

On Fair and Efficient Allocations of Indivisible Goods

Aniket Murhekar, Jugal Garg

University of Illinois, Urbana-Champaign
{aniket2, jugal}@illinois.edu

Abstract

We study the problem of fair and efficient allocation of a set of indivisible goods to agents with additive valuations using the popular fairness notions of envy-freeness up to one good (EF1) and equitability up to one good (EQ1) in conjunction with Pareto-optimality (PO). There exists a pseudo-polynomial time algorithm to compute an EF1+PO allocation, and a non-constructive proof of existence of allocations that are both EF1 and fractionally Pareto-optimal (fPO). We present a pseudo-polynomial time algorithm to compute an EF1+fPO allocation, thereby improving the earlier results. Our techniques also enable us to show that an EQ1+fPO allocation always exists when the values are positive, and that it can be computed in pseudo-polynomial time.

We also consider the class of k -ary instances where k is a constant, i.e., each agent has at most k different values for the goods. We show that for such instances an EF1+fPO allocation can be computed in polynomial time. When all values are positive, we show that an EQ1+fPO allocation for such instances can be computed in polynomial time. Next, we consider instances where the number of agents is constant, and show that an EF1+PO (also EQ1+PO) allocation can be computed in polynomial time. These results significantly extend the polynomial-time computability beyond the known cases of binary or identical valuations.

Further, we show that the problem of computing an EF1+PO allocation polynomial-time reduces to a problem in the complexity class PLS. We also design a polynomial-time algorithm that computes Nash welfare maximizing allocations when there are constantly many agents with constant many different values for the goods.

Introduction

The problem of fair division was formally introduced by Steinhaus (1948), and has since been extensively studied in various fields, including economics and computer science (Brams and Taylor 1996; Moulin 2003). It concerns allocating resources to agents in a *fair* and *efficient* manner, and has various practical applications such as rent division, division of inheritance, course allocation and government auctions. Much of the work has focused on *divisible* goods, which can be shared between agents. In this setting, several appealing

fairness concepts like *envy-freeness* have been defined (Foley 1967; Varian 1974), which requires that every agent prefer their own bundle of goods to that of any other. On the other hand, when the goods are *indivisible*, envy-free allocations need not even exist, e.g., in the simple case of one good and two agents. Other classical notions of fairness, like *equitability* and *proportionality*, may also be impossible to satisfy when goods are indivisible. However, fair division of indivisible goods remains an important problem since goods cannot always be shared, and also because it models several practical scenarios such as a course allocation (Othman et al. 2010). Since allocations satisfying standard fairness criteria fail to exist in the case of indivisible goods, several weaker fairness notions have been defined.

A relaxation of envy-freeness, called *envy-freeness up to one good* (EF1) was defined by Budish (2011). An allocation is said to be EF1 if every agent prefers their own bundle over the bundle of any other agent after removing at most one good from the other agent’s bundle. When the valuations of the agents for the goods are monotone, EF1 allocations exist and are polynomial time computable (Lipton et al. 2004).

The standard notion of economic efficiency is Pareto optimality (PO). An allocation is said to be PO if no other allocation makes an agent better off without making someone else worse off. A natural question to ask is whether EF1 can be achieved in conjunction with PO under additive valuations. The concept of *Nash welfare* provides a positive answer to this question. The Nash welfare is the geometric mean of the agents’ utilities, and by maximizing it we achieve a tradeoff between efficiency and fairness. Caragiannis et al. (2016) showed that any maximum Nash welfare (MNW) allocation is EF1 and PO. For binary valuations, the MNW allocation can be computed in polynomial time (Barman et al. (2018b)). However, in general, the problem of computing the MNW allocation is APX-hard (Lee 2015).

Bypassing this barrier, Barman, Krishnamurthy, and Vaish (2018a) devised a pseudo-polynomial time algorithm that computes an allocation that is both EF1 and PO. They also showed that allocations that are both EF1 and *fractionally* Pareto-optimal (fPO) always exist, where an allocation is said to be fPO if no fractional allocation exists that makes an agent better off without making anyone else worse off. They showed this result via a non-constructive convergence argument used in real analysis, and did not provide an al-

gorithm for computing such an allocation. Clearly, fPO is a stronger notion of economic efficiency, and hence the problem of computing EF1+fPO allocations is an important one. Another reason to prefer fPO allocations over PO allocations in practice is that the former property admits efficient verification; whereas checking if an allocation is PO is known to be co-NP-complete (de Keijzer et al. 2009). This means that in scenarios where a centralized entity is responsible for the allocation, all participants can efficiently verify if an allocation is fPO (and thus PO). However, in general, the same efficient verification is not possible for PO allocations.

In this paper, we present a pseudo-polynomial time algorithm that computes an allocation that is EF1+fPO. Not only does this improve the result of Barman et al., but it also provides other interesting insights. We consider the class of k -ary instances for a constant k , i.e., each agent has at most k different values for the goods. Our analysis shows that for such instances, an EF1+fPO allocation can be found in polynomial time. Our result becomes especially interesting in light of the fact that computing the MNW allocation remains NP-hard for such instances (Lee 2015), even for $k = 3$ (Amanatidis et al. 2020). Further, this is the first class apart from binary or identical valuations for which EF1+fPO allocations are polynomial time computable.

While k -ary instances are interesting from a theoretical point of view to understand the limits of tractability in computing fair and efficient allocations, they are also relevant from a practical perspective. Eliciting the values that agents have for goods is often a tricky task, as agents may not be able to assert exactly what values they have for different goods. A simple protocol that the entity in charge of the allocation can do is to ask each agent to “rate” the goods using a few (constantly many) values. Based on these responses, the valuations of the agents can be established.

Equitability up to one good (EQ1) is a generalization of the classical fairness notion of equitability. An allocation is said to be EQ1 if the utility an agent gets from her bundle is no less than the utility any other agent gets after removing one specific good from their bundle. Using similar techniques to that of Barman et al. (2018a), a pseudo-polynomial time algorithm to compute an EQ1+PO allocation was developed by Freeman et al. (2019), when all the values are positive. We show the stronger result that EQ1+fPO allocations always exist for positive-valued instances, and can in fact be computed in pseudo-polynomial time. Our techniques also show that for k -ary instances with positive values where k is a constant, an allocation that is EQ1 and fPO can be computed in polynomial time.

In many practical applications, the number of agents n is constant. We show that for constant n , an EF1+PO allocation can be computed in time polynomial in the number of goods for the case of general additive valuations. In contrast, computing the MNW allocation remains NP-hard for $n = 2$. Our techniques also show that for constant n , an EQ1+PO allocation can also be computed in polynomial time. Further, when we consider k -ary instances with constant n and k , we show that many fair division problems, including computing the MNW allocation, have polynomial time complexity. This improves the result of Bliem et al. (2016) who show

that the EF+PO is tractable in this case.

