Fair Allocation of Resources with Uncertain Availability

Jan Buermann
University of Southampton
Southampton, UK
j.buermann@soton.ac.uk

Enrico H. Gerding
University of Southampton
Southampton, UK
eg@ecs.soton.ac.uk

Baharak Rastegari
University of Southampton
Southampton, UK
b.rastegari@soton.ac.uk

ABSTRACT

Multi-agent resource allocation is an important and well-studied problem within AI and economics. It is generally assumed that the quantity of each resource is known a priori. However, in many real-world problems, such as the production of renewable energy which is typically weather dependent, the exact amount of each resource may not be known at the time of decision making. In this paper we investigate fair division of a homogeneous divisible resource where the available amount is given by a probability distribution. Specifically, we study the notion of ex-ante envy-freeness, where, in expectation, agents weakly prefer their allocation over every other agent's allocation. We analyse the trade-off between fairness and social welfare. We show that allocations satisfying ex-ante envy-freeness can result in higher social welfare compared to those satisfying ex-post envy-freeness. Nevertheless, the price of envy-freeness is at least $\Omega(n)$, where n is the number of agents, and this is tight under concave valuation functions. Principally, we show that the problem of optimising ex-ante social welfare subject to ex-ante envy-freeness is NP-hard in the strong sense. Finally, we devise an integer program to calculate the optimal ex-ante envy-free allocation for linear satiable valuation functions.

KEYWORDS

Fair allocation, Social choice theory

ACM Reference Format:
Jan Buermann, Enrico H. Gerding, and Baharak Rastegari. 2020. Fair Allocation of Resources with Uncertain Availability. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 9 pages.

1 INTRODUCTION

We consider the problem of dividing a homogeneous resource among interested agents in a fair manner without using payments. This problem has a wide range of applications such as the allocation of electricity, estate, storage space, bandwidth or processing time [12, 27, 33]. Among the various in the literature introduced notions of fairness, envy-freeness has received a lot of attention in social choice theory [9, 14]. This natural criterion requires that no agent prefers another agent's allocation over his/her own. While obtaining envy-freeness is trivial for homogeneous divisible goods, we are also interested in maximising social welfare (i.e. efficiency) and understanding the trade-offs between them. In addition, we consider, for the first time, these two problems when the amount of available resource is ex-ante uncertain. In this case the allocations are conditional on the events (i.e. the amount of available resource) and this allows us to distinguish between two notions of envy-freeness and efficiency: ex-ante, i.e. in expectation based on the probability distribution, and ex-post, i.e. at the time of consumption.

A particular example of interest is that of local energy exchange markets [24]. In these communal markets fairness plays a major role. Furthermore, the amount of available energy is uncertain due the variable production of renewables [22, 32].

The literature on fair division of resources is extensive but mainly focuses on the allocation of bundles of items or heterogeneous resources [6, 11]. One reason for this is that, without payments or uncertainty, the envy-free allocation of a homogeneous good is trivial and consists of giving every agent the same amount [12]. Even with uncertainty, ex-post envy-freeness would generally require equal distribution in all events and therefore inhibits ex-ante social welfare improvements. However, relaxing envy-freeness to only hold in expectation allows social welfare improvements in expectation. The following example shows that an ex-ante envy-free allocation can achieve more ex-ante social welfare than the ex-post envy-free solution.

Example 1.1. Consider a renewable energy setting where 2 agents share a photovoltaic system. They want to plan for the next day to be able to acquire further electricity from other sources if necessary. The weather forecast for the next day has the weather as one of two events; event ω_1 is that the day is cloudy which occurs with a probability of $^2/_3$, whereas event ω_2 is that the day is sunny which occurs with a probability of $^1/_3$. Based on these weather predictions, the photovoltaic system produces an amount of $0.2\,kWh$ and $0.4\,kWh$ of energy, respectively. The two agents are interested in electricity and value it according to the valuation functions $v_1(x) = \frac{5}{0.3} \cdot x$ for $0 \leq x < 0.3$, $v_1(x) = 5$ for $x \geq 0.3$, and $v_2(x) = \frac{1}{0.2} \cdot x$ for $0 \leq x < 0.2$ and $v_2(x) = 1$ for $x \geq 0.2$ (see Figure 1a). The objective is to find an ex-ante allocation. In all three cases, the next day the agents would get deterministically the allocation associated with the event (the actual weather).

In the social welfare maximising allocation, agent 1 would get everything in the first event and $0.3\,kWh$ in the second, while agent 2 gets the remaining $0.1\,kWh$ (see Figure 1b), resulting in an expected social welfare of $4\frac{1}{18}$. In contrast, giving equal amounts to both agents in both cases (see Figure 1c) achieves an expected social welfare of $2\frac{8}{9}$. Since this allocation gives both agents the same in both events, it is both ex-ante and ex-post envy-free. Now, consider the allocation where the first agent gets $0.075kWh$ in the first event and $0.3kWh$ in the second event, and the second agent gets the remaining energy each time (see Figure 1d). For the first event agent 1 values his and the other agent's allocation at $\frac{5}{4}$ and $2\frac{1}{12}$ respectively. For the second event, agent 2 values his and the other agent's allocation at $\frac{1}{2}$ and 1. Hence, the allocation for neither of

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

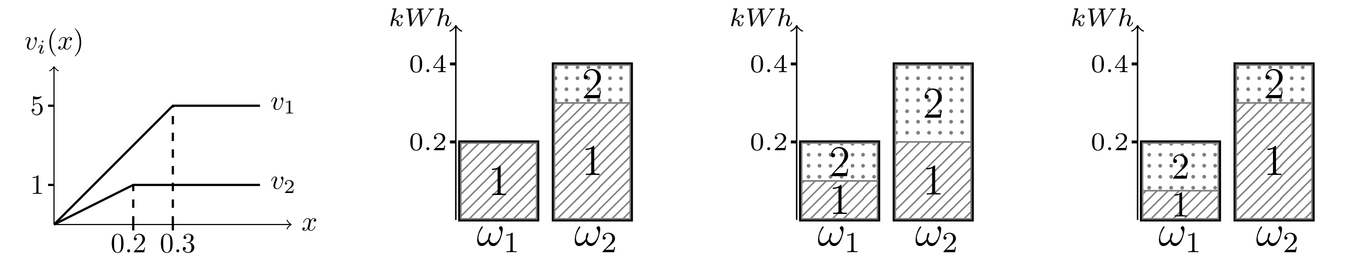

(a) The agents' valuation functions in the example. **(b) The efficient allocation unconstrained by envy-freeness.** **(c) The ex-post envy-free allocation.** **(d) An ex-ante but not ex-post envy-free allocation.**

Figure 1: The allocations of the example. In each case, the bars are the two events, ω_1 and ω_2 with an available amount of $0.2\,kWh$ and $0.4\,kWh$, respectively. The patterns and numbers inside the events refer to the allocations to the agents. The diagonal lines indicate an allocation to agent 1 and the dots indicate an allocation to agent 2.

the events is ex-post envy-free. In comparison, considering ex-ante valuations, agent 1 values his own and the other agent's allocation at $2\frac{1}{2}$ and $1\frac{51}{54}$ respectively, and agent 2 values his own allocation as well as the other agent's allocation both at $\frac{7}{12}$. Hence, the allocation is ex-ante envy-free and has with $3\frac{1}{12}$ an ex-ante social welfare higher than that of the ex-post envy-free allocation.

The example highlights how the allocation illustrated in Figure 1d utilises that the allocation can vary between the events. This allows increased ex-ante social welfare in comparison to the allocation illustrated in Figure 1c.

