On the Complexity of Pareto-Optimal and Envy-Free Lotteries

Ioannis Caragiannis
Aarhus University
Aarhus, Denmark
iannis@cs.au.dk

Kristoffer Arnsfelt Hansen
Aarhus University
Aarhus, Denmark
arnsfelt@cs.au.dk

Nidhi Rathi
Max Planck Institute for Informatics,
University of Saarland
Saarbrücken, Germany
nrathi@mpi-inf.mpg.de

ABSTRACT

We study the classic problem of dividing a collection of indivisible resources in a *fair* and *efficient* manner among a set of agents having varied preferences. *Pareto optimality* is a standard notion of economic efficiency, which states that it should be impossible to find an allocation that improves some agent's utility without reducing any other's. On the other hand, a fundamental notion of fairness in resource allocation settings is that of *envy-freeness*, which renders an allocation to be fair if every agent (weakly) prefers her own bundle over that of any other agent's bundle. Unfortunately, an envy-free allocation may not exist if we wish to divide a collection of indivisible items. Introducing randomness is a typical way of circumventing the non-existence of solutions, and therefore, *allocation lotteries*, i.e., distributions over allocations have been explored while relaxing the notion of fairness to *ex-ante* envy freeness.

We consider a general fair division setting with n agents and a family of admissible n-partitions of an underlying set of items. Every agent is endowed with *partition-based utilities*, which specify her cardinal utility for each bundle of items in every admissible partition. In such fair division instances, Cole and Tao (2021) have proved that an ex-ante envy-free and Pareto-optimal allocation lottery is always guaranteed to exist. We strengthen their result while examining the computational complexity of the above total problem and establish its membership in the complexity class PPAD. Furthermore, for instances with a constant number of agents, we develop a polynomial-time algorithm to find an ex-ante envy-free and Pareto-optimal allocation lottery. On the negative side, we prove that maximizing social welfare over ex-ante envy-free and Pareto-optimal allocation lotteries is NP-hard.

KEYWORDS

Fair Division, Efficiency, Randomization, Complexity, PPAD

ACM Reference Format:
Ioannis Caragiannis, Kristoffer Arnsfelt Hansen, and Nidhi Rathi. 2024. On the Complexity of Pareto-Optimal and Envy-Free Lotteries . In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

Fairly dividing a collection of resources among individuals (often dubbed as agents) with varied preferences forms a key concern in the design of many social institutions. Such problems arise naturally in many real-world scenarios such as assigning computational resources in a cloud computing environment, air traffic management, dividing business assets, allocation of radio and television spectrum, course assignments, and so on [2, 26, 36, 44]. The fundamental problem of fair division lies at the interface of economics, social science, mathematics, and computer science, and its formal study dates back about seven decades [24, 39]. In the last few decades, the area of fair division has witnessed a flourishing flow of research; see [3, 10, 11] for excellent expositions.

Traditionally, in early literature, fair division has been studied for a single *divisible* resource, classically known as *fair cake cutting*. Here, each agent specifies her valuations over a unit interval cake via a probability distribution over $[0, 1]$ and the problem is to divide the cake among agents in a fair manner. The quintessential notion of fairness in this line of work is that of *envy-freeness*, introduced by Foley [29] and Varian [42]. A cake division is said to be *envy-free* if every agent prefers her own share of the cake over any other agent's share. Stromquist [40] famously proved that an envy-free cake division (where every agent receives a connected interval of the cake) is always guaranteed to exist, under mild conditions. Later, Su [41] developed another existential proof using Sperner's Lemma and established a connection between the notion of envy-freeness and topology. Such strong existential results have arguably placed the notion of envy-freeness as the flagship bearer of fairness in resource allocation settings.

On the other hand, *Pareto optimality* is a standard notion of economic efficiency, which states that it should be impossible to find an allocation that improves some agent's utility without reducing any other's. Another important notion of (collective) efficiency measure of an allocation is that of *social welfare* [36] which is the sum of all the utilities derived by agents from their assigned bundle.

The goal of being fair towards the participating agents and achieving collective (economic) efficiency form the two important paradigms of resource allocation problems. Unfortunately, for an indivisible set of items, an envy-free allocation may not exist. For example, an instance with two agents having positive value for a single item admits no envy-free allocation.

A fair division instance consists of a set $[n] = \{1, 2 \ldots, n\}$ of n agents and a set M of items. In the most basic setting, every agent i has an additive utility function $u_i : 2^M \to \mathbb{R}$ that specifies her cardinal preferences for the items of a given bundle; in particular, $u_i(j) := u_i(\{j\})$ denotes agent i's utility for item $j \in M$. We say an *allocation* is a partition of items into n bundles, where every agent is assigned one bundle. The goal of simultaneously achieving fairness and efficiency is challenging for the problem of allocating indivisible items. Besides the mentioned fact that an envy-free allocation is not

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

guaranteed to exist, in cases where envy-free allocations do exist, envy-freeness may not be compatible with Pareto optimality [9].

The above discussion suggests that one should consider distributions over allocations (to be referred as *allocation lotteries*) in order to simultaneously achieve fairness and efficiency guarantees. In the random assignment literature in economics, the idea of constructing a fractional allocation and implementing it as a lottery over deterministic allocations was introduced by Hylland and Zeckhauser [31]. Introducing randomness is a typical way of circumventing the non-existence of various solution concepts, especially in computational social choice theory [1, 6, 14, 23]. In the process of exploring allocation lotteries, we appropriately relax the notion of fairness to *ex-ante envy-freeness*, which values the random bundles allocated agents in terms of expected utility. Recent works of [7, 13, 16, 30] deal with various computational aspects of allocation lotteries that are fair and efficient for fair division instances with additive utilities. Observe that, allocation lotteries that are just ex-ante envy-free or just ex-ante Pareto-optimal can be trivially computed in polynomial time. For the former, one can solve a linear program, while for the latter, one can assign each bundle to the agent that has the highest utility for it. That is, these notions of fairness and efficiency are tractable if dealt with individually. Therefore, the important question is to understand the computational complexity of computing allocation lotteries that are simultaneously ex-ante envy-free and Pareto-optimal. In this work, we consider this question for the most general setting of fair division, as detailed in the following section.

Context and Overview of Results. In this work, we consider a very general fair division setting with n agents and a family of admissible n-partitions of an underlying set of items. Every agent is endowed with *partition-based utilities* that specify her cardinal utility for different bundles in every partition. For such a broad class of fair division instances with partition-based utilities, including negative-valued utilities, the recent work of Cole and Tao [22] proves that an ex-ante envy-free and Pareto-optimal allocation lottery is always guaranteed to exist.

Note that, partition-based utilities provide a much broader way of expressing agents' utilities. In particular, it is possible that an agent may value the exact same bundle of items in two distinct partitions at two different values or, there may be a certain partition of items that is not favourable or suitable (depending on the context of application). This generalization allows us to remove unsuitable partitions from the family of admissible partitions, and still, the existence of ex-ante envy-free and Pareto optimal allocation lotteries is guaranteed.

In this work, we examine the computational complexity of the above total search problem and strengthen the work of Cole and Tao [22]. In particular, we establish that the problem of finding an ex-ante envy-free and Pareto optimal allocation lottery for fair division instances with partition-based utilities belongs to the complexity class PPAD. This containment result is even interesting for the special case of a single admissible partition. Namely, our PPAD membership result is for the *exact* search problem, of computing a rational valued lottery. This can be contrasted with the lottery provided by the Hylland-Zeckhauser (HZ) pseudo-market. Vazirani and Yannakakis [43] gave a simple example with four agents and four goods where the *unique* HZ equilibrium gives an irrational-valued lottery. This fact means that any algorithm for computing a HZ equilibrium exactly must overcome numerical challenges. Our result on the other hand gives hope for the possibility of developing a *practical* algorithm for computing an exact ex-ante envy-free and Pareto optimal allocation lottery, for instance by an adaptation of Lemke's algorithm [34].