We also make progress on the complexity-front. We prove that the problem of computing an EF1+PO allocation polynomial-time reduces to solving a problem in the complexity class Polynomial Local Search (PLS). For this, we carefully analyze Barman et al. (2018a)’s algorithm and show that it has the structure of local-search problem. Finally, we remark that our techniques also improve the results of (Chakraborty et al. 2020) and (Freeman et al. 2020) for the problems of computing weighted-EF1+fPO allocations of goods and EQ1+fPO allocations of *chores*, respectively.

Related work. Barman et al. (2018a) devised a pseudo-polynomial time algorithm that computes an allocation that is both EF1 and PO. This algorithm runs in time $\text{poly}(n, m, v_{max})$, where n is the number of agents, m the number of items, and v_{max} the maximum utility value. Their algorithm first perturbs the values to a desirable form, and then computes an EF1 and fPO allocation for the perturbed instance. Their approach is via *integral market-equilibria*, which guarantees fPO at every step. The *spending* of an agent, which is the sum of prices of the goods she owns in the equilibrium, works as a proxy for her utility. The returned allocation is approximately-EF1 and approximately-PO for the original instance, which for a fine-enough approximation is EF1 and PO. Our algorithm proceeds in a similar manner to their algorithm, with one main difference being that we do not need to consider any approximate instance and can work directly with the given valuations. Our algorithm returns an allocation that is not only PO, but is in fact fPO. Another key difference is the *run-time analysis*: while their analysis relies on bounding the number of steps using arguments about *prices*, our analysis is more direct and works with the *values*. This allows us to prove a general Theorem 1, of which polynomial run-time for k -ary instances with constant k is a consequence. Directly, such a conclusion cannot be drawn from the analysis of Barman et al. Another technical difference is that we raise the prices of multiple components of *least spenders* simultaneously, unlike in Barman et al. (2018a).

Barman et al. showed that there is a non-deterministic algorithm that computes an EF1+fPO allocation, since checking if an allocation is fPO can be done efficiently. In contrast, we present a deterministic algorithm computing an EF1+fPO allocation, albeit with worst-case pseudopolynomial run time. We also show that finding an EF1+PO allocation reduces to a problem in PLS. However this doesn’t place it in PLS, since the problem is not even in the class *Total Function NP* \supseteq PLS (Barman et al.).

Freeman et al. (2019) presented a pseudo-polynomial time algorithm for computing an EQ1+PO allocation for instances with positive values. Since they consider an approximate instance too, their algorithm does not achieve the guarantee of fPO. They also show that the leximin solution, i.e., the allocation that maximizes the minimum utility, and subject to this, maximizes the second minimum utility, and so on; is EQX (a stronger requirement than EQ1) and PO for positive utilities. However, as remarked in (Plaut and Rough-

garden 2018), computing a leximin solution is intractable.

Bredereck et al. (2019) presented a framework for fixed-parameter algorithms for many fairness concepts, including EF1+PO, parameterized by $n + v_{max}$. Our results improve their findings and show that the problem has polynomial time complexity for the cases of (i) constant n , (ii) constant number of utility values, (iii) v_{max} bounded by $\text{poly}(m, n)$.

Preliminaries

For $t \in \mathbb{N}$, let $[t]$ denote $\{1, \dots, t\}$.

Problem setting. A *fair division instance* is a tuple (N, M, V) , where $N = [n]$ is a set of $n \in \mathbb{N}$ agents, $M = [m]$ is the set of $m \in \mathbb{N}$ indivisible items, and $V = \{v_1, \dots, v_n\}$ is a set of utility functions, one for each agent $i \in N$. Each utility function $v_i : M \rightarrow \mathbb{Z}_{\geq 0}$ is specified by m numbers $v_{ij} \in \mathbb{Z}_{\geq 0}$, one for each good $j \in M$, which denotes the value agent i has for good j . We assume that the valuation functions are additive, that is, for every agent $i \in N$, and for $S \subseteq M$, $v_i(S) = \sum_{j \in S} v_{ij}$. Further, we assume that for every good j , there is some agent i such that $v_{ij} > 0$. Note that we can in general work with rational values, since they can be scaled to make them integral.

We call a fair division instance (N, M, V) a:

1. *Binary* instance if for all $i \in N$ and $j \in M$, $v_{ij} \in \{0, 1\}$.
2. *k-ary* instance, if $\forall i \in N, |\{v_{ij} : j \in M\}| \leq k$.
3. *Positive-valued* instance if $\forall i \in N, \forall j \in M, v_{ij} > 0$.

Note that the class of k -ary instances generalizes the class of k -valued instances as defined in Amanatidis et al. (2020), in which all values belong to a k -sized set, whereas we allow each agent to have k different values for the goods.

Allocation. An *allocation* \mathbf{x} of goods to agents is a n -partition of the goods $\mathbf{x}_1, \dots, \mathbf{x}_n$, where agent i is allotted $\mathbf{x}_i \subseteq M$, and gets a total value of $v_i(\mathbf{x}_i)$. A *fractional allocation* $\mathbf{x} \in [0, 1]^{n \times m}$ is a fractional assignment such that for each good $j \in M$, $\sum_{i \in N} x_{ij} \leq 1$. Here, $x_{ij} \in [0, 1]$ denotes the fraction of good j allotted to agent i .

For an agent $i \in N$, let U_i be the number of different utility values i can get in any allocation. Let $U = \max_{i \in N} U_i$.

Fairness notions. An allocation \mathbf{x} is said to be:

1. *Envy-free up to one good* (EF1) if for all $i, h \in N$, there exists a good $j \in \mathbf{x}_h$ s.t. $v_i(\mathbf{x}_i) \geq v_i(\mathbf{x}_h \setminus \{j\})$.
2. *Equitable up to one good* (EQ1) if for all $i, h \in N$, there exists a good $j \in \mathbf{x}_h$ s.t. $v_i(\mathbf{x}_i) \geq v_h(\mathbf{x}_h \setminus \{j\})$.

Pareto-optimality. An allocation \mathbf{y} dominates an allocation \mathbf{x} if $v_i(\mathbf{y}_i) \geq v_i(\mathbf{x}_i), \forall i$ and there exists h s.t. $v_h(\mathbf{y}_h) > v_h(\mathbf{x}_h)$. An allocation is said to be *Pareto optimal* (PO) if no allocation dominates it. Further, an allocation is said to be fractionally PO (fPO) if no fractional allocation dominates it. Thus, a fPO allocation is PO, but not vice-versa.