Related to our work, Feige and Tennenholtz [12] showed how to use a relaxation of fairness to hold in expectation to improve the ex-ante social welfare of a fair allocation by introducing randomisation of allocations. However, they not only use a different fairness measure but their mechanisms are uniform lotteries over equally sized allocations of the same resource amount, while in our setting the probability distribution is given exogenously and the events signify different amounts of available resource. Furthermore, in contrast to their lotteries, we use deterministic (conditional) allocations to improve ex-ante social welfare under ex-ante envy-freeness. While efficiency can be improved by considering ex-ante envy-freeness, there remain intrinsic trade-offs between envy-freeness and efficiency. To this end, we derive the price of envy-freeness [5] which measures the ratio of the (unconstrained) efficient allocation to the efficient allocation subject to ex-ante envy-freeness. Principally, we analyse the complexity of maximising ex-ante efficiency subject to ex-ante envy-freeness. Finally, we devise an integer program for linear satiable valuation functions using linearisation techniques.

Contributions. We present a new problem of fair division of a homogeneous resource with uncertain availability. For this setting, we show that:

(1) Ex-ante allocations are ex-ante efficient if and only if the allocations are ex-post efficient (Lemma 4.1). This means that efficient allocations can be easily calculated for reasonable valuation functions.
(2) Ex-post envy-free allocations are also ex-ante envy-free (Lemma 5.2). However, the opposite is not necessarily true which means that there are settings where ex-ante envy-free allocations can achieve a higher ex-ante efficiency than ex-post envy-free allocations (see Example 1.1).
(3) The social welfare of the ex-ante efficient allocation under ex-ante envy-freeness can be substantially smaller than the welfare of the (unconstrained) ex-ante efficient allocation. To be precise, the price of envy-freeness has a lower bound of $\Omega(n)$ (Theorem 6.4), where n is the number of agents, which is asymptotically tight for concave valuation functions (Theorem 6.5).
(4) The problem of maximising the ex-ante social welfare under ex-ante envy-freeness is strongly NP-hard even for continuous and concave valuation functions, and uniform probabilities (Theorem 7.2).
(5) We devise an integer program to calculate the optimal ex-ante envy-free allocation for linear but satiable functions (see Equations 1 - 4 in Section 8).

The first three results are obtained by relatively straightforward arguments and/or constructions. The insights they provide are nevertheless valuable. The proof of the intractability is rather involved and is the most interesting from a technical perspective.

In the remainder of this paper we discuss related work, formally introduce the model, and provide our results in the aforementioned order. Due to space limitations some proofs are omitted or sketched. Most of these proofs are either straightforward to obtain or technical but the meaning of the statements are intuitively clear and sufficient to explain the proofs of the main theorems.

2 RELATED WORK

Computational social choice is a prosperous research field and fair division is certainly within its core [6]. However, our problem is related to and has features of a wide array of problems and areas. From within the field of fair division these include cake cutting [28, 29] and estate or land division [27], and, from related fields, divisible auctions, divisible task scheduling [17, 25] and packing problems [30]. We have motivated our model by the requirements and constraints of local energy markets with renewable energy sources but this model can be applied to other important areas including emission permits for greenhouse gases [1], fair load shedding [26], and uncertain computational resources [3, 18].

As mentioned, the work of Feige and Tennenholtz [12] on a single homogeneous divisible resource is the most closely related to our work. However, in contrast to our consideration of envy-freeness, they are using a fairness criterion that adapts the concept of proportionality which ordinarily means each agent gets at least one n-th of his total utility. Additionally, there are other works on divisible goods, where these usually consider several divisible items [6, 20, 27] or auctioning the divisible resource [19, 21].

Gajdos and Tallon [15] study the relationship of ex-ante and ex-post envy-freeness under ex-ante efficiency when agents have different perceptions of the availability. They focus on two very simple cases where the amount of resource is equal for all events. They show that, for their setting, ex-ante optimal allocations are included in ex-post optimal ones and that ex-post envy-free allocations are a subset of ex-ante envy-free solutions. In contrast, we show that ex-post and ex-ante efficiency are the same when we do not require envy-freeness, and quantify the degradation of ex-ante efficiency from ex-ante envy-freeness as well as the complexity of calculating ex-ante envy-free efficient allocations.

The important area of cake cutting, in contrast to our homogeneous resource, considers a heterogeneous resource. Bei et al. [4] show that, for the measure of proportionality, efficiency is NP-hard to approximate within a factor of $\Omega(1/\sqrt{n})$ for general piecewise constant valuation functions and they give a PTAS for linear functions. Furthermore, Chen et al. [9] develop fair cake cutting algorithms focusing on piecewise uniform valuation functions.

Another extensive field is that of indivisible goods where whole items have to be assigned [31]. However, this area is different as envy-free allocations do not have to exist, and determining their existence is already hard [6].

For the welfare loss of envy-free allocations we use the price of envy-freeness. This measure was introduced by Bertsimas et al. [5] for resource allocation problems. Furthermore, it has been used to quantify the degradation of efficiency in cake cutting by, for example, Aumann and Dombb [2], and [8].

Finally, the problem considered in our work has similarities to several other areas. These include estate division [27] where parties have claims on the items, and land division [29] which is cutting across two dimensions. Furthermore, in scheduling, especially with divisible tasks, jobs have to be assigned to machines which is similar to agents getting resource from the events [10, 13, 25]. Additionally, allocating the resources is also close to packing problems [30], although this usually considers equal sized containers which differ from our events with different available amounts.

3 PRELIMINARIES

There are $n \in \mathbb{N}$ agents which are interested in a resource whose actual amount is uncertain. The uncertain amount of the resource is represented by a random variable $X \in [0,1]$ ($X : \Omega \to [0,1]$) with a finite number of events $m := |\Omega|$ with $\Omega \subsetneq [0,1]$ and probability mass function f. Note that, for simplicity, we have overloaded our notation for events and let $\omega := X(\omega)$ for $\omega \in \Omega$. The *allocation* of the resource to the agents is represented by the vector $A = (a_1, a_2, \ldots, a_n)$ where the allocation to an agent i is a function $a_i : \Omega \to [0,1]$. An allocation A is *valid* if it satisfies these two validity constraints: positivity that implies $a_i(\omega) \geq 0 \; \forall i \in [n], \forall \omega \in \Omega$, and respecting the maximal amount that implies $\sum_{i \in [n]} a_i(\omega) \leq \omega \; \forall \omega \in \Omega$. It is important to notice that allocation functions are conditional on the events. Let $\mathcal{F} \subset [0,1]^{\Omega}$ be the set of all valid allocation functions and let $\Lambda = \mathcal{F}^n$ be the set of all valid allocations.

An agent $i \in [n]$ values the amount of received resource according to a monotonically increasing *valuation function* $v_i : [0,1] \to \mathbb{R}$ which also satisfies non-negativity ($v(x) \geq 0 \; \forall x \in [0,1]$) and no valuation for zero ($v(0) = 0$). The monotonicity reflects that agents can drop excess resource which does not decrease their valuation. Let $\Theta \subset \mathbb{R}^{[0,1]}$ be the set of all valid valuation functions. Additionally, let $V_i : \mathcal{F} \to \mathbb{R}$ denote an agent i's *utility* given an allocation function which is equal to the expected valuation of the allocation function, that is, $V_i(a_j) := \sum_{\omega \in \Omega} v_i(a_j(\omega)) f(\omega)$ for any $j \in [n]$. We note here, for motivational reasons, the valuations are not scaled to not give all the agents the same weight. Moreover, normalisation is non-trivial in our setting and would not affect the negative results.