For instances with a constant number of agents, we develop a polynomial-time algorithm to compute an exact ex-ante envy-free and Pareto-optimal lottery. On the negative side, we prove that maximizing social welfare over ex-ante envy-free and Pareto optimal allocation lotteries is NP-hard.

Further Related Work. Fairness in resource-allocation settings is extensively studied in the economics, mathematics, and computer science literature (see [10, 11, 36]). As mentioned above, envy-free allocations may not exist for the case of indivisible items. Since envy-freeness is arguably a fundamental notion of fairness, as evident from its importance in fair cake cutting, there has been a significant body of research aimed towards finding ex-ante envy-free allocation lotteries in the indivisible setting. The work of Freeman et al. [30] addresses the key question of whether ex-ante envy-freeness can be achieved in combination with *ex-post envy-freeness up to one item.* They settle it positively by designing an efficient algorithm that achieves both properties simultaneously. Caragiannis et al. [16] explore the *interim allocation lotteries (iEF)* which provide fairness guarantees that lie between ex-post and ex-ante envy-freeness. They develop polynomial-time algorithms for computing iEF lotteries that maximize various efficiency notions.

Budish et al. [13] employ a general class of random allocation mechanisms to achieve ex-ante fairness and efficiency in the presence of real-world constraints. Several other works explore fairness and efficiency guarantees of allocation lotteries as well, but for ordinal utilities [1, 8, 21].

Another line of research has explored various relaxations of envy-freeness. The notion of *envy-freeness up to one item (EF1)* was introduced by Budish et al. [12] as one of the first 'good' relaxations of envy-freeness in the indivisible setting. We say an allocation is EF1 when every agent (weakly) prefers her own bundle over any other agent j's bundle after removing some item from j's bundle. EF1 allocations are guaranteed to always exist for general monotone valuations and can be computed efficiently [35]. Moreover, this fairness notion is compatible with the economic efficiency objective of Pareto-optimality [17]. Later, *envy-freeness up to any item (EFX)* was introduced by Caragiannis et al. [17] as a refinement of EF1 and is now considered as the most compelling fairness criterion while dividing indivisible items. We say an allocation is EFX when every agent (weakly) prefers her own bundle than any other agent j's bundle after removing her least positively-valued item from j's bundle. Recent works [4, 5, 18, 19] have shown existential guarantees for EFX in various special cases.

2 THE MODEL

Consider a set $[n] = \{1, 2, \ldots, n\}$ of n agents and a collection $\mathcal{P} = \{P^1, P^2, \ldots, P^m\}$ of admissible partitions of a set M of items. Every partition P^k for $k \in [m]$ consists of n bundles, i.e., $|P^k| = n$

and the union of those bundles is $\bigcup_{A \in P^k} A \subseteq M$. Agents are endowed with utility functions u_i's that specify their cardinal preferences for all bundles in every different partition. In particular, the function u_{ij}^k specifies *partition-based cardinal utilities* of agent $i \in [n]$ for the jth bundle (for $j \in [n]$) in partition $P^k \in \mathcal{P}$. It is important to note that an agent with partition-based utilities can have different utilities for the exact same bundle, occurring in two distinct partitions. We will denote a fair division instance by the tuple $\mathcal{I} = \langle [n], \mathcal{P}, \{u_{ij}^k\}_{i,j \in [n], k \in [m]} \rangle$.

For a given fair division instance, we define an *allocation* to be an assignment of the n bundles of a partition in $\mathcal{P}$ to the agents, such that every agent receives exactly one bundle. We assume that for any partition, the set of (admissible) allocations is specified by the $n!$ permutations of n bundles among n agents, and that the utility of an agent depends only on the partition and the bundle received and not to whom the remaining bundles are given. We refer to this property of a fair division instance as the *anonymity property*. Therefore, we have a total of $m \cdot n!$ many distinct admissible *allocations* in a given fair division instance. Furthermore, in a given fair division instance, we define a *lottery* to be a probability distribution over these allocations.

The overarching goal is to find a *fair* and *efficient* lottery among agents from the given set of admissible partitions. As mentioned, Cole and Tao [22] established the existence of fair and efficient lotteries for fair division instances with the anonymity property using Kakutani's fixed-point theorem [32]. Since there are a total of $m \cdot n!$ many allocations, one can specify probabilities with which every allocation occurs in a lottery. This leads to a very convenient but also very inefficient way of representing a lottery via an exponential-dimensional vector $(p_1, p_2, \dots, p_{m \cdot n!})$, where p_i represents the probability with which the i-th allocation is chosen. This representation was used by Cole and Tao [22] for their proof of existence, but it is clearly not suitable for studying the computational aspects of finding lotteries. Instead, we will represent a lottery in the following manner: Let $\mathbf{p} = \{p_k\}_{k \in [m]}$, where $p_k \in [0, 1]$ denotes the probability with which partition $P^k \in \mathcal{P}$ is selected in a lottery. The vector $\mathbf{q} = \{q_{ij}^k\}_{i,j \in [n], k \in [m]}$ of length $m \cdot n^2$ then specifies the full lottery, where q_{ij}^k is the probability with which the lottery $\mathbf{q}$ assigns the jth bundle in partition P^k to agent $i \in [n]$. The vectors $\mathbf{p}$ and $\mathbf{q}$ are characterized by the following constraints.

$$\sum_{i=1}^{n} q_{ij}^k = p_k \text{ for all } j \in [n], k \in [m]$$

$$\sum_{j=1}^{n} q_{ij}^k = p_k \text{ for all } i \in [n], k \in [m]$$

$$\text{and, } \sum_{k=1}^{m} p_k = 1$$

The constraints express that every bundle is assigned to one agent and that every agent receives one bundle, respectively. We can now express the expected utility, $\mathbb{E}[u_i(\mathbf{q})]$, for agent $i \in [n]$ in a lottery $\mathbf{q}$ as

$$\mathbb{E}[u_i(\mathbf{q})] \coloneqq \sum_{k=1}^{m} \sum_{j=1}^{n} u_{ij}^k \cdot q_{ij}^k \ .$$

More generally, let

$$u_i(\mathbf{q}; i') \coloneqq \sum_{k=1}^{m} \sum_{j=1}^{n} u_{ij}^k \cdot q_{i'j}^k$$

denote the expected utility of agent i for the bundle of agent i' in the lottery $\mathbf{q}$. Observe that, we have $u_i(\mathbf{q}; i) = \mathbb{E}[u_i(\mathbf{q})]$ for any agent $i \in [n]$.

Let us now define the standard notions of fairness and optimality in resource allocation settings. A lottery $\mathbf{q}$ is said to be *ex-ante envy-free* if $u_i(\mathbf{q}; i) \geq u_i(\mathbf{q}; i')$ holds for all $i, i' \in [n]$. Furthermore, we say that $\mathbf{q}$ is *ex-ante Pareto-optimal* if there does not exist any other lottery $\widetilde{\mathbf{q}}$ such that $u_i(\widetilde{\mathbf{q}}; i) \geq u_i(\mathbf{q}; i)$ holds for all $i \in [n]$, with a strict inequality for at least one agent $i \in [n]$. *Social welfare* is a standard notion of measuring the collective welfare of an allocation. We define *social welfare* of a lottery $\mathbf{q}$ as the sum of the expected utilities of all agents, i.e., $\mathrm{SW}(\mathbf{q}) = \sum_{i \in [n]} u_i(\mathbf{q}; i)$.

3 PPAD-MEMBERSHIP

In this section, we show that the problem of finding an exact ex-ante envy-free and Pareto-optimal lottery in a given fair division instance belongs to the class PPAD. Our proof is based on (i) a significant simplification of the existence proof of Cole and Tao [22], (ii) a characterization of PPAD in terms of computing fixed points of piecewise linear arithmetic circuits due to Etessami and Yannakakis [25] (i.e. PPAD = Linear-FIXP), and (iii) a framework for proving FIXP and PPAD-membership via convex optimization recently developed by Filos-Ratsikas et al. [27, 28]. Formally, we obtain the following theorem.