Maximum Nash Welfare. The Nash welfare of an allocation \mathbf{x} is given by $\text{NW}(\mathbf{x}) = (\prod_{i \in N} v_i(\mathbf{x}_i))^{1/n}$. An allocation that maximizes the NW is called a MNW allocation.

Fisher markets. A Fisher market or a *market instance* is a tuple (N, M, V, e) , where the first three terms are interpreted

as before, and $e = \{e_1, \dots, e_n\}$ is the set of agents' budgets, where $e_i \geq 0$, for each $i \in N$. In this model, goods can be allocated fractionally. Given prices $\mathbf{p} = (p_1, \dots, p_m)$ of goods, each agent aims to obtain the set of goods that maximizes her total value subject to her budget constraint.

A *market outcome* is a (fractional) allocation \mathbf{x} of the goods to the agents and a set of prices \mathbf{p} of the goods. The spending of an agent i under the market outcome (\mathbf{x}, \mathbf{p}) is given by $\mathbf{p}(\mathbf{x}_i) = \sum_{j \in M} p_j x_{ij}$. For an agent i , we define the *bang-per-buck* ratio α_{ij} of good j as v_{ij}/p_j , and the *maximum bang-per-buck* (MBB) ratio $\alpha_i = \max_j \alpha_{ij}$. We define $\text{mbb}_i = \{j \in M : v_{ij}/p_j = \alpha_i\}$, called the *MBB-set*, to be the set of goods that give MBB to agent i at prices \mathbf{p} . A market outcome (\mathbf{x}, \mathbf{p}) is said to be '*on MBB*' if for all agents i and goods j , $x_{ij} > 0 \rightarrow j \in \text{mbb}_i$. For integral \mathbf{x} , this means $\mathbf{x}_i \subseteq \text{mbb}_i$ for all $i \in N$.

A market outcome (\mathbf{x}, \mathbf{p}) is said to be a *market equilibrium* if (i) the market clears, i.e., all goods are fully allocated. Thus, for all j , $\sum_{i \in N} x_{ij} = 1$, (ii) budget of all agents is exhausted, for all $i \in N$, $\sum_{j \in M} x_{ij} p_j = e_i$, and (iii) agents only spend money on goods that give them maximum bang-per-buck, i.e., (\mathbf{x}, \mathbf{p}) is on MBB. Given a market outcome (\mathbf{x}, \mathbf{p}) with \mathbf{x} integral, we say it is *price envy-free up to one good* (pEF1) if for all $i, h \in N$ there is a good $j \in \mathbf{x}_h$ such that $\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_h \setminus \{j\})$. For integral market outcomes on MBB, the pEF1 condition implies the EF1 condition.

Lemma 1. *Let (\mathbf{x}, \mathbf{p}) be an integral market outcome on MBB. If (\mathbf{x}, \mathbf{p}) is pEF1 then \mathbf{x} is EF1 and fPO.*

Proof. We first show that (\mathbf{x}, \mathbf{p}) forms a market equilibrium for the Fisher market instance (N, M, V, e) , where for every $i \in N$, $e_i = \mathbf{p}(\mathbf{x}_i)$. It is easy to see that the market clears and the budget of every agent is exhausted. Further \mathbf{x} is on MBB as per our assumption. Now the fact that \mathbf{x} is fPO follows from the First Welfare Theorem (Mas-Colell et al. 1995), which shows that for any market equilibrium (\mathbf{x}, \mathbf{p}) , the allocation \mathbf{x} is fPO.

Since (\mathbf{x}, \mathbf{p}) is pEF1, for all pairs of agents $i, h \in N$, there is some good $j \in \mathbf{x}_h$ s.t. $\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_h \setminus \{j\})$. Since (\mathbf{x}, \mathbf{p}) is on MBB, $\mathbf{x}_i \subseteq \text{mbb}_i$. Let α_i be the MBB-ratio of i at the prices \mathbf{p} . By definition of MBB, $v_i(\mathbf{x}_i) = \alpha_i \mathbf{p}(\mathbf{x}_i)$, and $v_h(\mathbf{x}_h \setminus \{j\}) \leq \alpha_i \mathbf{p}(\mathbf{x}_h \setminus \{j\})$. Combining these, we get that \mathbf{x} is EF1. \square

Given a price vector \mathbf{p} , we define the MBB graph to be the bipartite graph $G = (N, M, E)$ where for an agent i and good j , $(i, j) \in E$ iff $j \in \text{mbb}_i$. Such edges are called *MBB*. Given an accompanying allocation \mathbf{x} , we supplement G to include *allocation edges*, an edge between i and j if $j \in \mathbf{x}_i$.

We call an agent i with minimum $\mathbf{p}(\mathbf{x}_i)$ a *least spender* (LS), and denote by L the set of least spenders. For agents i_0, \dots, i_ℓ and goods j_1, \dots, j_ℓ , consider a path $P = (i_0, j_1, i_1, j_2, \dots, j_\ell, i_\ell)$ in the supplemented MBB graph, where for all $1 \leq \ell' \leq \ell$, $j_{\ell'} \in \text{mbb}_{i_{\ell'-1}} \cap \mathbf{x}_{i_{\ell'}}$. Define the *level* of an agent h to be the length of the shortest such path from the LS to h , and to be n if no such path exists. Define *alternating paths* to be such paths where the edges are between agents at a lower level to agents at a strictly

Algorithm 1 Computing an EF1+fPO allocation of goods

Input: Fair division instance (N, M, V) **Output:** An integral allocation \mathbf{x}

- 1: $(\mathbf{x}, \mathbf{p}) \leftarrow$ initial welfare maximizing integral market allocation, where $p_j = v_{ij}$ for $j \in \mathbf{x}_i$.
 - 2: $L \leftarrow \{i \in N : i \in \operatorname{argmin}_{h \in N} \mathbf{p}(\mathbf{x}_h)\}$ \triangleright set of LS
 - 3: **if** $\exists i \in L, \exists$ alternating path $(i, j_1, i_1, \dots, j_\ell, i_\ell)$, s.t. $\mathbf{p}(\mathbf{x}_{i_\ell} \setminus \{j_\ell\}) > \mathbf{p}(\mathbf{x}_i)$ **then**
 - 4: Transfer j_ℓ from i_ℓ to $i_{\ell-1}$
 - 5: Repeat from Line 2
 - 6: **if** (\mathbf{x}, \mathbf{p}) is pEF1 **then return** \mathbf{x}
 - 7: **else**
 - 8: $\gamma_1 = \min_{h \in C_L \cap N, j \in M \setminus C_L} \frac{\alpha_h}{v_{hj}/p_j}$ \triangleright
Factor by which prices of goods in C_L are raised until a new MBB edge appears from an agent in C_L to a good outside C_L
 - 9: $\gamma_2 = \min_{i \in L, h \in N \setminus C_L} \frac{\mathbf{p}(\mathbf{x}_h)}{\mathbf{p}(\mathbf{x}_i)}$ \triangleright
Factor by which prices of goods in C_L are raised until a new agent outside C_L becomes a new LS
 - 10: $\beta = \min(\gamma_1, \gamma_2)$
 - 11: **for** $j \in C_L \cap M$ **do**
 - 12: $p_j \leftarrow \beta p_j$
 - 13: Repeat from Line 2
-