The goal is, given the agents' valuation functions, to find a valid ex-ante envy-free allocation A, that is an allocation where, in terms of utility, an agent weakly prefers his allocation over every other agent's allocation, $V_i(a_i) \geq V_i(a_j)$ for all $i, j \in [n]$, such that A maximises ex-ante social welfare $W(A) := \sum_{i \in [n]} V_i(a_i)$. This is in contrast to ex-post envy-freeness that implies $v_i(a_i(\omega)) \geq v_i(a_j(\omega))$ for all $i, j \in [n]$, and the maximisation of ex-post social welfare $\sum_{i \in [n]} v_i(a_i(\omega))$ with respect to one event $\omega \in \Omega$. We note that, henceforth, any reference to envy-freeness or social welfare without preposition refers to the respective ex-ante notion. Furthermore, an allocation of maximum social welfare is called *efficient*.

Finally, the price of envy-freeness is the ratio $\max_{\theta \in \Theta} \frac{W(A_E(\theta))}{W(A_{EF}(\theta))}$ where A_E is an unrestricted ex-ante efficient allocation and A_{EF} is an ex-ante envy-free and efficient solution. It expresses the degradation of efficiency due to the enforcement of ex-ante envy-freeness and a higher value indicates a higher efficiency loss.

4 EFFICIENT ALLOCATION

We begin by focusing on the efficiency of unconstrained allocations, i.e. when not requiring ex-ante envy-freeness. We show that any efficient allocation is also ex-post efficient and vice versa, which means that, under concave valuation functions, an efficient solution can be found in polynomial time. This gives us a reference for the price of envy-freeness and also for any future approximation.

Lemma 4.1. *An allocation is ex-ante efficient if and only if the allocation is ex-post efficient for every $\omega \in \Omega$.*

Proof Sketch. By linearity of expectation, ex-post efficiency implies ex-ante efficiency. For the reverse, we assume for contradiction that an allocation A is ex-ante efficient but not ex-post efficient for a number of events $\Psi \subseteq \Omega$. Using A and an allocation $A^e = (a_1^e, \ldots, a_n^e)$ that is ex-post efficient for every $\omega \in \Psi$, i.e. $\sum_{i \in [n]} v_i(a_i^e(\omega)) > \sum_{i \in [n]} v_i(a_i(\omega))$, we create allocation A' with $a_i'(\omega) = a_i(\omega)$ if $\omega \notin \Psi$ and $a_i'(\omega) = a_i^e(\omega)$ if $\omega \in \Psi$. Efficiency of A^e for all $\omega \in \Psi$ and linearity of expectation imply that A' has a higher social welfare than A. Hence, A cannot be ex-ante efficient. □

Therefore, we can construct an ex-ante efficient allocation by using the ex-post efficient allocations for every event. This calculation for one event can be represented as an optimisation problem with the ex-ante social welfare as the optimisation function and two linear constraints which require that the allocation does not exceed the available amount and that it is not negative for any agent. Hence, we have that, for concave valuations functions, the optimisation function is concave and the problem can be solved in polynomial time within some restrictions [7]. Since there are m events we can find the entire allocation in polynomial time.

5 ENVY-FREE ALLOCATION

In contrast to the unconstrained efficiency of the previous section, allocations for the envy-freeness constraint problem can vary significantly. Our result is that there is always an envy-free allocation where all resources are allocated, denoted as equal share allocation:

Definition 5.1 (Equal Share Allocation). The *equal share allocation* A_{ES} is the allocation with $a_i(\omega) = \frac{\omega}{n}$ for all $i \in [n]$ and $\omega \in \Omega$.

In order to establish that this allocation is envy-free we prove that, for envy-freeness, the set of ex-post allocations is a subset of the set of ex-ante allocations.

Lemma 5.2. *If an allocation function is ex-post envy-free for every $\omega \in \Omega$, it is also ex-ante envy-free.*

Proof. Let $A \in \Lambda$ be a valid allocation satisfying the lemmas statement, i.e. $v_i(a_i(\omega)) \geq v_i(a_j(\omega))$ for all $i, j \in [n]$ and $\omega \in \Omega$. Then, $V_i(a_i) = \sum_{\omega \in \Omega} v_i(a_i(\omega)) f(\omega)$
$\geq \sum_{i \in [n]} \sum_{\omega \in \Omega} v_i(a_j(\omega)) f(\omega) = V_i(a_j), \forall i, j \in [n]$. □

Note, the opposite is not necessarily true and this is exactly why it is possible to find allocations with increased ex-ante social welfare. Additionally, since the equal share allocation is trivially ex-post envy-free, it is also ex-ante envy-free by Lemma 5.2.

Corollary 5.3. *The equal share allocation is ex-ante envy-free.*

6 PRICE OF ENVY-FREENESS

We next consider the extent to which both efficiency and envy-freeness can be achieved. Specifically, we show in Theorem 6.4 that, without restricting the valuation functions, the price of envy-freeness is at least in the order of the number of agents. This follows since, in the case of linear and non-equal valuation functions, the most efficient solution is to give everything to one agent (see Lemma 6.2), and that, under ex-ante envy-freeness, it is not possible to achieve a higher efficiency than the equal share allocation (see Lemma 6.3). Finally, we show that the bound is asymptotically tight for concave valuation functions (see Theorem 6.5).

For this section, we assume all valuation functions are linear, i.e. $v_i(x) = c_i \cdot x$ with $c_i \in \mathbb{R}_+$ for $i \in [n]$. Then, an ex-ante efficient allocation, which we call *maximal slope allocation*, gives everything to an agent with the highest slope. The ex-ante social welfare is a function of the maximal slope and the probability distribution.

Definition 6.1 (Maximal Slope Allocation). A *maximal slope allocation* is the allocation where an agent $j \in \arg\max_{i \in [n]}\{c_i\}$ gets all of the resource.

Lemma 6.2. *A maximal slope allocation is efficient and has an ex-ante social welfare of $c_j \cdot \mathrm{E}[X]$ for $j \in \arg\max_{i \in [n]}\{c_i\}$.*

In comparison, the maximum possible social welfare under envy-freeness is achieved by the equal share allocation and is a function of the average of the slopes of all agents. This demonstrates that there is an inherent trade-off between social welfare and envy-freeness.

Lemma 6.3. *No ex-ante envy-free allocation can have more ex-ante social welfare than $\mathrm{mean}_{i \in [n]}\{c_i\} \cdot \mathrm{E}[X]$ which is matched by the equal share allocation.*

Proof sketch. The key point is that under linearity of the valuations ex-ante envy-freeness implies that the expected allocations are the same, i.e. $\mathrm{E}[a_i] = \mathrm{E}[a_j]$ for all $i, j \in [n]$. Hence, the ex-ante social welfare is determined by the allocations and the average slope, i.e. $W(A) = n \cdot \mathrm{mean}_{i \in [n]}\{c_i\} \cdot \mathrm{E}[a_k]$ for any $k \in [n]$. Then, since every allocation is limited by the available amount, we have that $n \cdot \mathrm{E}[a_k] \leq E[X]$ which implies the claimed limit on the achievable ex-ante social welfare. It is straightforward to show that this is matched by the equal share allocation. □

Finally, we construct an instance for the price of envy-freeness.

Theorem 6.4. *The division of a homogeneous resource has a price of envy-freeness of $\Omega(n)$.*

Proof. Let X be arbitrary but have at least one event with positive amount and probability. Let the valuation function of the first agent be $v_1(x) = 2x$ and the remaining valuation functions $v_i(x) = \frac{1}{n-1} \cdot x$ for $i \in \{2, \ldots, n\}$. Then, the ratio of the efficient to the envy-free allocation is, by Lemma 6.2 and 6.3, $(c_j \cdot \mathrm{E}[X])/(\mathrm{mean}_{i \in [n]}\{c_i\} \cdot \mathrm{E}[X]) = 2/3 \cdot n$ which implies the claim. □

For arbitrary valuation functions the upper bound can be unbounded. However, for the realistic assumption of concave valuation functions the bound is asymptotically tight.