Theorem 3.1. *The problem of finding an ex-ante envy-free and Pareto-optimal lottery in a fair division instance belongs to* PPAD.

It is possible to adapt the existence proof of Cole and Tao (by changing to our succinct representation of lotteries) to obtain a proof of FIXP membership using the framework of Filos-Ratsikas et al [27]. The proof of Cole and Tao employs Kakutani's fixed point theorem to a correspondence defined on pairs consisting a lottery $\mathbf{q}$ and a vector of positive weights $\mathbf{w} \in W_\varepsilon$ for the agents, from a closed set W_ε. This correspondence maps $(\mathbf{q}, \mathbf{w})$ to pairs $(\mathbf{q}', \mathbf{w}')$ such that $\mathbf{q}'$ is a lottery maximizing the weighted sum of utilities of the agents and where $\mathbf{w}'$ is obtained from $\mathbf{w}$ by translating each coordinate by a nonlinear function of the lottery $\mathbf{q}$ followed by a projection to the set W_ε.

The maximization of the weighted sum of utilities may be phrased as a linear program and the projection may be phrased as a convex quadratic program. While both of these fall in the scope of the framework of Filos-Ratsikas et al. [28] for proving PPAD-membership, the nonlinear transformation involved cannot be computed be a piecewise linear arithmetic circuit.

Our simplified proof involves only optimization of a linear program and the solution of a feasibility program with conditional linear constraints, together with operations computable by linear arithmetic circuits. In this case the framework Filos-Ratsikas et al. applies to give PPAD-membership [28].

Another framework for proving PPAD-membership was also recently introduced by Papadimitriou, Vlatakis-Gkaragkounis and Zampetakis [38]. With this framework, however, it would only be possible to directly prove PPAD-membership for an *approximate* version of the problem, rather than the exact problem.

In the remainder of this section, we let $\mathcal{I} = \langle [n], \mathcal{P}, \{u_{ij}^k\}_{i,j,k} \rangle$ denote a fair division instance with n agents and $m = |\mathcal{P}|$ partitions, where the utilities u_{ij}^k are given as rational numbers.

3.1 Fixed Point Formulation

We first present our fixed point formulation for ex-ante envy-free and Pareto-optimal lotteries; afterwards we consider the implications for the computational complexity of the problem.

A standard technique for expressing the Pareto frontier of an optimization problem, also employed by Cole and Tao, is the *weighted sum method* [45]. Let $w_1, \ldots, w_n > 0$ be strictly positive weights. Then, any lottery $\mathbf{q}$ maximizing the weighted sum of utilities $\sum_{i=1}^n w_i u_i(\mathbf{q}; i)$ must be Pareto-optimal. Conversely, if $\mathbf{q}$ is a Pareto-optimal lottery, there are strictly positive weights such that $\mathbf{q}$ maximizes the weighted sum of utilities.

The task of maximizing the weighted sum of utilities can be expressed by the following linear program with decision variables q_{ij}^k and p_k, and parameterized by the variables w_i.

$$\begin{array}{lll} \max & \sum_{i=1}^n w_i \sum_{k=1}^m \sum_{j=1}^n u_{ij}^k q_{ij}^k & \\ \text{s.t.} & \sum_{i=1}^n q_{ij}^k = p_k & \text{for all } j \in [n], k \in [m] \\ & \sum_{j=1}^n q_{ij}^k = p_k & \text{for all } i \in [n], k \in [m] \\ & \sum_{k=1}^m p_k = 1 & \\ & p_k \geq 0 & \text{for all } k \in [m] \\ & q_{ij}^k \geq 0 & \text{for all } i, j \in [n], k \in [m] \end{array} \tag{1}$$

The proof of Cole and Tao [22] shows the existence of positive weights such that any lottery $\mathbf{q}$ maximizing the corresponding weighted sum of utilities is also ex-ante envy-free.

We next define the following key quantity $0 < \rho \leq \frac{1}{2}$ and state Lemma 3.3 (proved in [22]) that will be used to place restrictions of weights.

Definition 3.2. Let $J = \{(k, l, h, a, b) \in [m] \times [n]^4 \mid u_{la}^k < u_{lb}^k \text{ and } u_{ha}^k < u_{hb}^k\}$. We define ρ as follows,

$$\rho = \begin{cases} \frac{1}{2} \min_{(k,l,h,a,b) \in J} (u_{lb}^k - u_{la}^k)/(u_{hb}^k - u_{ha}^k) & \text{if } J \neq \emptyset \\ \frac{1}{2} & \text{otherwise} \end{cases}$$

Lemma 3.3 (cf. [22, Claim 4.13]). *Suppose that $(\mathbf{q}, \mathbf{p})$ is an optimal solution of LP (1). If $0 < w_h \leq \rho w_l$ it follows that $u_l(\mathbf{q}; l) \geq u_l(\mathbf{q}; h)$ (i.e. that agent l does not envy agent h).*

Define $\varepsilon = \rho^n/n$ and let $W_\varepsilon = \{w \in \mathbb{R}^n \colon \sum_{i=1}^n w_i = 1 \text{ and } w_i \geq \varepsilon \; \forall\, i \in [n]\}$. We shall restrict the weights to belong to W_ε, which in particular, ensure that they are strictly positive. We consider the following feasibility problem with conditional linear constraints having decision variables w_i, and parameterized by variables q_{ij}^k.

$$\begin{array}{ll} [u_l(\mathbf{q}; h) - u_l(\mathbf{q}; l) > 0] \Rightarrow [w_h - \rho w_l \leq 0] & l, h \in [n] \\ \sum_{i=1}^n w_i = 1 & \\ w_i \geq \varepsilon & i \in [n] \end{array} \tag{2}$$

Here, the conditional constraint $[u_l(\mathbf{q}; h) - u_l(\mathbf{q}; l) > 0] \Rightarrow [w_h - \rho w_l \leq 0]$ is satisfied if either $u_l(\mathbf{q}; h) - u_l(\mathbf{q}; l) \leq 0$ or $w_h - \rho w_l \leq 0$. In words, whenever agent l envies agent h in the lottery q, a solution $\mathbf{w}$ of the system must satisfy $0 < w_h \leq \rho w_l$, which is precisely the antecedent stated in Lemma 3.3. We can think of the feasibility problem as a system of inequalities in variables $\mathbf{w}$, some of which may be "disabled" by inequalities expressed in the variables $\mathbf{q}$.

In order to characterize the solvability of this feasibility program, it is convenient to introduce the *envy graph* of the lottery $\mathbf{q}$. Filos-Ratsikas et al. [28] consider feasibility programs as above in a general form and characterizes their solvability in terms of a *feasibility graph.* In our case, this feasibility graph is exactly the same as the *envy graph* defined next.

Definition 3.4 (Envy graph). For a given lottery $\mathbf{q}$, denote by $\mathcal{G}(\mathbf{q})$ the *envy graph* with nodes $[n]$ and an arc (l, h) whenever $u_l(\mathbf{q}; l) < u_l(\mathbf{q}; h)$, for all $l, h \in [n]$. We let $A(\mathcal{G}(\mathbf{q}))$ denote the set of arcs of $\mathcal{G}(\mathbf{q})$.

We can then precisely characterize the solvability of the feasibility problem (2) by the graph structure of $\mathcal{G}(\mathbf{q})$.

Lemma 3.5. *Suppose that $\mathbf{q}$ is a lottery such that $\mathcal{G}(\mathbf{q})$ is acyclic. Then the feasibility program (2) is solvable.*

Proof. First note that the condition $u_l(\mathbf{q}; h) - u_l(\mathbf{q}; l) > 0$ is satisfied precisely when $(l, h) \in A(\mathcal{G}(\mathbf{q}))$. Thus we are to find weights w_i such that $w_h \leq \rho w_l$, whenever $(l, h) \in A(\mathcal{G}(\mathbf{q}))$.

