by the allocation π' . For each $i, j \in \{w, \ell\}$ with $i \neq j$, we write

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

Without loss of generality, we assume that, in π , the winner has utility which is at least as high as in π' , while the loser is strictly better off. Taking into account 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

$$u_w(G_{\ell w}) + u_w(C_{\ell w}) - u_w(G_{w\ell}) - u_w(C_{w\ell}) \geq 0; \text{ and} \quad (1)$$

$$u_\ell(G_{w\ell}) + u_\ell(C_{w\ell}) - u_\ell(G_{\ell w}) - u_\ell(C_{\ell w}) > 0 \quad (2)$$

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 (this is why the allocation of the items in the first two bundles changes while the allocation of the items in the last two bundles does not). Let α be such that

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

This definition implies the inequalities,

$$u_\ell(G_{w\ell}) \leq \alpha u_w(G_{w\ell}); u_\ell(G_{\ell w}) \geq \alpha u_w(G_{\ell w});$$

$$-u_\ell(C_{w\ell}) \geq -\alpha u_w(C_{w\ell}); -u_\ell(C_{\ell w}) \leq -\alpha u_w(C_{\ell w}),$$

which, together with inequality (2), yields

$$0 < u_\ell(G_{w\ell}) + u_\ell(C_{w\ell}) - u_\ell(G_{\ell w}) - u_\ell(C_{\ell w})$$

$$\leq -\alpha(u_w(G_{\ell w}) + u_w(C_{\ell w}) - u_w(G_{w\ell}) - u_w(C_{w\ell})) \leq 0,$$

a contradiction. The last inequality follows by (1) and by the fact that α is non-negative.

Now observe that at the final allocation π^* , at most one agent 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, a contradiction. Thus, π^* is EF1 when the loser envies the winner at π^* . 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, and $X = \pi'(w) \cap \pi^*(w)$ and $Y = \pi'(\ell) \cap \pi^*(\ell)$. By construction, the loser envies the winner more than one item at π' , which implies $u_\ell(Y) < u_\ell(X)$. Suppose towards a contradiction that the winner envies the loser more than one item at π^* , which implies $u_w(X) < u_w(Y)$. If g is the last good that has been moved from the winner to the loser, then allocating X to ℓ and $Y \cup \{g\}$ to w would be a Pareto-improvement

of π' , a contradiction. Similarly, if c is the last chore that has been moved from the loser to the winner, then allocating $X \cup \{c\}$ to ℓ and Y to w would be a Pareto-improvement of π' , a contradiction. Hence, the winner envies the loser up to one item; we conclude that π^* is EF1. \square

The following example illustrates the generalized AW.

Example 4 (Example of Generalized AW) Consider two agents, Alice and Bob, and five items with the following additive utilities where the grey circles correspond to goods and the white circles correspond to 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

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 up to one item. Hence the final allocation gives the items 2 and 4 to Alice and the rest to Bob.

A natural question is whether PO and EF1 allocation exists for three agents with general additive utilities; we leave this as an interesting open question. We note that Pareto-optimality by itself is easy to achieve, by taking a permutation of agents and applying a variant of ‘serial dictatorship’.

Proposition 4 A Pareto-optimal allocation can be computed in linear time.

5 Finding Contiguous PROP1 Allocation

We saw that an EF1 allocation always exists for subjective goods/chores. If we weaken EF1 to PROP1, one can achieve another requirement besides fairness, that is, *connectivity*. Specifically, we will consider a situation when the items are placed on a path, and each agent gets a connected bundle of the path. Finding a connected set of items is relevant in many scenarios, in particular when the items have a spatial or temporal structure. For example, the items can be a set of rooms in a corridor and agents will often value being allocated a contiguous chunk of rooms (see e.g., [Bouveret et al., 2017]).

We will show that a connected PROP1 allocation exists and can be found efficiently. Biló et al. [2018] recently consider related questions regarding EF1 allocations of goods. Without loss of generality, we assume that the path is given by a sequence of items (o_1, o_2, \dots, o_m) . We first prove a result for the case of cake cutting 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 $j \in [m]$. An *allocation* of a mixed cake assigns each agent a disjoint sub-interval of the cake where the union of the intervals equals the entire cake $[0, m]$; it satisfies *proportionality* if each agent i gets an interval of value at least the proportional fair share $u_i(O)/n$.