Theorem 6.5. *For concave valuation functions, the division of a homogeneous resource has a price of envy-freeness of at most n.*

The results of the theorems 6.4 and 6.5 are the same as the bounds for every event ex-post. However, the results of the theorems extend the ex-post results by showing that irrespective of the probability distribution, the results also holds ex-ante. In other words, the bounds confirm that there is still no worst-case improvement.

7 COMPLEXITY

The results from the previous section show that, for linear valuation functions, equal share already attains maximum possible efficiency. However, Example 1.1 shows that, if the valuation functions of some of the agents are non-linear, there are allocations with higher ex-ante social welfare. Hence we ask: how difficult would it be to find more efficient allocations in general. We show that the problem of maximising ex-ante social welfare under ex-ante envy-freeness is strongly NP-hard. In order to prove this we consider the decision version of our problem which we call *decision version of uncertain amount fair division (D-UAFD).*

Definition 7.1 (Decision Version of Uncertain Amount Fair Division (D-UAFD)).

Instance: $\langle X, (v_i)_{i \in [n]}, B \rangle$ where $X : \Omega \to [0, 1]$ with finite $\Omega \subsetneq [0, 1]$ is a discrete random variable, $(v_i)_{i \in [n]}$ with $v_i \in \Theta$ are valuation functions and $B \in \mathbb{R}$ is a bound.

Problem: Does there exist a valid allocation $A = (a_i)_{i \in [n]} \in \mathcal{F}^n$ such that $V_i(a_i) \geq V_i(a_j)$ for all $i, j \in [n]$ and $W(A) \geq B$.

Θ, $\mathcal{F}$, V_i and $W(A)$ are defined as in the preliminaries.

Theorem 7.2. *D-UAFD is NP-complete in the strong sense.*

We prove this by reducing from the 3-partition problem, that is known to be NP-complete in the strong sense [16].

Definition 7.3 (3-Partition Problem).

Instance: $\langle S, B \rangle$, where $S \subset \mathbb{N}$ is a finite multiset of $3m$ elements $s_1 \dots s_{3m}$ and $B \in \mathbb{Z}^+$ is a bound such that $B/4 < s_i < B/2 \; \forall i \in [3m]$ and $\sum_{i \in [3m]} s_i = mB$.

Problem: Can S be partitioned into m disjoint subsets $S_1, \dots, S_m$ such that $\sum_{x \in S_i} x = B$ for all $i \in [m]$.

(Notice that the item size requirements imply that each subset must contain exactly 3 elements.)

We reduce a given instance of 3-partition problem to an instance of our fair distribution problem which we call *envy partition instance*. In short, we scale the 3-partition instance down and transform it into an instance of D-UAFT where the sets are the events and the 3-partition elements are the agents. To achieve this, the agents' valuation functions are chosen so that the opposing goals of envy-freeness and efficiency require that every agent gets allocated exactly the amount specified by the corresponding 3-partition element in a single event. Therefore, allocating this amount to the agents corresponds to choosing a set for the corresponding 3-partition element. Additionally, to address that every event has to have a different amount, we slightly increase the size of the events and add additional agents who desire exactly these amounts.

Definition 7.4 (Envy Partition Instance). An *envy partition instance* $\langle X, (v_i)_{i \in [n]}, B' \rangle$ is an instance of D-UAFD constructed given a 3-partition instance $\langle S, B \rangle$ in the following way. We assume without loss of generality that the 3-partition elements are ordered in a non-decreasing way, i.e. $s_1 \le s_2 \le \dots \le s_{3m}$.

Firstly, there are $4m$ agents in total; i.e. $n = 4m$. The agents $m+1, \dots, 4m$ correspond to the elements of the 3-partition instance. The remaining m agents correspond to both the partitioning subsets of the 3-partition instance and the events as defined for this instance. In particular, let $s'_i := s_i/(mB)$ for all $i \in [3m]$, let ε be chosen such that $0 < \varepsilon < 2/(m^2 B(m+1))$, and let $Z := \sum_{i=1}^{m} i \cdot \varepsilon$. Then the valuation functions are defined as follows.

$$v_i(x) = \begin{cases} 0 & \text{if } x \le 0 \\ i \cdot x & \text{if } 0 < x \le \hat{s}_i \\ i \cdot \hat{s}_i & \text{if } x > \hat{s}_i \end{cases}$$

$$\text{with } \hat{s}_i = \begin{cases} i \cdot \varepsilon & \text{if } i \le m \\ s'_{i-m} & \text{if } m+1 \le i \le 4m \end{cases}$$

for $i \in [4m]$. We refer to $\hat{s}_i$ as agent i's saturation amount. Note that $\hat{s}_i$'s are weakly increasing in i.

The second step is the construction of the random variable X. The set of events consists of m elements which are of size $\frac{1}{m}$ plus an additional small and increasing offset. Formally, $\Omega = \{\omega_1, \dots, \omega_m\} = \left\{\frac{1}{m} + \varepsilon, \frac{1}{m} + 2 \cdot \varepsilon, \dots, \frac{1}{m} + m \cdot \varepsilon\right\}$. Additionally, our random variable is uniformly distributed, i.e. has the probability mass function $f(\omega) = \frac{1}{|\Omega|} = \frac{1}{m}$ for all $\omega \in \Omega$. Note that $\sum_{\omega \in \Omega} \omega = 1 + Z = \sum_{i \in [4m]} \hat{s}_i$.

Finally, $B' = \frac{1}{m} \cdot \sum_{i=1}^{3m} i \cdot s'_i + \frac{\varepsilon}{6}\left(2m^2 + 3m + 1\right) + 1$.

This construction can be done in polynomial time, the constructed instance is of polynomial size, and the constructed values are polynomial in the values of the 3-partition instance.

Given the constructed envy partition instance, we first show that an allocation A achieves the required social welfare B' and satisfies ex-ante envy-freeness only if it allocates to each agent, except maybe agent 1, his saturation amount in exactly one event and zero in all other events; that is, $a_i(\omega') = \hat{s}_i$ for one $\omega' \in \Omega$ and $a_i(\omega) = 0$ for all $\omega \in \Omega \backslash \{\omega'\}$. We then show that such an allocation, satisfying envy-freeness and having social welfare of at least B', exists if and only if a 3-partition exists.

Intuitively speaking, to maximise social welfare we need to give the agents with higher index larger allocations, as they have higher valuations, but we cannot allocate to them too much as then the agents with lower index will envy them.

For the purpose of our proof, we introduce a new representation of an allocation which we call *allocation by pieces*. In this representation, we allocate up to the entire amount of resource over all events $\sum_{\omega \in \Omega} \omega = 1 + Z$ to agents as *pieces*, independent of the events. Each agent receives m pieces. We then map events to pieces.

Definition 7.5 (Allocation by Pieces). An allocation by pieces consists of functions $a_i^p : [m] \to [0, 1]$ for each agent $i \in [4m]$ such that $0 \le a_i^p(j) \le \hat{s}_i, \forall j \in [m]$ and $\sum_{i \in [n]} \sum_{j \in [m]} a_i^p(j) \le 1 + Z$, and a function $\Phi : [n] \times \Omega \to [m]$ where $\Phi(i, \omega) \ne \Phi(i, \omega'), \forall i \in [n], \forall \omega \ne \omega' \in \Omega$.