higher level. The edges in an alternating path alternate between MBB edges and allocation edges. For a least spender i , define C_i^ℓ to be the set of all goods and agents which lie on alternating paths of length ℓ . Call $C_i = \bigcup_\ell C_i^\ell$ the *component* of i , the set of all goods and agents reachable from the least spender i through alternating paths.

Finding EF1+fPO Allocations of Goods

We now present the main algorithm of our paper. Given a fair division instance (N, M, V) , our algorithm returns an allocation \mathbf{x} that is EF1 and fPO. We show Algorithm 1 terminates in time $\operatorname{poly}(n, m, U)$ with an EF1+fPO allocation. Detailed proofs are available in the full version of the paper.

Algorithm 1 starts with a welfare maximizing integral allocation (\mathbf{x}, \mathbf{p}) , where $p_j = v_{ij}$ for $j \in \mathbf{x}_i$. The algorithm then explores if there is an alternating path $P = (i, j_1, i_1, \dots, j_\ell, i_\ell = h)$, from some LS agent $i \in L$, such that $\mathbf{p}(\mathbf{x}_h \setminus \{j_\ell\}) > \mathbf{p}(\mathbf{x}_i)$, i.e., an alternating path along which the pEF1 condition is violated for the LS agent w.r.t the good j_ℓ . When such a path is encountered, the algorithm transfers j_ℓ from h to $i_{\ell-1}$. When there is no such path from i , the component C_i of the LS agent i is pEF1. We denote by C_L the union of all components of LS agents. Suppose the overall allocation is not pEF1, then the algorithm raises the prices of all goods in the C_L until either (i) a new MBB edge gets added from an agent $h \in C_L$ to a good $j \notin C_L$ (corresponding to a price rise of γ_1), or (ii) the spending of an agent $h \notin C_L$ becomes equal to the spending of the agents in L (corresponding to a price rise of γ_2). The algorithm then proceeds as before from Line 2.

We will use the terms *time-step* or *iteration* interchangeably to denote either a transfer or a price rise step. We say ‘at

time-step t' , to refer to the state of the algorithm *just before* the event at t happens. We denote by $(\mathbf{x}^t, \mathbf{p}^t)$ the allocation and price vector at time-step t . First we note that:

Lemma 2. *At any time-step t , $(\mathbf{x}^t, \mathbf{p}^t)$ is on MBB.*

This follows from the fact that goods are always transferred along MBB edges, and the prices are raised only until a new MBB edge is created. Thus, the MBB condition is never violated for any agent and the allocation is always on MBB throughout the run of the algorithm.

If Algorithm 1 terminates, then the final outcome (\mathbf{x}, \mathbf{p}) is pEF1. Since it is also on MBB, by Lemma 1, \mathbf{x} is EF1+fPO. We now proceed towards the run-time analysis of Algorithm 1. First we observe that since prices are raised only until the spending of a new agent becomes equal to the spending of the least spenders:

Lemma 3. *The spending of the least spender(s) does not decrease as the algorithm progresses. Further at any price rise event t with price-rise factor β , the spending of the least spender(s) increases by a factor of β .*

Next we argue:

Lemma 4. *The number of iterations with the same set of least spenders is $\operatorname{poly}(n, m)$.*

Proof. Let us fix a set L of least spenders. We count the number of alternating paths from some $i \in L$ to an agent k who owns a good j which is then transferred to an agent h . The number of such paths is at most $n \cdot n \cdot m \cdot n = n^3 m$, thus there are at most $\operatorname{poly}(n, m)$ transfers with the same set of LS in the absence of price-rise steps. Further, there can be at most n price-rise steps without any change in L . \square

The next lemma is key. We argue that between the time-steps at which an agent i ceases to be a LS and subsequently becomes a LS again, her utility strictly increases.

Lemma 5. *Let t_0 be a time-step where agent i ceases to be a LS, and let t_ℓ be the first subsequent time step just after which i becomes the LS again. Then:*

$$v_i(\mathbf{x}_i^{t_\ell+1}) > v_i(\mathbf{x}_i^{t_0})$$

Note here that $v_i(\mathbf{x}_i^{t_0})$ is the utility of agent i just before time-step t_0 , and $v_i(\mathbf{x}_i^{t_\ell+1})$ her utility just after time-step t_ℓ .

Proof. From Lemma 3, since i ceases to be a LS after time-step t_0 , i must have received some good j at time step t_0 . Since $j \in mbb_i$ at t_0 , $v_{ij} > 0$. Suppose i does not lose any good in any subsequent iterations until t_ℓ , then $\mathbf{x}_i^{t_\ell+1} \supseteq \mathbf{x}_i^{t_0} \cup \{j\}$, and hence $v_i(\mathbf{x}_i^{t_\ell+1}) \geq v_i(\mathbf{x}_i^{t_0} \cup \{j\}) = v_i(\mathbf{x}_i^{t_0}) + v_{ij} > v_i(\mathbf{x}_i^{t_0})$, using additivity of valuations.

On the other hand suppose i does lose some goods between t_0 and t_ℓ . Let $t_k \in (t_0, t_\ell]$ be the last time-step when i loses a good, say j' . Let t_1, \dots, t_{k-1} be time-steps (in order) between t_0 and t_k when i experiences price-rise, and $t_{k+1}, \dots, t_{\ell-1}$ be time-steps (in order) between t_k and t_ℓ when i experiences price-rise, until finally after the event t_ℓ agent i becomes the LS again. Let us define β_t to be the price-rise factor at the time-step t . If t is a price-rise step, $\beta_t > 1$, else we set $\beta_t = 1$. Hence

$\beta_{t_1}, \dots, \beta_{t_{k-1}}, \beta_{t_{k+1}}, \dots, \beta_{t_{\ell-1}}$ are price-rise factors at the corresponding events $t_1, \dots, t_{k-1}, t_{k+1}, \dots, t_{\ell-1}$ and are all greater than 1. If t_ℓ is a price-rise event, let the price-rise factor be $\beta_{t_\ell} > 1$; and if not let $\beta_{t_\ell} = 1$. Note that t_k is not a price-rise event and hence $\beta_{t_k} = 1$.