Theorem 5 *A contiguous proportional allocation of a mixed cake exists and can be computed in polynomial time.*

Proof sketch: Let N^+ be the set of agents with strictly positive total value for O . We combine the moving-knife algorithms for goods and chores. First, if there is an agent who has positive proportional fair share, i.e., $N^+ \neq \emptyset$, 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 value of exactly the proportional fair share. The first agent j who shouts is allocated to the left bundle $[0, x_j]$, and the algorithm recurs on the remaining instance. Second, if no agent has positive proportional fair share, our algorithm moves a knife from right to left, and agents shout whenever the left part of the cake has value exactly proportional fair share. Again, the first agent j who shouts is allocated to the left bundle $[0, x_j]$, and the algorithm recurs on the remaining instance.

We will prove by induction on $|N|$ that the allocation of a mixed cake produced by the algorithm satisfies the following:

- if $N^+ \neq \emptyset$, then each agent in N^+ receives an interval of value at least proportional fair share and each agent not in N^+ receives an empty piece; and
- if $N^+ = \emptyset$, then each agent receives an interval of value at least proportional fair share.

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 positive proportional fair share, i.e., $N^+ \neq \emptyset$. If $|N^+| = 1$, the claim is trivial; thus assume otherwise. Clearly, a shouter j receives an interval of value at least the proportional fair share. Further, all other agents in N^+ have value at most the proportional fair share for the left piece $[0, x_j]$. Indeed, if some agent $i \in N^+$ values the left piece greater than the proportional fair share, then i would have shouted when the knife reaches before x_j by the continuity of \hat{u}_i , and $\hat{u}_i([0, x_j]) > \hat{u}_i([0, 0])$, contradicting the minimality of x_j . Thus, the remaining agents in N^+ have at least $(n - 1) \cdot \hat{u}_i([0, m])/n$ value for $[x_j, m]$. By the induction hypothesis, each agent in N^+ gets an interval of value at least proportional fair share, and the rest gets an empty piece.

Suppose that no agent has positive proportional fair share. Again, if some agent i values $[0, x_j]$ greater than the proportional fair share, then i would have shouted when the knife reaches before x_j by the continuity of \hat{u}_i and $\hat{u}_i([0, x_j]) > \hat{u}_i([0, m])$, contradicting the maximality of x_j . Thus all the remaining agents have at least $(n - 1) \cdot \hat{u}_i([0, m])/n$ value for $[x_j, m]$, and hence by the induction hypothesis each agent gets an interval of value at least proportional fair share. \square

The theorem stated above also applies to a general cake-cutting model in which information about agent’s utility function over an interval can be inferred by a series of queries. We note that a contiguous envy-free allocation of a mixed cake is known to exist only when the number of agents is four or a prime number [Segal-Halevi, 2018; Meunier and Zerbib, 2018]. A fractional proportional allocation can be then used to achieve a contiguous PROP1 division of indivisible items.

Theorem 6 *A connected PROP1 allocation of a path always exists and can be computed in polynomial time.*

We note that our PROP1 existence result extends to any *traceable* graph that admits a Hamiltonian path: we can run our algorithm on the Hamiltonian path and the resulting output can be easily seen to satisfy both contiguity and PROP1.

6 Discussion

In this paper, we formally studied fair allocation when the items are a combination of subjective goods and chores. Our work paves the way for detailed examination of allocation of goods/chores, and opens up an interesting line of research, with many problems left open to explore. In particular, there are further fairness concepts that could be studied from both existence and complexity issues, most notably *envy-freeness up to the least valued item* (EFX) [Caragiannis et al., 2016]. In our setting, one can define an allocation π to be EFX if for any pair of agents i, j , the following two hold:

- (i) $\forall o \in \pi(i)$ s.t. $u_i(o) < 0$: $u_i(\pi(i) \setminus \{o\}) \geq u_i(\pi(j))$; and
- (ii) $\forall o \in \pi(j)$ s.t. $u_i(o) > 0$: $u_i(\pi(i)) \geq u_i(\pi(j) \setminus \{o\})$.

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. [2016] mentioned the following ‘enigmatic’ problem: does an EFX allocation exist for goods? It would be intriguing to investigate the same question for subjective or objective goods/chores.

We also note that while the relationship between Pareto-optimality and most of fairness notions is still unclear, Conitzer, Freeman, and Shah [2017] proposed a fairness concept called *Round Robin Share* that can be achieved together 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 utilities as agent i . In that case, the minimum utility achieved by any of the agents is $\text{RRS}_i(I)$. Formally for each agent i , we order the items in a non-increasing order of her utilities, i.e., $u_i(o^{(1)}) \geq u_i(o^{(2)}) \geq \dots \geq u_i(o^{(m)})$; the round robin share for i is given by $\text{RRS}_i(I) = \min_{j \in [n]} \sum_{k=0}^{p-1} u_i(o^{(j+k \cdot n)})$ where $p = \lfloor m/n \rfloor$. An allocation satisfies RRS if each agent i gets utility at least $\text{RRS}_i(I)$. It would be then very natural to ask what is the computational complexity of finding an allocation satisfying both properties.