For $i \in [n]$, let d_i denote the length of a longest path in $\mathcal{G}(\mathbf{q})$ from node i to a sink node, and define

$$w_i = \frac{\rho^{d_i}}{\sum_{j=1}^n \rho^{d_j}} \text{ for all } i \in [n] \ .$$

Clearly $\sum_{i=1}^n w_i = 1$, and since $d_i \leq n$ and $\rho \leq 1$ we also have $w_i \geq \varepsilon$. Suppose now that $(l, h) \in A(\mathcal{G}(\mathbf{q}))$. This means that $d_l \geq d_h + 1$ and thus also $w_l \leq \rho w_h$. In conclusion, we have that the weights w_i are a solution to the feasibility program (2). □

We can note that acyclicity of $\mathcal{G}(\mathbf{q})$ is also necessary for the solvability of the feasibility program (2), since the inequalities $w_l \leq \rho w_h$ given by the arcs (l, h) of a cycle in $\mathcal{G}(\mathbf{q})$ are contradictory. But note also that if $\mathcal{G}(\mathbf{q})$ contains a cycle, all agents in the cycle will increase their utility if the lottery is shifted along the cycle. We thus have the following simple but crucial observation.

Observation 3.6 (cf. [22, Claim 4.8]). *If $\mathbf{q}$ is Pareto-optimal, the envy graph $\mathcal{G}(\mathbf{q})$ is acyclic.*

We can now conclude with the following fixed-point formulation, showing that a pair $(\mathbf{q}, \mathbf{w})$ that is simultaneously solving the linear program (1) and the feasibility problem (2) give an ex-ante envy-free and Pareto-optimal lottery.

Proposition 3.7. *Suppose $\mathbf{q}$ is a lottery and $\mathbf{w} \in W_\varepsilon$ are weights such that $\mathbf{q}$ is an optimal solution of the linear program LP (1) w.r.t. the weights $\mathbf{w}$, and $\mathbf{w}$ is a solution of the feasibility program of conditional linear constraints (2) with conditions given by $\mathbf{q}$ (note that the system is in fact solvable by the optimality of $\mathbf{q}$, Observation 3.6 and Lemma 3.5). Then $\mathbf{q}$ is an ex-ante envy-free and Pareto-optimal lottery.*

Proof. Since the weights $\mathbf{w}$ are strictly positive and $\mathbf{q}$ is an optimal solution of LP (1) it follows that $\mathbf{q}$ is Pareto-optimal. Suppose now for contradiction that there exists agents l and h such that agent l envies agent h, that is, $u_l(\mathbf{q}; h) > u_l(\mathbf{q}; l)$. Since $\mathbf{w}$ is a solution to the system (2) with conditions given by $\mathbf{q}$ given it follows that $w_h \leq \rho w_l$. But then Lemma 3.3 gives $u_l(\mathbf{q}; h) \leq u_l(\mathbf{q}; l)$, contradicting the assumption. It follows that $\mathbf{q}$ must also be ex-ante envy-free. ◻

3.2 PPAD, FIXP, and Linear-FIXP

The complexity class PPAD was originally defined in seminal work of Papadimitriou [37] as the class of total NP search problems reducible to a concrete problem called End-Of-Line. As mentioned above, to obtain result, we shall instead make use of a characterization of PPAD in terms of computation of fixed points of functions computed by piecewise linear arithmetic circuits. Below we briefly introduce this characterization and refer to [25] for further details.

An arithmetic circuit is a circuit C with gates computing binary operations belonging to the set $\{+, -, *, \div, \max, \min\}$ together with rational constants. The size of C refers to the size of an encoding of C. A *piecewise linear* arithmetic circuit C restricts the allowable binary operations to the set $\{+, -, \max, \min\}$, but allows also for multiplication by rational constants.

The class FIXP consists of (real-valued) search problems that reduce to finding a fixed point of a function $F \colon D \to D$, where D is an explicitly given convex polytope and F is a function computable by an algebraic circuit. By Brouwer's fixed point theorem such a fixed point is guaranteed to exist, thus making the search problem a total search problem. Linear-FIXP is the subclass obtained by restricting the arithmetic circuits to be piecewise linear.

As defined above, the classes FIXP and Linear-FIXP consist of real-valued search problems, which means that reductions must specify a real-valued function mapping fixed points of the function F to solutions of the search problem. In the case when F is computed by a piecewise linear arithmetic circuit C, there exists rational-valued fixed points of polynomial bitsize in the size of C [25, Theorem 5.2], which allows the use of ordinary polynomial-time reductions. With this convention, Etessami and Yannakakis [25] showed that PPAD = Linear-FIXP [25, Theorem 5.4].

3.3 PPAD-Membership via Convex Optimization

From the characterization PPAD = Linear-FIXP, in order to prove Theorem 3.1, it is sufficient to reduce the task of computing an ex-ante envy-free and Pareto-optimal lottery to that of computing a fixed point of a piecewise linear arithmetic circuit defined on an explicitly given convex polytope.

Constructing such a suitable circuit from scratch can potentially be a very challenging task, as many existing proofs of PPAD-membership in the literature give evidence of. Recently however, Filos-Ratsikas et al. [27, 28] introduced a general technique for proving FIXP and PPAD-membership, by which the arithmetic circuit defining the fixed point search problem can be augmented with *pseudo-gates* that solve very general convex optimization problems. By a pseudo-gate, we mean a (multi-input and multi-output) gate that is only required to compute the correct output at a fixed point of the full circuit. More precisely, the pseudo-gate is implemented by an arithmetic circuit using auxiliary variables, and when these auxiliary variables are in a fixed point, the pseudo-gate computes the correct output.

Definition 3.8 (Pseudo-circuit)**.** A pseudo-circuit with n inputs and m outputs is an arithmetic circuit C computing a function $F \colon \mathbb{R}^n \times [0,1]^\ell \to \mathbb{R}^m \times [0,1]^\ell$. The output of C on input $x \in \mathbb{R}^n$ is any $y \in \mathbb{R}^m$ such that there exists $z \in [0,1]^\ell$ such that $F(x,z) = (y,z)$. The variables $z \in [0,1]^\ell$ are called auxiliary variables.

By a pseudo-gate, we mean the use of a pseudo-circuit as a sub-circuit of larger pseudo-circuit, and where the auxiliary variables of the pseudo-gate are augmented to the auxiliary variables of the larger pseudo-circuit. The simple but crucial observation about pseudo-circuits is that, for the purpose of proving FIXP and PPAD-membership they are just as good as normal arithmetic circuits.

In the setting of PPAD, Filos-Ratsikas et al. [28] developed a pseudo-gate, coined the *linear-OPT-gate*, implemented as a piecewise linear arithmetic circuit, that in particular can be used to solve both the linear program (1) and the feasibility problem (2). For the linear program (1) this is possible since the coefficients of all linear constraints are constants and that the coefficients of the objective function are linear functions of the parameter variables $\mathbf{w}$. For the feasibility program (2) this is possible since the coefficients of all linear constraints are constants and the antecedents of the conditional linear constraints are given by a strict linear inequalities for functions computable by piecewise linear circuits applied to the parameter variables $\mathbf{q}$.

3.4 Proof of Theorem 3.1

We finally show how our fixed point formulation for ex-ante envy-free and Pareto-optimal lotteries in conjunction with the framework of Filos-Ratsikas et al. [28] allows for a simple proof of PPAD membership for the problem of computing such lotteries.

The fixed point formulation of Proposition 3.7 amounts to finding $(\mathbf{q}, \mathbf{p}, \mathbf{w})$ such that $(\mathbf{q}, \mathbf{p})$ is an optimal solution of the LP (1), parametrized by $\mathbf{w}$, and such that $\mathbf{w}$ is a solution to the feasibility program of conditional linear constraints (2), parametrized by $\mathbf{q}$.