Using Lemma 3, together with the fact that i does not lose any good after t_k , we have:

$$\mathbf{p}^{t_\ell+1}(\mathbf{x}_i^{t_\ell+1}) \geq (\beta_{t_\ell}\beta_{t_{\ell-1}} \cdots \beta_{t_{k+1}})\mathbf{p}^{t_k}(\mathbf{x}_i^{t_k} \setminus \{j'\}). \quad (1)$$

The above may not be an equality because in addition to experiencing price-rises during t_{k+1}, \dots, t_ℓ , agent i may also gain some new good. If i_k is a LS at t_k , then for agent i to lose the good j' it must be the case that:

$$\mathbf{p}^{t_k}(\mathbf{x}_i^{t_k} \setminus \{j'\}) > \mathbf{p}^{t_k}(\mathbf{x}_i^{t_k}). \quad (2)$$

Let i_t be a LS at time-step t . Then by repeatedly applying Lemma 3, we get:

$$\begin{aligned} \mathbf{p}^{t_k}(\mathbf{x}_i^{t_k}) &\geq \beta_{t_k-1}\mathbf{p}^{t_k-1}(\mathbf{x}_i^{t_k-1}) \\ &\geq \cdots \geq (\beta_{t_k-1}\beta_{t_k-2} \cdots \beta_1)\mathbf{p}^{t_0}(\mathbf{x}_i^{t_0}) \\ &\geq (\beta_{t_k-1}\beta_{t_k-2} \cdots \beta_{t_1})\mathbf{p}^{t_0}(\mathbf{x}_i^{t_0}), \end{aligned} \quad (3)$$

where the last transition follows from the facts that (i) each $\beta_t \geq 1$, (ii) $\{t_1, \dots, t_{k-1}\} \subseteq \{1, \dots, t_k-1\}$, and (iii) $i_{t_0} = i$, since i is a least spender at t_0 . Putting (1), (2) and (3) together, we get:

$$\mathbf{p}^{t_\ell+1}(\mathbf{x}_i^{t_\ell+1}) > (\prod_{r=1}^{\ell} \beta_{t_r})\mathbf{p}^{t_0}(\mathbf{x}_i^{t_0}). \quad (4)$$

Let α_i^t denote the MBB-ratio of i at the time step t . Observe that in every price rise event with price rise factor β , the MBB ratio of any agent experiencing the price rise decreases by a factor β . Further, the MBB ratio of any agent does not change unless she experiences a price-rise step. Thus:

$$\alpha_i^{t_\ell+1} = \frac{\alpha_i^{t_0}}{(\beta_{t_\ell}\beta_{t_{\ell-1}} \cdots \beta_{t_{k+1}})(\beta_{t_k-1}\beta_{t_k-2} \cdots \beta_{t_1})}. \quad (5)$$

Therefore using the fact that the allocation is on MBB edges, and with (4) and (5), we have:

$$\begin{aligned} v_i(\mathbf{x}_i^{t_\ell+1}) &= \alpha_i^{t_\ell+1}\mathbf{p}^{t_\ell+1}(\mathbf{x}_i^{t_\ell+1}) && (\mathbf{x}_i^{t_\ell+1} \text{ is on MBB}) \\ &> \frac{\alpha_i^{t_0}}{(\prod_{r=1}^{\ell} \beta_{t_r})}(\prod_{r=1}^{\ell} \beta_{t_r})\mathbf{p}^{t_0}(\mathbf{x}_i^{t_0}) && (\text{From (4) and (5)}) \\ &= \alpha_i^{t_0}\mathbf{p}^{t_0}(\mathbf{x}_i^{t_0}) = v_i(\mathbf{x}_i^{t_0}), && (\mathbf{x}_i^{t_0} \text{ is on MBB}) \end{aligned}$$

as claimed. \square

Using the above lemmas, we show:

Lemma 6. *Algorithm 1 terminates in time $\text{poly}(n, m, U)$.*

Proof. Consider any agent i . From Lemma 5, it is clear that every time i becomes the LS again her utility has strictly increased compared to her utility the last time she was a LS. The number of utility values that i can have is U_i , and hence we conclude that the number of times she stops being an LS and becomes LS again is at most U_i . Since there are n agents, and each agent i can become the LS again at most U_i times, we have that after $\text{poly}(n, \max_{i \in N} U_i)$ changes in the

set of least spenders, there will be no changes further in the set of least spenders. After this, in at most n more price-rise steps, either the allocation becomes pEF1 or all agents get added to C_L , since no new agent becomes a LS on raising prices. Further, the number of transfers with the same set of least spenders is at most $\text{poly}(n, m)$ (Lemma 4). This shows that Algorithm 1 terminates in time $\text{poly}(n, m, U)$. \square

Putting it all together, we conclude:

Theorem 1. *Let $I = (N, M, V)$ be a fair division instance. Then an allocation that is both EF1 and fPO can be computed in time $\text{poly}(n, m, U)$.*

Observe that in any allocation and for any agent, the minimum utility is 0, and the maximum utility is mv_{max} , where $v_{max} = \max_{i,j} v_{ij}$. Since the utility values are integral, we have $U \leq mv_{max} + 1$. Thus, Algorithm 1 computes an EF1+fPO allocation in pseudo-polynomial time.

Theorem 2. *Given a fair division instance $I = (N, M, V)$, an allocation that is both EF1 and fPO can be computed in time $\text{poly}(n, m, v_{max})$, where $v_{max} = \max_{i,j} v_{ij}$. In particular, when $v_{max} \leq \text{poly}(n, m)$, an EF1+fPO allocation can be computed in $\text{poly}(n, m)$ time.*

The guarantee of EF1+fPO offered by our algorithm is stronger than the guarantee of EF1+PO provided by the algorithm of Barman et al. (2018a). We next turn our attention to k -ary instances where k is a constant. First we observe that for such instances, the maximum number of different utility values any agent can get is at most $\text{poly}(m)$.