Finally, recent works of Bouveret et al. [2017] and Biló et al. [2018] showed that a connected allocation satisfying several fairness notions, such as MMS and EF1, is guaranteed to exist for some restricted domains. Although these existence results crucially rely on the fact that the agents have monotonic valuations, it remains open whether similar results can be obtained in fair division of indivisible goods and chores.

Acknowledgements. We thank the reviewers at IJCAI-19 for helpful feedback. Ayumi Igarashi is supported by the KAKENHI Grant-in-Aid for JSPS Fellows number 18J00997. Haris Aziz is supported by a UNSW Scientia Fellowship. Toby Walsh is funded by the European Research Council under the Horizon 2020 Programme via AMPLify 670077.

References

- [Aziz et al., 2015] Aziz, H.; Gaspers, S.; Mackenzie, S.; and Walsh, T. 2015. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* 227:71–92.
- [Aziz et al., 2016] Aziz, H.; Biro, P.; Lang, J.; Lesca, J.; and Monnot, J. 2016. Optimal reallocation under additive and ordinal preferences. In *Proc. of the 15th AAMAS Conference*, 402–410.
- [Aziz et al., 2018] Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2018. Fair allocation of combinations of indivisible goods and chores. *CoRR* abs/1807.10684.
- [Aziz, 2016] Aziz, H. 2016. Computational social choice: Some current and new directions. In *Proc. of the 25th IJCAI*, 4054–4057.
- [Barman and Murthy, 2017] Barman, S., and Murthy, S. K. K. 2017. Approximation algorithms for maximin fair division. In *Proc. of the 18th ACM-EC Conference*, 647–664.
- [Barman, Murthy, and Vaish, 2018] Barman, S.; Murthy, S. K. K.; and Vaish, R. 2018. Finding fair and efficient allocations. In *Proc. of the 19th ACM-EC Conference*. ACM Press.
- [Biló et al., 2018] Biló, V.; Caragiannis, I.; Flammini, M.; Igarashi, A.; Monaco, G.; Peters, D.; Vinci, C.; and Zwicker, W. S. 2018. Almost Envy-Free Allocations with Connected Bundles. In Blum, A., ed., *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, volume 124 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 14:1–14:21. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [Bogomolnaia et al., 2017] Bogomolnaia, A.; Moulin, H.; Sandomirskyi, F.; and Yanovskaya, E. 2017. Competitive division of a mixed manna. *Econometrica* 85(6):1847–1871.
- [Bogomolnaia et al., 2019] Bogomolnaia, A.; Moulin, H.; Sandomirskyi, F.; and Yanovskaya, E. 2019. Dividing bads under additive utilities. *Social Choice and Welfare* 52(3):395–417.
- [Bouveret et al., 2017] Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 135–141.
- [Bouveret, Chevaleyre, and Maudet, 2016] Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair allocation of indivisible goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 12.
- [Brams and Taylor, 1996a] Brams, S. J., and Taylor, A. D. 1996a. *Fair division: From cake-cutting to dispute resolution*. Cambridge University Press.
- [Brams and Taylor, 1996b] Brams, S. J., and Taylor, A. D. 1996b. A procedure for divorce settlements. *Issue Mediation Quarterly Mediation Quarterly* 13(3):191–205.
- [Budish, 2011] Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.
- [Caragiannis et al., 2012] Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2012. The efficiency of fair division. *Theory of Computing Systems* 50(4):589–610.
- [Caragiannis et al., 2016] Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum Nash welfare. In *Proc. of the 17th ACM-EC Conference*, 305–322. ACM Press.
- [Conitzer, Freeman, and Shah, 2017] Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair public decision making. In *Proc. of the 18th ACM-EC Conference*, 629–646. ACM Press.
- [de Keijzer et al., 2009] de Keijzer, B.; Bouveret, S.; Klos, T.; and Zhang, Y. 2009. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *Proc. of the 1st International Conference on Algorithmic Decision Theory*, 98–110.
- [Lipton et al., 2004] Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proc. of the 5th*, 125–131. ACM Press.
- [Meunier and Zerbib, 2018] Meunier, F., and Zerbib, S. 2018. Envy-free divisions of a partially burnt cake. *CoRR* abs/1804.00449.
- [Segal-Halevi, 2018] Segal-Halevi, E. 2018. Fairly dividing a cake after some parts were burnt in the oven. In *Proc. of the 17th AAMAS Conference*, 1276–1284. Extended version available as arXiv:1704.00726v5 [math.CO].