We thus build a piecewise linear arithmetic pseudo-circuit C accomplishing both tasks. The circuit C takes as input the variables $(\mathbf{q}, \mathbf{p}, \mathbf{w})$. Using the linear-OPT-gate of [28] we let C output $(\mathbf{q}', \mathbf{p}', \mathbf{w}')$ such that:

(1) $(\mathbf{q}', \mathbf{p}')$ is an optimal solution of LP (1) parametrized by $\mathbf{w}$.
(2) If the feasibility program (2) parametrized by $\mathbf{q}$ is feasible, then $\mathbf{w}'$ is a solution.

Suppose now that $(\mathbf{q}, \mathbf{p}, \mathbf{w})$ is a fixed point of the circuit C (where also the auxiliary inputs of C are assumed to be in a fixed point). Since $(\mathbf{q}, \mathbf{p})$ is then an optimal solution of LP (1) parametrized by $\mathbf{w}$, this means that $\mathbf{q}$ is Pareto-optimal by the weighted sum method. From Observation 3.6 and Lemma 3.5 we then have that the feasibility program (2) parametrized by $\mathbf{q}$ is in fact feasible, and this then means that $\mathbf{w}$ is a solution. By Proposition 3.7 we can then conclude that $\mathbf{q}$ is an ex-ante envy-free and Pareto-optimal lottery.

We have thus reduced the task of computing an ex-ante envy-free and Pareto-optimal lottery to the task of computing a fixed point of a piecewise linear arithmetic pseudo-circuit defined on a explicitly given convex polytope, thereby completing the proof.

4 AN EFFICIENT ALGORITHM FOR CONSTANT NUMBER OF AGENTS

In this section, we develop a very simple polynomial time algorithm for computing an ex-ante envy-free and Pareto-optimal lottery when the number of agents is constant. Consider a fair division instance $\mathcal{I}$ consisting of n agents, a set of m partitions $\mathcal{P} = \{P^1, P^2, ..., P^m\}$, and agent utilities u_{ij}^k for $i, j \in [n]$ and $k \in [m]$. The algorithm begins with evaluating the agents' valuations in the $n!$ possible allocations for each partition P^k for $k \in [m]$. That is, we obtain $n!$ utility profiles in $\mathbb{R}^n$ for each partition and $m \cdot n!$ utility profiles overall. The Pareto-optimal lotteries are formed by faces of the convex hull of these utility profiles.

Since the dimension n is constant, the convex hull can be computed in polynomial time [20]. We may then enumerate over the faces forming the Pareto-frontier. For each of these faces, we compute a hyperplane H that contains the face. For such a hyperplane $H = \{x \in \mathbb{R}^n \mid w_1x_1 + w_2x_2 + \cdots + w_nx_n = w_0\}$, we can determine whether it contains an envy-free lottery by linear programming.

$$
\begin{array}{lll}
\text{find} & (\mathbf{q}, \mathbf{p}) & \\
\text{s.t.} & u_i(\mathbf{q}; i) \geq u_i(\mathbf{q}; i') & \text{for all } i, i' \in [n] \\
& \sum_{i=1}^{n} w_i u_i(\mathbf{q}; i) = w_0 & \\
& \sum_{i=1}^{n} q_{ij}^k = p_k & \text{for all } j \in [n], k \in [m] \\
& \sum_{j=1}^{n} q_{ij}^k = p_k & \text{for all } i \in [n], k \in [m] \\
& \sum_{k=1}^{m} p_k = 1 & \\
& p_k \geq 0 & \text{for all } k \in [m] \\
& q_{ij}^k \geq 0 & \text{for all } i, j \in [n], k \in [m]
\end{array}
$$

Since we know that there does exist an ex-ante envy-free and Pareto-optimal lottery, at least one of these linear programs must be feasible. The next statement summarizes the discussion above.

Theorem 4.1. *For fair division instances with a constant number of agents, an ex-ante envy-free and Pareto-optimal allocation can be computed in polynomial time.*

5 ENVY-FREE AND PARETO-OPTIMAL LOTTERIES OF HIGH SOCIAL WELFARE

As our last technical contribution, we study the problem of optimizing social welfare over ex-ante envy-free and Pareto-optimal allocation lotteries and prove the following statement for its decision version. A few proofs are omitted; these may be found in the full version of the paper [15].

Theorem 5.1. *Given a fair division instance with partition-based utilities and $K > 0$, the problem of deciding whether there exists an ex-ante envy-free and Pareto-optimal allocation lottery of social welfare at least K is* NP*-complete.*

It is easy to see that the above problem belongs to the complexity class NP. First, notice that it is trivial to check whether a given lottery $\mathbf{q}$ is ex-ante envy-free and has social welfare at least K. To verify Pareto optimality, it suffices to search for another lottery $\widetilde{\mathbf{q}}$ which gives to any agent expected utility at least as high as her expected utility in $\mathbf{q}$, maximizing the total excessive utility through the following linear program:

$$
\begin{array}{lll}
\max & \sum_{i=1}^{n} t_i & \\
\text{s.t.} & u_i(\widetilde{\mathbf{q}}; i) \geq u_i(\mathbf{q}; i) + t_i & \text{for all } i \in [n] \\
& \sum_{i=1}^{n} \widetilde{q}_{ij}^k = \widetilde{p}_k & \text{for all } j \in [n], k \in [m] \\
& \sum_{j=1}^{n} \widetilde{q}_{ij}^k = \widetilde{p}_k & \text{for all } i \in [n], k \in [m] \\
& \sum_{k=1}^{m} \widetilde{p}_k = 1 & \\
& \widetilde{q}_{ij}^k \geq 0 & \text{for all } i, j \in [n], k \in [m] \\
& \widetilde{p}_k \geq 0 & \text{for all } k \in [m] \\
& t_i \geq 0 & \text{for all } i \in [n]
\end{array}
$$

Clearly, the lottery $\widetilde{\mathbf{q}}$ Pareto-dominates $\mathbf{q}$ if and only if the objective value of the above linear program is strictly positive.

For proving NP-hardness, we will develop a polynomial-time reduction from the classic NP-complete problem *Exact Cover by 3-Sets* (X3C) [33] to our problem. X3C is defined as follows:

Instance: A universe $\mathcal{E} = \{e_1, e_2, \ldots, e_r\}$ of r elements, a family $\mathcal{S} = \{S_1, S_2, \ldots, S_t\}$ of triplets from $\mathcal{E}$, i.e., $S_j \subseteq \mathcal{E}$ with $|S_j| = 3$ for all $j \in [t]$.

Question: Does there exist an exact cover, i.e., a set of $r/3$ triplets from $\mathcal{S}$ that includes all elements of the universe $\mathcal{E}$?

For $i \in [r]$, we let f_i denote the frequency of occurrence of element e_i, i.e., $f_i := |\{j : e_i \in S_j\}|$.

5.1 The Reduction

Starting with an instance ϕ of X3C, our reduction constructs a fair division instance $\mathcal{I}(\phi)$ as follows. Instance $\mathcal{I}(\phi)$ has the following set of $n = t + 1 + 2t^2 + 3r$ agents.

- $t + 1$ *base agents* $b_0, b_1, b_2, \ldots, b_t$,
- $2t$ *set agents* $h_{j,1}, h_{j,2}, \ldots, h_{j,2t}$, for every $j \in [t]$,
- three *element agents* v_i, w_i, and z_i for every $i \in [r]$

The set $\mathcal{P}$ of admissible partitions of an underlying set of items consists of $m = 3t$ partitions $P_{j,c}$ for $j \in [t]$ and $c \in [3]$. We identify the n bundles of partitions in accordance to the type of agents. So, each partition has $t+1$ bundles $B_0, B_1, \ldots, B_t$, $2t$ bundles $H_{j,1}, \ldots, H_{j,2t}$ for every $j \in [t]$, and three bundles V_i, W_i, and Z_i for every $i \in [r]$. The utilities of the agents for the bundles of partition $P_{j,c}$ for $j \in [t]$ and $c \in [3]$ are given in Table 1. The table includes only non-zero utilities; any utility that is not specified in the table is equal to zero. In our reduction, we use parameters $\varepsilon = \frac{1}{12t^2}$, $R = \frac{6t^3}{\varepsilon}$, and $Q = \frac{6t}{\varepsilon}$.