Social welfare and envy-freeness for a_i^p's are defined as in the preliminaries. Furthermore, an allocation by pieces is valid if $a_i^p(\Phi(i, \omega))$ is valid $\forall i \in [n], \forall \omega \in \Omega$, where validity is defined as in the preliminaries. Given a_i^p's and given a Φ, we say that Φ is valid if $a_i^p(\Phi(i, \omega))$'s are valid.

We will show that taking any allocation by pieces and modifying it so that the pieces for a given agent i are reduced to a number of pieces of size $\hat{s}_i$ and maximum one piece of size less than $\hat{s}_i$ (see Definition 7.7) does not change social welfare (see Lemma 7.8). Furthermore, if a_i^p's are ex-ante envy-free and a valid Φ exists then $a_i^p(\Phi(i, \omega))$'s are ex-ante envy-free (see Lemma 7.9).

Note that an agent's valuation in any given event does not increase after reaching his saturation amount. Therefore, reducing his allocation to his saturation amount and changing nothing else maintains the social welfare and envy-freeness.

Proposition 7.6. *It suffices to consider allocations A with $a_i(\omega) \le \hat{s}_i$ for all $i \in [n]$.*

From now on we only consider allocations in which each agent receives in each event at most his saturation amount.

Any allocation $A = (a_1, \dots, a_n)$ is equivalent to the allocation by pieces A^p with $a_i^p(j) = a_i(\omega_j)$ and $\Phi(i, \omega_j) = j$. Considering the whole amount of the allocation, we can observe that the pieces can be reduced to a number of pieces allocating exactly the saturation amount and at most one piece with less than the saturation amount. This means that we can create a new allocation by pieces which allocates to each agent the same total amount as in A. We denote this allocation as *total amount allocation*.

Definition 7.7 (Total Amount Allocation). An agent i's allocation can be represented as a total amount allocation $A_i = n_i \cdot \hat{s}_i + d_i$ with $0 \le d_i < \hat{s}_i$. This corresponds to the allocation in pieces $a_i^p(j) = \hat{s}_i$ for $j \in [n_i]$ with $n_i \in ([m] \cup \{0\})$ and $a_i^p(n_i + 1) = d_i$.

Given an allocation by pieces, a total amount allocation can be created in a stepwise manner. For example, assume that agent i has two pieces $(1, 2)$ below his saturation amount and that his allocation under 1 is weakly less than his allocation under 2; i.e. $a_i^p(1) \le a_i^p(2) < \hat{s}_i$. Then it is possible to reduce the allocation by $d := \min\{a_i^p(2), \hat{s}_i - a_i^p(1)\}$ in piece 2 ($a_i^{p'}(2) := a_i^{p'}(2) + d$) and increase it in piece 1 ($a_i^{p'}(1) := a_i^{p'}(1) + d$).

We remark again that we ignore Φ here. Generally, a valid Φ does not have to exist for all total amount allocations. However, the aim is to show that for any allocation to achieve a social welfare of B', any agent i's allocation has to be exactly $\hat{s}_i$ in one event and zero in all other events. This is independent of the choice of the event in which the agent gets his saturation amount. In this respect, the total amount allocation is suitable since its social welfare is the same as the social welfare of the allocation it is based on, and if the original allocation is ex-ante envy-free and a valid Φ exists, then the allocation $(a_i^p(\Phi(i, \omega)))_{i \in [n], \omega \in \Omega}$ is also ex-ante envy-free. The social welfare is unaffected since the valuation functions are linear.

Lemma 7.8. *The ex-ante social welfare of an allocation and its representing total amount allocation are equivalent.*

The envy-freeness of $(a_i^p(\Phi(i, \omega)))_{i \in [n], \omega \in \Omega}$ can be argued with the shifts to create a total amount allocation. An agent i is indifferent between the allocation before and after a shift since he receives the same amount in total and all events are equally likely. Every other agent j's utility of agent i's allocation after the shift is weakly less than before the shift, since it is the same amount but all utilities are subject to agent j's saturation amount. Hence, the difference in the ex-ante envy-freeness inequalities can only increase.

Lemma 7.9. *If a_i^p's are ex-ante envy-free then the total amount allocation is ex-ante envy-free.*

Lemma 7.10. *If a_i^p's are ex-ante envy-free and a valid Φ exists then $(a_i(\Phi(i, \omega)))_{i \in [n], \omega \in \Omega}$ is ex-ante envy-free.*

The following lemmata are used in the shifting procedure in the proof of Theorem 7.2. If any shift based on the following lemmata ends in an allocation for which a valid Φ exists then by Lemmata 7.9 and 7.10 we have reached an envy-free allocation.

Lemma 7.11. *Fixing an allocation A and an agent $1 \le k < 4m$, if all agents i where $i < k$ have a total amount allocation of $A_i = \hat{s}_i$ and agent k's total amount allocation is $A_k > \hat{s}_k$, then shifting allocation to create allocation A' such that any excess is shifted to the next agent, i.e. $A'_k = \hat{s}_k$ and $A'_{k+1} = A_{k+1} + A_k - \hat{s}_k$, increases social welfare.*

Envy-freeness stipulates that, if an agent receives less than his saturation amount, then all agents of higher index also receive less than their saturation amounts.

Lemma 7.12. *For two agents $i, k \in [4m]$ with $i > k$, if agent k has a total amount allocation of less than, or equal to, his saturation amount, i.e. $A_k \le \hat{s}_k$, then the same must be true for agent i, i.e. $A_i \le \hat{s}_i$. The same holds for the strict case, i.e. if $A_k < \hat{s}_k$ then $A_i < \hat{s}_i$.*

Further, envy-freeness imposes the following conditions on Φ.

Lemma 7.13. *If for an envy-free allocation every agent has an allocation of once his saturation amount then agents $2, \ldots, 4m$ have to have their saturation amount in exactly one event or piece.*

Proof sketch. If the allocation of $\hat{s}_i$ for $i \in [n] \backslash \{1\}$ is in one event, the utility of agent 1 for the allocation of agent i will be $\hat{s}_1$ (the remainder $\hat{s}_i - \hat{s}_1$ is of no value to him). If $\hat{s}_i$ is split between two or more events, then agent 1 will receive some additional utility from $\hat{s}_i - \hat{s}_1$ and hence his utility for agent i's allocation will be more than $\hat{s}_1$, implying that agent 1 envies agent i's allocation. □

Finally, we are able to prove the main theorem.

Proof of Theorem 7.2. It is easy to see that given an allocation, we can calculate the social welfare and verify the envy-freeness in polynomial time. Therefore D-UAFD is in NP. To prove NP-completeness, we show that the constructed envy partition instance (see Definition 7.4) is a yes-instance if and only if the given 3-partition instance is a yes-instance.

First, consider the case where there is a valid partition $S_1, \ldots, S_m$ for the 3-partition instance. Consider the allocation A where for $s_i \in S_j$ agent $i + m$ is allocated $\hat{s}_{i+m}$ in ω_j for all $i \in [3m]$ and for $i \in [m]$ agent i gets assigned $\hat{s}_i$ in ω_i. Recall that $\hat{s}_{i+m} = s'_i = \frac{s_i}{mB}$ $\forall i \in [3m]$, and $\hat{s}_i = i.\varepsilon$ $\forall i \in [m]$. It is thus easy to verify that allocation A is valid as $\sum_{s_i \in S_j} s'_i = \sum_{s_i \in S_j} \frac{s_i}{mB} = \frac{1}{m}$. The utility for an agent $i \in [n]$ with respect to his own allocation is $V_i(a_i) = \frac{i}{m} \cdot \hat{s}_i$, and his utility with respect to the allocation of agent $j \in [n]$ is $V_i(a_j) = \frac{i}{m} \cdot \hat{s}_i$ if $j > i$ and $V_i(a_j) = \frac{i}{m} \cdot \hat{s}_j \le \frac{i}{m} \cdot \hat{s}_i$ if $i > j$. Hence, A is ex-ante envy-free. Furthermore, the sum of the utilities is $\sum_{i \in [n]} V_i(a_i) = \frac{1}{m} \left(\sum_{i=1}^{3m} (i + m) \cdot s'_i + \sum_{i=1}^{m} i^2 \cdot \varepsilon \right) = B'$, which concludes this case.