Lemma 7. *In a k -ary fair division instance (N, M, V) with constant k , $U \leq \text{poly}(m)$.*

Proof. For any agent i , let $\{v_i^\ell\}_{\ell \in [k]}$ be the different utility values i has for the goods. In an allocation \mathbf{x} , let $m_i^\ell \in \mathbb{Z}_{\geq 0}$ be the number of goods in \mathbf{x}_i with value v_i^ℓ . Then agent i 's utility is simply: $v_i(\mathbf{x}_i) = m_i^1 v_i^1 + \cdots + m_i^k v_i^k$. Since each $0 \leq m_i^\ell \leq m$, the number of possible utility values that i can get in any allocation is at most $(m+1)^k$, which is $\text{poly}(m)$ since k is constant. Thus $U \leq \text{poly}(m)$. \square

Therefore, using Lemma 7, Theorem 1 gives:

Theorem 3. *Given a k -ary fair division instance $I = (N, M, V)$ where k is a constant, an allocation that is both EF1 and fPO can be computed in time $\text{poly}(n, m)$.*

Remark 1. *We note that our techniques can also be used to show that an allocation that is weighted-EF1 (wEF1) and fPO exists and can be computed in pseudo-polynomial time.*

Finding EQ1+fPO Allocations of Goods

We now show that Algorithm 2 finds an EQ1+fPO allocation given a fair division instance with positive values. We require the values to be positive because instances with zero values might not even admit an allocation that is EQ1+PO (Freeman et al. 2019). Algorithm 2 is similar to Algorithm 1, except that it works with values instead of spendings of agents since EQ1 allocations are desired as opposed to EF1.

Algorithm 2 starts with a welfare maximizing integral allocation (\mathbf{x}, \mathbf{p}) , where $p_j = v_{ij}$ for $j \in \mathbf{x}_i$. We refer to the

Algorithm 2 Computing an EQ1+fPO allocation of goods

Input: Positive-valued fair division instance (N, M, V) **Output:** An integral allocation \mathbf{x}

- 1: $(\mathbf{x}, \mathbf{p}) \leftarrow$ initial welfare maximizing integral market allocation, where $p_j = v_{ij}$ for $j \in \mathbf{x}_i$.
 - 2: $L \leftarrow \{i \in N : i \in \operatorname{argmin}_{h \in N} v_h(\mathbf{x}_h)\}$ \triangleright set of LU
 - 3: **if** $\exists i \in L, \exists$ alt. path $(i, j_1, i_1, \dots, j_\ell, i_\ell)$, s.t. $v_{i_\ell}(\mathbf{x}_{i_\ell} \setminus \{j_\ell\}) > v_i(\mathbf{x}_i)$ **then**
 - 4: Transfer j_ℓ from i_ℓ to $i_{\ell-1}$
 - 5: Repeat from Line 2
 - 6: **if** \mathbf{x} is EQ1 **then return** \mathbf{x}
 - 7: **else**
 - 8: $\beta = \min_{h \in C_L \cap N, j \in M \setminus C_L} \frac{\alpha_h}{v_{hj}/p_j} \triangleright$ Factor by which prices of goods in C_L are raised until a new MBB edge appears from an agent in C_L to a good outside C_L
 - 9: **for** $j \in C_L \cap M$ **do**
 - 10: $p_j \leftarrow \beta p_j$
 - 11: Repeat from Line 2
-

agent(s) with the least utility as the LU agent(s), and let L be the set of LU agents. The algorithm first explores if there is an alternating path $P = (i = i_0, j_1, i_1, \dots, j_\ell, i_\ell = h)$, from some LU agent $i \in L$, such that $v_h(\mathbf{x}_h \setminus \{j_\ell\}) > v_i(\mathbf{x}_i)$, i.e., an alternating path along which the EQ1 condition is violated for the LU agent. When such a path is encountered, the algorithm *transfers* j_ℓ from h to $i_{\ell-1}$. When there is no such path from i , the component C_i of the LS agent i is EQ1. Let C_L be the union of components of the LU agents. Suppose the overall allocation is not EQ1, then the algorithm *raises the prices* of all goods in the C_L until a new MBB edge gets added from an agent $h \in C_L$ to a good $j \notin C_L$. The algorithm then proceeds as before from Line 2.

By arguments similar to Lemma 2, we can show that the allocation (together with associated prices) is always on MBB, and hence is fPO. Further the algorithm only terminates if the allocation is EQ1 (Line 7). Similar to Lemma 5, we can show the following:

Lemma 8. *Let t_0 be a time-step where agent i ceases to be an LU agent, and let t_ℓ be the first subsequent time step just after which i becomes the LU agent again. Then:*

$$v_i(\mathbf{x}_i^{t_\ell+1}) > v_i(\mathbf{x}_i^{t_0}).$$

Using the above, as argued in Lemma 6, we can show:

Lemma 9. *Algorithm 2 terminates in time $\operatorname{poly}(n, m, U)$.*

We conclude:

Theorem 4. *Let $I = (N, M, V)$ be a positive-valued fair division instance. Then an allocation that is both EQ1 and fPO can be computed in time $\operatorname{poly}(n, m, U)$.*

As argued before, we have $U \leq mv_{max} + 1$. This gives:

Theorem 5. *Given a fair division instance $I = (N, M, V)$, an allocation that is EQ1 and fPO can be computed in time $\operatorname{poly}(n, m, v_{max})$, where $v_{max} = \max_{i,j} v_{ij}$. In particular, when $v_{max} \leq \operatorname{poly}(n, m)$, an EQ1+fPO allocation can be computed in $\operatorname{poly}(n, m)$ time.*

Finally using Lemma 7, Theorem 4 becomes:

Theorem 6. *Given a k -ary fair division instance $I = (N, M, V)$ where k is a constant, an allocation that is EQ1 and fPO can be computed in time $\operatorname{poly}(n, m)$.*

Remark 2. *We remark that our techniques can be used to show that EQ1+fPO allocations of chores can be computed in pseudo-polynomial time, and in polynomial-time for k -ary instances with constant k .*

Finding EF1+PO Allocations for Constant n

Our algorithm relies on the fact that there *exists* an integral allocation that is EF1+fPO, even when the values are rational numbers and not integers, since valuations are scale-free.

We call an instance *non-degenerate* if there are no multiplicative relationships between the v_{ij} 's. For such an instance, the MBB graph of market equilibrium is *acyclic* (Orlin 2010). In particular, this means that in any fractional allocation respecting the MBB graph, at most $n - 1$ goods are shared between agents, and remaining goods are allocated integrally. Since each shared good is shared between at most n agents, there are at most $(n - 1)^n = O(1)$ many integral allocations, obtained by rounding all possible fractional allocations that respect a given MBB graph. Note that an MBB graph is uniquely defined by the prices of the goods.