The reduction is clearly computable in polynomial time. We shall, without loss of generality, assume in the following that $t \geq 9$ and $r \leq 3t$; otherwise, it is trivial to decide ϕ.

Definition 5.2 (Canonical allocation)**.** For any partition $P_{j,c}$ with $j \in [t]$ and $c \in [3]$, we define the *canonical* allocation as follows: bundle B_k is assigned to base agent b_k for $k \in \{0, 1, \ldots, t\}$, bundle $H_{j,\ell}$ is assigned to set agent $h_{j,\ell}$ for $\ell \in [2t]$, and, finally, bundle V_i is assigned to element agent v_i, bundle W_i is assigned to element agent w_i, and Z_i is assigned to element agent z_i for $i \in [r]$.

Table 1: The reduction in the proof of Theorem 5.1.

c	agent	bundle	utility
any	b_0	B_0	R/t
	b_0	B_j	R
	b_j	B_j	R
	z_i for $i \in [r] : e_i \in S_j$	Z_i	$1/f_i$
1	$h_{j,1}$	$H_{j,1}$	Q
	v_i for $i \in [r] : e_i \in S_j$	V_i	2
	z_i for $i \in [r] : e_i \in S_j$	V_i	2/3
2	$h_{j,2}$	$H_{j,2}$	Q
	w_i for $i \in [r] : e_i \in S_j$	W_i	2
	z_i for $i \in [r] : e_i \in S_j$	W_i	$(1 + 1/f_i)/f_i$
3	$h_{j,1}$	$H_{j,1}$	$Q(1 - \varepsilon)$
	$h_{j,2}$	$H_{j,2}$	$Q(1 - \varepsilon)$
	$h_{j,\ell}$ for $\ell = 3, ..., t + 1$	$H_{j,1}$	ε
	$h_{j,\ell}$ for $\ell = t + 2, ..., 2t$	$H_{j,2}$	ε
	v_i for $i \in [r] : e_i \in S_j$	V_i	2
	w_i for $i \in [r] : e_i \in S_j$	W_i	2
	z_i for $i \in [r] : e_i \in S_j$	V_i	2/3
	z_i for $i \in [r] : e_i \in S_j$	W_i	$(1 + 1/f_i)/f_i$

5.2 Proof of Theorem 5.1

We now prove the correctness of our reduction. We remark that when we refer to the expected social welfare achieved by a set F of agents in a lottery $\mathbf{q}$, we refer to the sum of the expected utilities of agents in F in $\mathbf{q}$. We begin two two simple technical lemmas.

Lemma 5.3. *Consider an ex-ante envy-free lottery of instance $\mathcal{I}(\phi)$. For $j \in [t]$, the expected utility the set and element agents can get from each of the partitions $P_{j,1}$ or $P_{j,2}$, conditioned on the partition being the outcome of the lottery, is at most $Q + 9$. Similarly, the expected utility the set and element agents can get from the partition $P_{j,3}$, conditioned on the partition being the outcome of the lottery, is at most $Q/2 + 9$.*

Lemma 5.4. *In instance $\mathcal{I}(\phi)$, for any partition $P_{j,c}$ with $j \in [t]$ and $c \in [3]$, any allocation in the support of a Pareto-optimal lottery, either assigns bundle B_j to agent b_0 or assigns bundle B_0 to agent b_0 and bundle B_j to agent b_j.*

Proof. Consider a Pareto-optimal lottery $\mathbf{q}$ and assume, for the sake of contradiction, that it has in its support an allocation in partition $P_{j,c}$ for $j \in [t]$ and $c \in [3]$ which assigns to the base agent b_0 neither bundle B_0 nor bundle B_j. Then, since the base agent b_0 is the only one who can get positive utility from bundle B_0, the lottery $\widetilde{\mathbf{q}}$, which moves probability mass from the above allocation to the one in which the agent who gets bundle B_0 and the base agent b_0 have their bundles swapped, Pareto-dominates $\mathbf{q}$, contradicting its Pareto-optimality.

Now, assume that $\mathbf{q}$ has in its support an allocation in partition $P_{j,c}$ for $j \in [t]$ and $c \in [3]$, in which the base agent b_0 is assigned to bundle B_0 but bundle B_j is not assigned to the base agent b_j. Then, since the base agent b_j is the only agent besides b_0 who has positive utility for bundle B_j at partition $P_{j,c}$, the lottery $\widetilde{\mathbf{q}}$, which moves probability mass from this allocation to the one in which the agent who gets bundle B_j and the base agent b_j have their bundles swapped, Pareto-dominates $\mathbf{q}$, again contradicting its Pareto-optimality. The lemma follows. □

In the statements and proofs below, for a given lottery, we denote by $p_{j,c}$ the probability of partition $P_{j,c}$ being the outcome of the lottery, for $j \in [t]$ and $c \in [3]$. The next lemma shows that ex-ante envy-free lotteries of high social welfare must place close to total probability $\frac{1}{t}$ on the partitions $P_{j,1}$, $P_{j,2}$ and $P_{j,3}$, for each $j \in [t]$.

Lemma 5.5. *In any ex-ante envy-free and Pareto-optimal lottery of instance $\mathcal{I}(\phi)$ in which the base agents have social welfare at least $R+R/t+r/t-3$, it holds that $p_{j,1}+p_{j,2}+p_{j,3} \in [\frac{1-\varepsilon}{t}, \frac{1+\varepsilon}{t}]$, for each $j \in [t]$.*

Our next technical lemma shows that in Pareto-optimal lotteries, almost all of the total probability given to the two partitions $P_{j,1}$ and $P_{j,2}$ is given to one of them.

Lemma 5.6. *Any Pareto-optimal lottery in instance $\mathcal{I}(\phi)$ satisfies $\max\{p_{j,1}, p_{j,2}\} \geq (1 - \varepsilon)(p_{j,1} + p_{j,2})$, for all $j \in [t]$.*

Together, the lemmas above allow us to show that any ex-ante envy-free and Pareto-optimal lottery of high social welfare has an *almost combinatorial structure*. We remark that this is the crucial property of our reduction that essentially allows us to embed the combinatorial search space of X3C into the continuous space of allocation lotteries. Namely, for each $j \in [t]$, the lottery must give a probability mass of almost $1/t$ to one of the partitions $P_{j,1}$ or $P_{j,2}$, and a probability mass of almost 0 to the other. This is stated more precisely in the following lemma.

Lemma 5.7. *In instance $\mathcal{I}(\phi)$, any ex-ante envy-free and Pareto-optimal lottery, in which the expected social welfare of the set and element agents is at least $Q+6+r/t$ and the expected social welfare of the base agents is at least $R+R/t+r/t-3$, satisfies $\max\{p_{j,1}, p_{j,2}\} \geq \frac{1-3\varepsilon}{t}$, $\min\{p_{j,1}, p_{j,2}\} \leq \frac{2\varepsilon}{t}$, and $p_{j,3} \leq \frac{\varepsilon}{t}$ for each $j \in [t]$.*

Our arguments in the next two lemmas use a particular type of non-canonical allocations.

Definition 5.8. An allocation in partition $P_{j,c}$ for $j \in [t]$ and $c \in [3]$ is called *defective* if there is $i \in [r]$ such that $e_i \in S_j$ and agent z_i is not assigned bundle Z_i.

Our next lemma proves an upper bound on the probability mass put by any Pareto-optimal lottery with high enough social welfare on defective allocations.