Now, consider the case where there is an envy-free allocation A with social welfare at least B' for the constructed envy partition instance. By Observation 7.6, we can assume that every agent receives in every event at most his saturation amount. Therefore, A can be represented as an allocation by pieces (Definition 7.5), and hence has a total amount allocation representation $(A_1, \ldots, A_{4m})$ (Definition 7.7) that has the same ex-ante social welfare as A (by Lemma 7.8) and is ex-ante envy-free (by Lemma 7.9). This gives us a framework in which we show that, independent of Φ, to achieve B' every agent, except maybe agent 1, has to get his saturation amount exactly once; i.e. $a_i(\omega') = \hat{s}_i$ for one $\omega' \in \Omega$ and $a_i(\omega) = 0$ for all $\omega \in \Omega \setminus \{\omega'\}$, for all agents i, $2 \le i \le 4m$.

Starting from the total amount allocation $(A_1, \ldots, A_{4m})$ and considering agents one by one in increasing order of indices, we construct another allocation A' using the following procedure. If the current agent i is agent $4m$ or has a total amount allocation $A_i < \hat{s}_i$, we stop. If the current agent i has a total amount allocation $A_i > \hat{s}_i$ the additional amount is shifted to the next agent.

If the procedure reaches agent $4m$ and no amount has been shifted, then for each agent $i < 4m$ we have that $A_i = \hat{s}_i$ and, by Lemma 7.12, we have that $A_{4m} \le \hat{s}_{4m}$. An allocation where every agent is allocated exactly his saturation amount has social welfare of B'. Thus, if $A_{4m} < \hat{s}_{4m}$ then $W(A) < B'$ which contradicts our assumption. If $A_{4m} = \hat{s}_{4m}$ then, by Lemma 7.13, every agent in A, except maybe agent 1, is allocated his saturation amount exactly once in one event. If the procedure reaches agent $4m$ and some amount has been shifted during the procedure, then by Lemma 7.11, $W(A') > W(A)$. Furthermore, each agent $i < 4m$ is allocated exactly his saturation amount in A' and the total amount of resource available stipulates that agent $4m$ is allocated at most his saturation amount. An allocation where every agent is allocated exactly his saturation amount has social welfare of B'. Therefore $W(A') \leq B'$, and hence $W(A) < B'$ which contradicts our assumption. If the procedure stops at an agent $i < 4m$, then we have that (1) every agent $j < i$ is allocated exactly his saturation amount in A', and (2) since $A_i < \hat{s}_i$ then, by Lemma 7.12, for each agent $j \geq i$ we have that $A_j < \hat{s}_j$ and hence each such agent is allocated less than his saturation amount in A'. An allocation where every agent is allocated exactly his saturation amount has social welfare of B'. Therefore $W(A') < B'$. By Lemma 7.11, $W(A') \geq W(A)$, and hence $W(A) < B'$ which contradicts our assumption. Hence we can conclude that either each agent, except maybe agent 1, is allocated his saturation amount exactly once, or $W(A) < B'$ which is a contradiction.

So far we have established that allocation A is equivalent to the allocation by pieces A^p where $a_i^p(1) = \hat{s}_i$ and $a_i^p(j) = 0$ for $j \in ([m]\backslash\{1\})$, for all $i \in [4m] \setminus \{1\}$. In the remainder of the proof we show that a valid Φ exists only if there exists a 3-partition and, moreover, agent 1 too receives his saturation amount exactly once in A. Since by assumption A is valid, hence a valid Φ must exist and thus a 3-partition exists and, moreover, every agent receives his saturation amount exactly once in A.

For easier presentation, we multiply saturation amounts and the amount of events by mB which results in saturation amounts of agents $m + 1 \ldots 4m$ to be equal to 3-partition elements, which are all integers, and the amount of resource in each event $j \in [m]$ to be equal to $B + B \cdot m \cdot j \cdot \varepsilon$. Recall that integer B denotes the size of each set. Furthermore, by the choice of ε, $B \cdot m \cdot j \cdot \varepsilon < 1$, $\forall j \in [m]$.

We now investigate what must hold for a valid Φ to exist. We first claim that each agent i with $2 \leq i \leq m$ must be assigned to event i. For if not, then there must be an event $j \in [m]$ with a remaining unallocated amount of less than B. Hence, since no event has an amount of $B + 1$ or more, this implies that we have a total amount of less than mB that we can assign to agents $m + 1 \ldots 4m$. But then these agents' saturation amounts add up to mB so it impossible to have these agents assigned. A similar argument holds if agent 1 is not assigned to event 1, and hence we conclude that each agent $i \in [m]$ is assigned to event i. Therefore, the remaining amount in each event is exactly B. We have that the saturation amounts of agents $m + 1 \ldots 4m$ are equal to their corresponding 3-partition elements, hence the existence of a valid Φ implies that there exists a 3-partition, which can be directly derived from Φ.

Therefore, the envy partition instance is a yes-instance if and only if the 3-partition instance is a yes-instance. Finally, since the envy partition instance can be constructed in polynomial time and is of polynomial size this concludes the strong completeness. □

8 INTEGER PROGRAM

Unfortunately, the instances of both the NP-hardness proof and the price of envy-freeness together imply that there is no reasonable relaxation which allows easy solutions. Essentially, linear functions do not allow any increase above equal share and, in all other cases, the envy-freeness is a non-linear constraint. Nevertheless, in this section we consider linear satiable valuations and formulate an integer program to calculate the efficient envy-free allocation.

It is straightforward to represent the problem in this work as an optimisation problem. However, this program, depending on the valuation functions, may be non-linear and non-concave. In the setting of linear satiable valuation functions we can rewrite the utility functions and envy-freeness constraint with minima functions. These can be transformed into a series of constraints to reformulate the mathematical program into an integer program to calculate the optimal envy-free solution.

More explicitly, the *linear satiable valuation* for agent $i \in [n]$ is defined as $v_i(x) = {}^{u_i}/_{q_i} \cdot x$ if $x \leq q_i$ and $v_i(x) = u_i$ otherwise, with *saturation amount* $q_i \in [0, 1]$ and *maximal value* $u_i \in \mathbb{R}^+$.

Like in the proof of Theorem 7.2 the utility of any allocation depends on the minimum of the agent's saturation amount and the allocated amounts. This fact allows us to rewrite the utility as well as the equations representing envy-freeness.

Proposition 8.1. *The utility for agent $i \in [n]$ with $(u_i, q_i) \in ([0, 1], \mathbb{R}^+)$ is* $V_i(a) = \frac{u_i}{q_i} \cdot \sum_{\omega \in \Omega} \min\{a(\omega_j), q_i\} f(\omega)$.