Thus, for a non-degenerate instance, our algorithm first enumerates all possible price vectors that correspond to market equilibria. As shown by Devanur and Kannan (2008), it is possible via cell-enumeration to enumerate all possible equilibrium price vectors in time $\operatorname{poly}(m)$ when n is constant. Then, for each such price vector, since we know the MBB graph is acyclic, we can enumerate all possible integral equilibria corresponding to it. Finally, we check if the allocation is EF1 in $\operatorname{poly}(m)$ time. Since we know that there exists some EF1+fPO allocation, and since we enumerate all possible price vectors corresponding to fPO allocations, and then compute all possible integral allocations at those prices, our algorithm is guaranteed to find one such allocation.

We now show how to adapt our algorithm for all instances, not just non-degenerate ones. Given a fair division instance $I = (N, M, V)$, we construct a non-degenerate instance $I' = (N, M, V')$, where V' is defined as follows:

$$v'_{ij} = \begin{cases} v_{ij} + \delta_{ij}, & \text{if } v_{ij} > 0 \\ 0, & \text{if } v_{ij} = 0 \end{cases},$$

where for all i, j , δ_{ij} is a small rational perturbation that ensures that the instance I' is non-degenerate. For instance, (Orlin 2010) set $\delta_{ij} = \epsilon^{in} + \epsilon^j$, where ϵ is arbitrarily close to 0. Let $\delta = \max_{i,j} \delta_{ij}$. We can ensure that $\delta < \frac{1}{2mv_{max}}$.

We now run our algorithm on the non-degenerate instance I' and we are guaranteed to find a market equilibrium (\mathbf{x}, \mathbf{p}) that is EF1 and fPO. Since the prices can be scaled, we can ensure for all $j \in M$, $1 \leq p_j \leq p_{max}$, where $p_{max} = \frac{1}{2mv_{max}\delta}$. We can then show that the allocation \mathbf{x} which is EF1+fPO for I' , is EF1+PO for I .

Lemma 10. *If \mathbf{x} is EF1 for I' , then \mathbf{x} is EF1 for I .*

Lemma 11. *If \mathbf{x} is fPO for I' , then \mathbf{x} is PO for I .*

Thus we have shown:

Theorem 7. *Given a fair division instance $I = (N, M, V)$, where n is constant, an EF1+PO allocation can be computed in $\text{poly}(m)$ time.*

Finally, we note that the techniques in Lemma 10 also extend easily to EQ1. Since we showed that an EQ1+fPO allocation is guaranteed to exist for positive instances:

Theorem 8. *Given a positive fair division instance $I = (N, M, V)$, where n is constant, an EQ1+PO allocation can be computed in $\text{poly}(m)$ time.*

PLS and Finding an EF1+PO Allocation

We closely follow the algorithm \mathcal{A} of Barman et al. (2018a) that finds an EF1+PO allocation in time $\text{poly}(n, m, v_{max})$, where $v_{max} = \max_{i,j} v_{ij}$, and show that it has the structure of a local search problem.

We first define some terms relevant to \mathcal{A} . For a small perturbation $\epsilon > 0$, a market outcome (\mathbf{x}, \mathbf{p}) is said to be ϵ -price envy-free up to one good (ϵ -pEF1) if for all agents i, h there is a good $j \in \mathbf{x}_h$ such that $(1 + \epsilon)\mathbf{p}(\mathbf{x}_i) \geq \mathbf{p}(\mathbf{x}_h \setminus \{j\})$. We can show like in Lemma 1, that if (\mathbf{x}, \mathbf{p}) is ϵ -pEF1, then \mathbf{x} is ϵ -EF1. Further, for sufficiently small ϵ , if \mathbf{x} is ϵ -EF1, then \mathbf{x} is also EF1 by using that v_{ij} 's are integers. The algorithm \mathcal{A} first perturbs all valuations to powers of $1 + \epsilon$, and then proceeds similarly to Algorithm 1, except that it transfers goods along alternating paths when the ϵ -pEF1 condition is violated. Further, prices are always maintained as powers of $1 + \epsilon$, bounded by p_{max} and a single agent is chosen as the LS, with ties broken lexicographically. The algorithm and its analysis is further described in the full version of the paper.

By the above arguments, computing an EF1+PO allocation is equivalent to computing an ϵ -pEF1+fPO market outcome (\mathbf{x}, \mathbf{p}) for the corresponding ϵ -perturbed instance, for small enough ϵ . We now reduce the EF1+PO problem to a PLS problem. First we describe its solution space, cost function, and the neighborhood structure.

Solution space. For an allocation \mathbf{x} and a vector of prices $\mathbf{p} = (p_1, \dots, p_m)$, call a configuration (\mathbf{x}, \mathbf{p}) *valid* if (i) \mathbf{x} is integral (ii) (\mathbf{x}, \mathbf{p}) is on MBB (iii) all prices p_j are of the form $(1 + \epsilon)^{q_j}$, where $q_j \in \mathbb{Z}$ is between 0 and $\log_{1+\epsilon} p_{max}$. Let the solution space be given by $S = \{(\mathbf{x}, \mathbf{p}) | (\mathbf{x}, \mathbf{p}) \text{ is valid}\}$. Since allocations are integral and prices are bounded integral powers of $(1 + \epsilon)$, S is bounded.

Cost function. Let $\delta(\mathbf{x}, \mathbf{p}) = 1$ if a valid allocation (\mathbf{x}, \mathbf{p}) is ϵ -pEF1, else 0. The cost function is a lexicographic function given by $\text{cost}(\mathbf{x}, \mathbf{p}) = (\delta(\mathbf{x}, \mathbf{p}), \min_{i \in N} \mathbf{p}(\mathbf{x}_i))$ if $(\mathbf{x}, \mathbf{p}) \in S$, and equal to $(-1, -1)$ if $(\mathbf{x}, \mathbf{p}) \notin S$.

Neighborhood structure. The structure is described by a polynomial time algorithm D . Each configuration (\mathbf{x}, \mathbf{p}) has a single neighbor. If $(\mathbf{x}, \mathbf{p}) \notin S$, then its neighbor $D(\mathbf{x}, \mathbf{p})$ is $(\mathbf{x}_0, \mathbf{p}_0)$. If $(\mathbf{x}, \mathbf{p}) \in S$, then $D(\mathbf{x}, \mathbf{p}) = (\mathbf{x}', \mathbf{p}')$, which is the allocation obtained by running \mathcal{A} starting with the allocation (\mathbf{x}, \mathbf{p}) until the spending of the least spender strictly increases. We show that:

Lemma 12. *Algorithm D terminates in polynomial time.*

Membership in PLS. We need to show the existence of three polynomial time algorithms: A – which outputs a solution $(\mathbf{x}_0, \mathbf{p}_0)$; B – which on input (\mathbf{x}, \mathbf{p}) computes the

$\text{cost}(\mathbf{x}, \mathbf{p})$; and C – which on input (\mathbf{x}, \mathbf{p}) computes a neighbor which has a strictly larger cost.