Lemma 5.9. *In instance $\mathcal{I}(\phi)$, any ex-ante envy-free and Pareto-optimal lottery with social welfare at least $R + R/t + r/t - 3$ for the base agents and at least $Q + 6 + r/t$ for the set and element agents, must put a probability mass of at most 5ε on defective allocations.*

We are now ready to prove the soundness and completeness of our reduction. This is done in Lemmas 5.10 and 5.11, respectively, which complete the proof of Theorem 5.1. In the proof of the next lemma, for a given lottery, we will denote by $\gamma_{j,c}$ the probability mass put on defective allocations in partition $P_{j,c}$ for $j \in [t]$ and $c \in [3]$. We denote by γ the total probability mass put on defective allocations, i.e., $\gamma = \sum_{j\in[t]} (\gamma_{j,1} + \gamma_{j,2} + \gamma_{j,3})$.

Lemma 5.10. *If instance $\mathcal{I}(\phi)$ admits an ex-ante envy-free and Pareto-optimal lottery of social welfare at least $R + R/t + Q + 6 + r/t$, then instance ϕ has an exact cover.*

Proof. Let $\mathbf{q}$ be an ex-ante envy-free and Pareto-optimal lottery in instance $\mathcal{I}(\phi)$ with the stated social welfare guarantee. By Lemma 5.3, the expected utility set and element agents have is at most $Q + 9$. Hence, the social welfare of the base agents is at least $R + R/t + r/t - 3$ and the conditions of Lemma 5.5 are satisfied. Now, observe that for $i \in [r]$, the utility of agent z_i is at most $1/f_i$ in any non-defective allocation in partitions $P_{j,c}$ for $j \in [t]$ such that $e_i \in S_j$ and $c \in [3]$, while it is at most $\max\{1/f_i, 2/3, (1 + 1/f_i)/f_i\} \leq 1 + 1/f_i$ in any defective allocation in partition $P_{j,c}$ for $j \in [t]$ such that $e_i \in S_j$ and $c \in [3]$. Clearly, the utility of agent z_i is zero in any allocation in partition $P_{j,c}$ for $j \in [t]$ such that $e_i \notin S_j$. Thus, the expected utility of agent z_i for $i \in [r]$ is

$$\begin{aligned}
&u_{z_i}(\mathbf{q}; z_i) \\
&\leq \sum_{j\in[t]:e_i\in S_j} \Bigg((p_{j,1} + p_{j,2} + p_{j,3} - \gamma_{j,1} - \gamma_{j,2} - \gamma_{j,3}) \cdot \frac{1}{f_i} \\
&\qquad + (\gamma_{j,1} + \gamma_{j,2} + \gamma_{j,3}) \cdot \left(1 + \frac{1}{f_i}\right)\Bigg) \\
&= \sum_{j\in[t]:e_i\in S_j} (p_{j,1} + p_{j,2} + p_{j,3}) \cdot \frac{1}{f_i} \\
&\qquad + \sum_{j\in[t]:e_i\in S_j} (\gamma_{j,1} + \gamma_{j,2} + \gamma_{j,3}) \\
&\leq \frac{1+\varepsilon}{t} \sum_{j\in[t]:e_i\in S_j} \frac{1}{f_i} + \gamma \leq \frac{1}{t} + \frac{\varepsilon}{t} + 5\varepsilon < \frac{1}{t} + \frac{1}{2t^2}.
\end{aligned} \tag{3}$$

The second inequality follows by Lemma 5.5 which asserts that $p_{j,1} + p_{j,2} + p_{j,3} \leq \frac{1+\varepsilon}{t}$, the third one by the definition of f_i and Lemma 5.9, and the last one by the definition of ε (recall that $\varepsilon = \frac{1}{12t^2}$) and since $t \geq 9$.

On the other hand, notice that the bundle B_0 gives utility R/t only to the base agent b_0, while for $j \in [t]$ and $c \in [3]$, the only bundle among $B_1, B_2, ..., B_t$ that gives non-zero utility (equal to R) to some base agent is bundle B_j. Thus, the social welfare of the base agents is at most $R + R/t$ and, hence, the social welfare of the set and element agents in lottery $\mathbf{q}$ is at least $Q + 6 + r/t$. Together with the properties of ex-ante envy-freeness and Pareto-optimality and the lower bound on the social welfare of the base agents claimed above, the conditions of Lemma 5.7 are satisfied, meaning that the lottery $\mathbf{q}$ has the combinatorial structure indicated by it.

Define $C = \{j \in [t] : p_{j,1} \geq \frac{1-3\varepsilon}{t}\}$. We will show that C forms an exact cover of ϕ. For the sake of contradiction, assume otherwise that there exists an element e_{i^*} for some $i^* \in [r]$ that is included in either none or in at least two sets S_j such that $j \in C$. We distinguish between two cases:

Case 1. If e_{i^*} is not included in any set S_j such that $j \in C$, then $p_{j,1} < \frac{1-3\varepsilon}{t}$ and, by Lemma 5.7, $p_{j,2} \geq \frac{1-3\varepsilon}{t}$ for all $j \in [t]$ such that $e_{i^*} \in S_j$. Now, notice that agent w_{i^*} is assigned bundle W_{i^*} (for which agent z_{i^*} has utility $(1 + 1/f_{i^*})/f_{i^*}$) in every non-defective allocation in partition $P_{j,2}$ for $j \in [t]$ such that $e_{i^*} \in S_j$. Thus, the expected utility agent z_{i^*} has for bundle assigned to agent w_{i^*} is

$$\begin{aligned}
&u_{z_{i^*}}(\mathbf{q}; w_{i^*}) \\
&\geq \sum_{j\in[t]:e_{i^*}\in S_j} (p_{j,2} - \gamma_{j,2}) \cdot \left(1 + \frac{1}{f_{i^*}}\right) \cdot \frac{1}{f_{i^*}} \geq \left(1 + \frac{1}{t}\right) \\
&\qquad \cdot \sum_{j\in[t]:e_{i^*}\in S_j} (p_{j,2} - \gamma_{j,2}) \cdot \frac{1}{f_{i^*}} \\
&\geq \left(1 + \frac{1}{t}\right) \cdot \frac{1-3\varepsilon}{t} \cdot \sum_{j\in[t]:e_{i^*}\in S_j} \frac{1}{f_{i^*}} - \left(1 + \frac{1}{t}\right) \cdot \sum_{j\in[t]:e_{i^*}\in S_j} \frac{\gamma_{j,2}}{f_{i^*}} \\
&\geq \left(1 + \frac{1}{t}\right) \cdot \frac{1-3\varepsilon}{t} - \left(1 + \frac{1}{t}\right) \cdot \gamma \geq \frac{1}{t} + \frac{1}{t^2} - \frac{3\varepsilon}{t^2} - \frac{8\varepsilon}{t} - 5\varepsilon \\
&> \frac{1}{t} + \frac{1}{2t^2}.
\end{aligned} \tag{4}$$

The second inequality follows since $f_{i^*} \leq t$ by definition, the third one since $p_{j,2} \geq \frac{1-3\varepsilon}{t}$, the fourth one by the definitions of γ and f_{i^*}, the fifth one by Lemma 5.9, and the last one by the definition of ε (recall that $\varepsilon = \frac{1}{12t^2}$ and since $t \geq 9$. By inequalities (3) and (4), we obtain that $u_{z_{i^*}}(\mathbf{q}; z_{i^*}) < u_{z_{i^*}}(\mathbf{q}; w_{i^*})$, meaning that agent z_{i^*} is envious of agent w_{i^*}, a contradiction.