Proposition 8.2. *The envy-freeness constraint for agents $i, k \in [n]$ is represented by the equation*

$$EF(i, k) := \sum_{\omega \in \Omega} (\min\{a_i(\omega), q_i\} - \min\{a_k(\omega), q_i\}) f(\omega) \geq 0$$

Altogether, we can formulate the problem as the following optimisation program with decision variables x_{ij} for $i \in [n]$ and $j \in [m]$.

$$\max \quad \sum_{i \in [n]} V_i \left((x_{ij})_{j \in [m]}\right) \tag{1}$$

$$\text{s.t.} \quad \sum_{i \in [n]} x_{ij} \leq \omega_j \qquad \forall j \in [m] \tag{2}$$

$$EF(i, k) \geq 0 \qquad \forall i, k \in [n] \tag{3}$$

$$x_{ij} \geq 0 \qquad \forall i \in [n], j \in [m] \tag{4}$$

The optimisation function is the social welfare rewritten using Proposition 8.1. The first constraint limits the allocation x_{ij} for agent $i \in [n]$ in event $j \in [m]$ to the maximal available amount. The second constraint is envy-freeness and the last constraint ensures only positive allocations are attained.

In this formulation neither the optimisation function nor the envy-freeness constraint appear linear. However, in a series of replacements we can replace those with linear constraints and integer variables.

Firstly, similar to Proposition 7.6 allocating more to an agent than the saturation amount does not increase the value. Furthermore, it can only negatively affect envy-freeness since the agents' valuations do not increase but another agent might be envious of the increased amount. Consequently, we can replace one of the minimum functions in the optimisation program.

Lemma 8.3. *The expression* $\min\{x_{ij}, q_i\}$ *can be replaced with* x_{ij} *and constraint* $x_{ij} \leq q_i$ *for all* $i \in [n]$ *and* $j \in [m]$.

The other minimum in the envy-freeness equation, $\min\{x_{kj}, q_i\}$, cannot be replaced that easily. We apply linearisation techniques and replace the minimum function with three more types of variables and a number of constraints. An overview of linearisation techniques can be found for example in the work by Liberti [23].

Firstly, we substitute the minimum function $\min\{x_{kj}, q_i\}$ with a variable x^i_{kj}. Secondly, we use a second integer variable y^i_{kj} to ensure that the substitution is valid.

Lemma 8.4. *The equation* $x^i_{kj} = \min\{x_{kj}, q_i\}$ *for all* $i, k \in [n], j \in [m]$ *holds for constraints* $q_i \cdot y^i_{kj} \leq x^i_{kj} \leq q_i$ *and* $x_{kj} \cdot (1 - y^i_{kj}) \leq x^i_{kj} \leq x_{kj}$ *with* $y^i_{kj} \in \{0, 1\}$.

The general idea behind the technique is as follows: for the equation to hold, the variable x^i_{kj} has to be smaller than both values in the minimum; yet at the same time it also has to be greater than one of the two, i.e. be exactly of that value. By setting y^i_{kj} to one or zero we can tighten one or the other constraint to be of the minimal value. Considering the four cases of the two possible values of the minimum and the two possible values of y^i_{kj}, one can see that y^i_{kj} can only be chosen so that the equation $x^i_{kj} = \min\{x_{kj}, q_i\}$ holds.

However, the constraint $x_{kj} \cdot (1 - y^i_{kj}) \leq x^i_{kj}$ still contains a non-linear expression in the form of a product. We replace that product $x_{kj} \cdot y^i_{kj}$ with a new variable z^i_{kj} which yields constraint $x^i_{kj} \geq x_{kj} - z^i_{kj}$ and further constraints.

Lemma 8.5. *The equation* $z^i_{kj} = x_{kj} \cdot y^i_{kj}$ *for all* $i, k \in [n], j \in [m]$ *holds for constraints* $0 \leq z^i_{kj} \leq x_{kj}$ *and* $x_{kj} + y^i_{kj} - 1 \leq z^i_{kj} \leq y^i_{kj}$.

The main observations for the linearisation are that $x_{kj} \in [0, 1]$ and that y^i_{kj} is binary. This implies that z^i_{kj} has to be smaller than both factors. Furthermore, the constraint $x_{kj} + y^i_{kj} - 1$ is either x_{kj} or non-positive and therefore assures that the result holds.

Finally, replacing the minimum functions with the variables and constraints from the Lemmata 8.3, 8.4 and 8.5 yields an integer linear program. We can verify this by observing that all of the resulting constraints are linear and we have continuous variables x_{kj}, x^i_{kj} and z^i_{kj} as well as the binary variables y^i_{kj}. Due to space limitations we omit the full integer program here.

8.1 Example

We ran a preliminary experiments on random instances to see the efficiency of the ex-ante envy-free solution. Firstly, the main constraint of the optimisation program is the envy-freeness. Computationally, this is apparent from the fact that increasing the number of agents increases the runtime significantly more than increasing the number of events. We selected one example (see Figure 2) where we show the social welfare of the equal share allocation, the unconstrained efficient allocation and the ex-ante efficient envy-free allocation for 30 events and 2 to 9 agents. The events are drawn uniformly from $(0, 1]$, the agents' maximal values are drawn uniformly from $[1, 20]$ and, in the case of satiable valuations, the agents' saturation amounts are drawn uniformly from $(0, 1]$. The plot shows that the social welfare difference of the unconstrained efficient allocation and the envy-free allocations increases. However, in the case of satiable valuations, in comparison to the linear valuations case, the social welfare is increasing for the envy-free allocations.

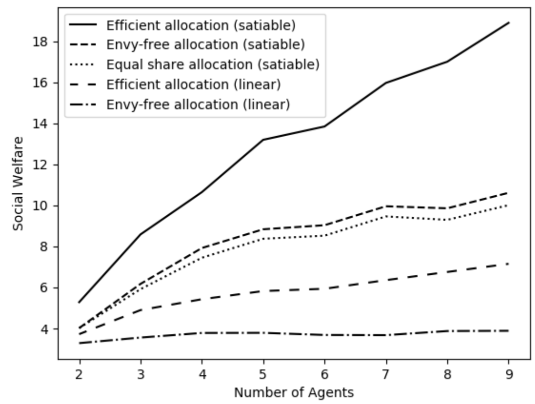

Figure 2: The social welfare of the equal share allocation, the efficient ex-ante envy-free allocation and the unconstrained efficient allocation for 30 **events and an increasing number of agents with linear or linear satiable valuation functions.**

9 CONCLUSIONS

We consider a fair division variant where the amount of a homogeneous resource is uncertain which is reflected by a random variable over a finite set of discrete events. We show that, while unconstrained efficiency optimisation can be solved in polynomial time for concave and other reasonable valuation functions, this is not no longer the case if ex-ante envy-freeness is required. In this case, an ex-ante envy-free allocation always exists but might have a significantly worse social welfare than the ex-ante efficient allocation. More specifically, the price of envy-freeness is tightly bounded by n for concave valuation functions, where n is the number of agents. Principally, we show that the problem of finding an ex-ante efficient allocation under ex-ante envy-freeness is strongly NP-complete, even under simple continuous valuation functions and with uniform probability over the events. Finally, we devise an integer program for the optimal ex-ante envy-free solution for linear satiable valuations.

The setting presented in this paper invites various directions for future work. Firstly, our NP-hardness result calls for a polynomial-time algorithm that approximates efficiency under ex-ante envy-freeness. Secondly, it would be interesting to investigate the price of envy-freeness in other restricted classes of valuation functions (besides linear valuation functions). Thirdly, we have assumed that the valuations are known. While this makes sense in some settings it would be of interest to examine the strategic case where agents can misrepresent their valuation.

ACKNOWLEDGMENTS

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) platform grant EP/R029563/1 and doctoral training grant EP/M508147/1.