Algorithms A and B are trivial and in polynomial time. Observe that each solution (\mathbf{x}, \mathbf{p}) has only one neighbor $(\mathbf{x}', \mathbf{p}')$, and that it has a strictly larger cost since spending of the least spender at $(\mathbf{x}', \mathbf{p}')$ is strictly more than the spending of the least spender at (\mathbf{x}, \mathbf{p}) , or the latter is EF1. Thus algorithm D itself is algorithm C . Finally note that any local maxima of (S, cost, D) is an integral market allocation (\mathbf{x}, \mathbf{p}) where $\delta(\mathbf{x}, \mathbf{p}) = 1$, i.e., it is ϵ -pEF1, and thus EF1, even for the original valuations. Similarly, as is argued for the analysis of \mathcal{A} , the allocation is also PO for the original valuations. Therefore, computing an EF1+PO allocation for integral, additive valuations polynomial-time reduces to a PLS problem. However since it is not in TFNP \supseteq PLS, this does not show membership in PLS.

k -ary Instances with Constant n and k

We now consider k -ary fair division instances (N, M, V) where both k and n , the number of agents is constant.

Let \mathcal{X} be the set of all allocations for the instance I . For each agent $i \in N$, let $T_i = \{v_i(\mathbf{x}_i) : \mathbf{x} \in \mathcal{X}\}$, the set of different utility values i can get from any allocation. Let $U = \max_{i \in N} |T_i|$. From Lemma 7, we know U is at most $\text{poly}(m)$. Define $T = T_1 \times \dots \times T_n$. We note that $|T| \leq (\text{poly}(m))^n = \text{poly}(m)$, since n is constant, and can be computed in $\text{poly}(m)$ -time.

To solve certain fair division problems for such instances, we enumerate over each entry (u_1, \dots, u_n) of T , and check if there is a feasible allocation \mathbf{x} in which each agent i gets utility exactly u_i . The next Lemma shows that the latter can be done efficiently.

Lemma 13. *Given a vector $(u_1, \dots, u_n) \in T$, it can be checked in $\text{poly}(m)$ -time whether there is a feasible allocation \mathbf{x} s.t. for all agents i , $v_i(\mathbf{x}_i) = u_i$.*

By iterating through T we can in $\text{poly}(m)$ -time prepare a list of feasible utility vectors (and corresponding allocations) that satisfy our fairness and efficiency criteria.

Theorem 9. *For k -ary instance $I = (N, M, V)$ where both k and the number of agents n are constants, we can compute in $\text{poly}(m)$ time (i) an MNW allocation, (ii) a leximin optimal allocation, (iii) a \mathcal{F} +fPO allocation (when it exists) where \mathcal{F} is any polynomial-time checkable property.*

Discussion

In this paper, we showed that an EF1+fPO allocation can be computed in pseudo-polynomial time, thus improving upon the result of Barman et al (2018a). Our work also implies polynomial time algorithms for two special cases: (i) computing EF1+fPO allocation for k -ary instances where k is a constant, (ii) computing EF1+PO allocation when n (number of agents) is constant. Settling the complexity of the problem for general k and n remains a challenging open problem. Our results extend to the fairness notions of EQ1 and wEF1 as well. We showed that computing an EF1+PO allocation reduces to a problem in the complexity class PLS. Showing the existence of EF1+PO allocations for chores is another interesting research direction.

Acknowledgements

Work on this paper is supported by NSF Grant CCF-1942321 (CAREER).

References

- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2020. Maximum Nash Welfare and Other Stories About EFX. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, (IJCAI)*, 24–30.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018a. Finding Fair and Efficient Allocations. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, 557–574.
- Barman, S.; Krishnamurthy, S. K.; and Vaish, R. 2018b. Greedy Algorithms for Maximizing Nash Social Welfare. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 7–13.
- Bliem, B.; Brederick, R.; and Niedermeier, R. 2016. Complexity of Efficient and Envy-Free Resource Allocation: Few Agents, Resources, or Utility Levels. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, 102–108. AAAI Press. ISBN 9781577357704.
- Brams, S. J.; and Taylor, A. D. 1996. Fair division - from cake-cutting to dispute resolution.
- Brederick, R.; Kaczmarczyk, A.; Knop, D.; and Niedermeier, R. 2019. High-Multiplicity Fair Allocation: Lenstra Empowered by N-Fold Integer Programming. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC '19*, 505–523. New York, NY, USA: Association for Computing Machinery. ISBN 9781450367929.
- Budish, E. 2011. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy* 119(6): 1061 – 1103.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC)*, 305–322.
- Chakraborty, M.; Igarashi, A.; Suksompong, W.; and Zick, Y. 2020. Weighted Envy-Freeness in Indivisible Item Allocation. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 231–239.
- de Keijzer, B.; Bouveret, S.; Klos, T.; and Zhang, Y. 2009. On the Complexity of Efficiency and Envy-Freeness in Fair Division of Indivisible Goods with Additive Preferences. In *Algorithmic Decision Theory*, 98–110.
- Devanur, N. R.; and Kannan, R. 2008. Market Equilibria in Polynomial Time for Fixed Number of Goods or Agents. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 45–53.
- Foley, D. 1967. Resource Allocation and the Public Sector. *Yale Economic Essays* 7(1): 45–98.
- Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable Allocations of Indivisible Goods. In Kraus, S., ed., *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2020. Equitable Allocations of Indivisible Chores. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 384–392.
- Lee, E. 2015. APX-Hardness of Maximizing Nash Social Welfare with Indivisible Items. *Information Processing Letters* 122.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, 125–131.
- Mas-Colell, A.; Mas-Colell, P.; D, W.; Whinston, M.; Green, J.; Hara, C.; Green, P.; Segal, I.; Press, O. U.; and Tadelis, S. 1995. *Microeconomic Theory*. Oxford University Press.
- Moulin, H. 2003. Fair Division and Collective Welfare.
- Orlin, J. B. 2010. Improved Algorithms for Computing Fisher’s Market Clearing Prices: Computing Fisher’s Market Clearing Prices. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing (STOC)*, 291–300.
- Othman, A.; Sandholm, T.; and Budish, E. 2010. Finding Approximate Competitive Equilibria: Efficient and Fair Course Allocation. In *Proceedings of the 9th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 873–880.
- Plaut, B.; and Roughgarden, T. 2018. Almost Envy-Freeness with General Valuations. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2584–2603.
- Steinhaus. 1948. The Problem of Fair Division. *Econometrica* 16(1): 33–111.
- Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9(1): 63 – 91.