Case 2. Let $D = \{j \in C : e_{i^*} \in S_j\}$ and assume that $|D| \geq 2$. Since $D \subseteq C$, we have $p_{j,1} \geq \frac{1-3\varepsilon}{t}$ for every $j \in D$. Notice that agent v_{i^*} is assigned bundle V_{i^*} (for which agent z_{i^*} has utility 2/3) in every non-defective allocation in partition $P_{j,1}$ for $j \in [t]$ such that $e_{i^*} \in S_j$. Thus, the expected utility agent z_{i^*} has for the bundle assigned to agent v_{i^*} is

$$\begin{aligned}
u_{z_{i^*}}(\mathbf{q}; v_{i^*}) &\geq \sum_{j\in D} \frac{2}{3} \cdot (p_{j,1} - \gamma_{j,1}) \geq \frac{4}{3} \cdot \frac{1-3\varepsilon}{t} - \frac{2}{3} \cdot \gamma \\
&\geq \frac{4}{3t} - \frac{4\varepsilon}{t} - \frac{10}{3}\varepsilon \geq \frac{1}{t} + \frac{1}{2t^2}.
\end{aligned} \tag{5}$$

The third inequality follows by Lemma 5.9 and the last one by the definition of ε (recall that $\varepsilon = \frac{1}{12t^2}$) and since $t \geq 9$. By inequalities (3) and (5), we obtain that $u_{z_{i^*}}(\mathbf{q}; z_{i^*}) < u_{z_{i^*}}(\mathbf{q}; w_{i^*})$, again meaning that agent z_{i^*} is envious of agent w_{i^*}, a contradiction. □

Lemma 5.11. *If instance ϕ has an exact cover, then instance $\mathcal{I}(\phi)$ admits an ex-ante envy-free and Pareto-optimal lottery of social welfare at least $R + R/t + Q + 6 + r/t$.*

6 CONCLUSION

In this work, we considered the general setting of the problem of dividing indivisible items in a fair and efficient manner to agents having partition-based utilities. We have shown membership of the total problem of finding ex-ante envy-free and Pareto-optimal allocation lotteries in the class PPAD. We consider settling the precise computational complexity of the problem an important question. From an algorithmic perspective it would also be very interesting to see if Lemke's algorithm [34] could be adapted to solve the problem, as this would likely lead to a practical algorithm.

ACKNOWLEDGMENTS

IC and KAH were partially supported by the Independent Research Fund Denmark under grant nos. 2032-00185B and 9040-00433B, respectively.

REFERENCES

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica* 66, 3 (1998), 689–701.

[2] Adjusted Winner [n.d.]. Adjusted Winner. http://www.nyu.edu/projects/adjustedwinner/. Accessed: October 2023.

[3] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence* 322 (2023), 103965.

[4] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. 2021. Maximum Nash welfare and other Stories about EFX. *Theoretical Computer Science* 863 (2021), 69–85.

[5] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. 2020. Multiple Birds with One Stone: Beating 1/2 for EFX and GMMS via Envy Cycle Elimination. *Theoretical Computer Science* 841 (2020), 94–109.

[6] Haris Aziz. 2019. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. In *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers*. Springer International Publishing, 45–50.

[7] Haris Aziz. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE) (Lecture Notes in Computer Science, Vol. 12495)*. 341–355.

[8] Anna Bogomolnaia and Hervé Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295–328.

[9] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing Conflicts in Fair Division of Indivisible Goods using a Scale of Criteria. *Autonomous Agents and Multi Agent Systems* 30, 2 (2016), 259–290.

[10] Steven J Brams and Alan D Taylor. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press.

[11] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia. 2016. *Handbook of Computational Social Choice.* Cambridge University Press.

[12] Eric Budish. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[13] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review* 103, 2 (2013), 585–623.

[14] Ioannis Caragiannis, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, and Rohit Vaish. 2021. Stable Fractional Matchings. *Artificial Intelligence* 295 (2021), 103416.

[15] Ioannis Caragiannis, Kristoffer Arnsfelt Hansen, and Nidhi Rathi. 2023. On the complexity of Pareto-optimal and envy-free lotteries. *arXiv preprint arXiv:2307.12605* (2023).

[16] Ioannis Caragiannis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2021. On Interim Envy-Free Allocation Lotteries. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. 264–284.

[17] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12:1–12:32.

[18] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. 2020. EFX Exists for Three Agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 1–19.

[19] Bhaskar Ray Chaudhury, Kavitha Telikepalli, Kurt Mehlhorn, and Alkmini Sgouritsa. 2021. A Little Charity Guarantees Almost Envy-Freeness. *SIAM J. Comput.* 50, 4 (2021), 1336–1358.

[20] Bernard Chazelle. 1993. An Optimal Convex Hull Algorithm in Any Fixed Dimension. *Discrete & Computational Geometry* 10, 4 (1993), 377–409.

[21] Yan Chen and Tayfun Sönmez. 2002. Improving Efficiency of On-Campus Housing: An Experimental Study. *American Economic Review* 92, 5 (2002), 1669–1686.

[22] Richard Cole and Yixin Tao. 2021. On the Existence of Pareto Efficient and Envy-Free Allocations. *Journal of Economic Theory* 193 (2021), 105207.

[23] Battal Doğan and Kemal Yıldız. 2016. Efficiency and Stability of Probabilistic Assignments in Marriage Problems. *Games and Economic Behavior* 95 (2016), 47–58.

[24] L. E. Dubins and E. H. Spanier. 1961. How to Cut A Cake Fairly. *The American Mathematical Monthly* 68, 1 (1961), 1–17.

[25] Kousha Etessami and Mihalis Yannakakis. 2010. On the Complexity of Nash Equilibria and Other Fixed Points. *SIAM J. Comput.* 39, 6 (2010), 2531–2597.

[26] Raul Etkin, Abhay Parekh, and David Tse. 2007. Spectrum Sharing for Unlicensed Bands. *IEEE Journal on Selected Areas in Communications* 25, 3 (2007), 517–528.

[27] Aris Filos-Ratsikas, Kristoffer Arnsfelt Hansen, Kasper Høgh, and Alexandros Hollender. 2021. FIXP-membership via Convex Optimization: Games, Cakes, and Markets. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 827–838.

[28] Aris Filos-Ratsikas, Kristoffer Arnsfelt Hansen, Kasper Høgh, and Alexandros Hollender. 2023. PPAD-membership for Problems with Exact Rational Solutions: A General Approach via Convex Optimization. *arXiv preprint arXiv:2312.01237* (2023).

[29] Duncan Karl Foley. 1966. *Resource Allocation and the Public Sector*. Vol. 7:45-98. Yale Economic Essays.

[30] Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2020. Best of Both Worlds: Ex-ante and Ex-post Fairness in Resource Allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 21–22.

[31] Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy* 87, 2 (1979), 293–314.

[32] Shizuo Kakutani. 1941. A generalization of Brouwer's Fixed Point Theorem. *Duke Mathematical Journal* 8, 3 (1941), 457–459.

[33] Richard M. Karp. 1972. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations (The IBM Research Symposia Series)*, Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (Eds.). 85–103.

[34] C. E. Lemke. 1965. Bimatrix Equilibrium Points and Mathematical Programming. *Management Science* 11, 7 (1965), 681–689.

[35] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*. 125–131.

[36] Hervé Moulin. 2004. *Fair Division and Collective Welfare.* MIT press.

[37] Christos H. Papadimitriou. 1994. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *J. Comput. System Sci.* 48, 3 (1994), 498–532.

[38] Christos H. Papadimitriou, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Manolis Zampetakis. 2023. The Computational Complexity of Multi-player Concave Games and Kakutani Fixed Points. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC).*

[39] Hugo Steihaus. 1948. The Problem of Fair Division. *Econometrica* 16 (1948), 101–104.

[40] Walter Stromquist. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly* 87, 8 (1980), 640–644.

[41] Francis Edward Su. 1999. Rental Harmony: Sperner's Lemma in Fair Division. *The American Mathematical Monthly* 106, 10 (1999), 930–942.

[42] Hal R. Varian. 1974. Equity, Envy, and Efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.

[43] Vijay V. Vazirani and Mihalis Yannakakis. 2021. Computational Complexity of the Hylland-Zeckhauser Scheme for One-Sided Matching Markets. In *Proceedings of the 12th Innovations in Theoretical Computer Science Conference (ITCS)*. 59:1–59:19.

[44] Thomas Vossen. 2002. *Fair Allocation Concepts in Air Traffic Management.* Ph.D. Dissertation. University of Maryland.

[45] L. Zadeh. 1963. Optimality and Non-Scalar-Valued Performance Criteria. *IEEE Trans. Automat. Control* 8, 1 (1963), 59–60.