REFERENCES

[1] Marcelo Wilson Furlan Matos Alves and Enzo Barberio Mariano. 2018. Climate justice and human development: A systematic literature review. *Journal of Cleaner Production* 202 (2018), 360–375. https://doi.org/10.1016/J.JCLEPRO.2018.08.091

[2] Yonatan Aumann and Yair Dombb. 2015. The Efficiency of Fair Division with Connected Pieces. *ACM Trans. Econ. Comput.* 3, 4 (2015), 23:1–23:16. https://doi.org/10.1145/2781776

[3] Yonatan Aumann, Yair Dombb, and Avinatan Hassidim. 2016. Auctioning Time: Truthful Auctions of Heterogeneous Divisible Goods. *ACM Trans. Econ. Comput.* 4, 1 (2016), 3:1–3:16. https://doi.org/10.1145/2833086

[4] Xiaohui Bei, Ning Chen, Xia Hua, Biaoshuai Tao, and Endong Yang. 2012. Optimal Proportional Cake Cutting with Connected Pieces. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, Vol. 2. 1263–1269.

[5] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2011. The Price of Fairness. *Operations Research* 59, 1 (2011), 17–31. https://doi.org/10.1287/opre.1100.0865

[6] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of Computational Social Choice.* Cambridge University Press, New York, NY, USA. 1–535 pages. https://doi.org/10.1017/CBO9781107446984

[7] Sébastien Bubeck. 2015. Convex Optimization: Algorithms and Complexity. *Foundations and Trends® in Machine Learning* 8, 3-4 (2015), 231–357. https://doi.org/10.1561/2200000050 arXiv:1405.4980

[8] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The Efficiency of Fair Division. *Theory of Computing Systems* 50, 4 (2012), 589–610. https://doi.org/10.1007/s00224-011-9359-y

[9] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2013. Truth, Justice, and Cake Cutting. *Games and Economic Behavior* 77, 1 (2013), 284–297. https://doi.org/10.1016/J.GEB.2012.10.009

[10] George Christodoulou, Elias Koutsoupias, and Annamária Kovács. 2010. Mechanism Design for Fractional Scheduling on Unrelated Machines. *ACM Transactions on Algorithms* 6, 2 (2010), 38:1 – 38:18. https://doi.org/10.1145/1721837.1721854

[11] Vincent P. Crawford. 2018. Fair Division. In *The New Palgrave Dictionary of Economics* (3rd ed.). Palgrave Macmillan UK, London, 4408–4410. https://doi.org/10.1057/978-1-349-95189-5_252

[12] Uriel Feige and Moshe Tennenholtz. 2014. On fair division of a homogeneous good. *Games and Economic Behavior* 87 (2014), 305–321. https://doi.org/10.1016/j.geb.2014.02.009

[13] Amos Fiat and Ariel Levavi. 2012. Tight Lower Bounds on Envy-Free Makespan Approximation. In *Internet and Network Economics*, Paul W. Goldberg (Ed.), Vol. 7695 LNCS. Springer Berlin Heidelberg, Berlin, Heidelberg, 553–558. https://doi.org/10.1007/978-3-642-35311-6_46

[14] Duncan K. Foley. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 7, 1 (1967), 45–98.

[15] Thibault Gajdos and Jean-Marc Tallon. 2002. Fairness under Uncertainty. *Economics Bulletin* 4, 18 (2002), 1–7. https://halshs.archives-ouvertes.fr/halshs-00086032

[16] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman & Co., New York, NY, USA.

[17] Yiannis Giannakopoulos, Elias Koutsoupias, and Maria Kyropoulou. 2016. The Anarchy of Scheduling Without Money. In *Algorithmic Game Theory*, Martin Gairing and Rahul Savani (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 302–314. https://doi.org/10.1007/978-3-662-53354-3_24

[18] Ramesh Johari and John N. Tsitsiklis. 2004. Efficiency Loss in a Network Resource Allocation Game. *Mathematics of Operations Research* 29, 3 (2004), 407–435. https://doi.org/10.1287/moor.1040.0091

[19] Ramesh Johari and John N. Tsitsiklis. 2009. Efficiency of Scalar-Parameterized Mechanisms. *Operations Research* 57, 4 (2009), 823–839. https://doi.org/10.1287/opre.1080.0638

[20] Frank Kelly. 1997. Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications* 8, 1 (1997), 33–37. https://doi.org/10.1002/ett.4460080106

[21] Ilan Kremer and Kjell G. Nyborg. 2004. Divisible-good auctions: the role of allocation rules. *RAND Journal of Economics* 35, 1 (2004), 147 – 159.

[22] Chaojie Li, Xinghuo Yu, Peter Sokolowski, Nian Liu, and Guo Chen. 2016. A Stochastic Game for Energy Resource Trading in the Context of Energy Internet. In *2016 IEEE Power and Energy Society General Meeting (PESGM).* IEEE, 1–5. https://doi.org/10.1109/PESGM.2016.7741891

[23] Leo Liberti. 2007. *Reformulation Techniques in Mathematical Programming.* Thèse d'Habilitation à diriger des Recherches. Université de Paris IX - Dauphine. http://www.lix.polytechnique.fr/Labo/Leo.Liberti/hdr.pdf

[24] Esther Mengelkamp, Benedikt Notheisen, Carolin Beer, David Dauer, and Christof Weinhardt. 2018. A Blockchain-Based Smart Grid: Towards Sustainable Local Energy Markets. *Computer Science - Research and Development* 33, 1-2 (2018), 207–214. https://doi.org/10.1007/s00450-017-0360-9

[25] Ahuva Mu'Alem. 2009. On Multi-dimensional Envy-Free Mechanisms. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory.* Springer-Verlag, Berlin, Heidelberg, 120–131. https://doi.org/10.1007/978-3-642-04428-1_11

[26] Olabambo Ifeoluwa Oluwasuji, Obaid Malik, Jie Zhang, and Sarvapali Ramchurn. 2018. Algorithms for Fair Load Shedding in Developing Countries. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI).* 1590–1596. https://doi.org/10.24963/ijcai.2018/220

[27] Dénes Pálvölgyi, Hans Peters, and Dries Vermeulen. 2014. A strategic approach to multiple estate division problems. *Games and Economic Behavior* 88 (2014), 135–152. https://doi.org/10.1016/j.geb.2014.09.005

[28] J. H. Reijnierse and J. A. M. Potters. 1998. On finding an envy-free Pareto-optimal division. *Mathematical Programming* 83, 1 (1998), 291–311. https://doi.org/10.1007/BF02680564

[29] E. Segal-Halevi, A. Hassidim, and Y. Aumann. 2015. Envy-free Cake-cutting in Two Dimensions. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, Vol. 2. AAAI Press, Austin, Texas, 1021–1028. http://dl.acm.org/citation.cfm?id=2887007.2887149

[30] Hadas Shachnai and Tami Tamir. 2001. On Two Class-Constrained Versions of the Multiple Knapsack Problem. *Algorithmica* 29, 3 (2001), 442–467. https://doi.org/10.1007/s004530010057

[31] Koichi Tadenuma and William Thomson. 1993. The fair allocation of an indivisible good when monetary compensations are possible. *Mathematical Social Sciences* 25, 2 (1993), 117–132. https://doi.org/10.1016/0165-4896(93)90047-M

[32] A.G.R. Vaz, B. Elsinga, W.G.J.H.M. van Sark, and M.C. Brito. 2016. An Artificial Neural Network to Assess the Impact of Neighbouring Photovoltaic Systems in Power Forecasting in Utrecht, the Netherlands. *Renewable Energy* 85 (2016), 631–641. https://doi.org/10.1016/J.RENENE.2015.06.061

[33] Sichao Yang and Bruce Hajek. 2007. VCG-Kelly Mechanisms for Allocation of Divisible Goods: Adapting VCG Mechanisms to One-Dimensional Signals. *IEEE Journal on Selected Areas in Communications* 25, 6 (2007), 1237–1243. https://doi.org/10.1109/JSAC.2007.070817