Bicriteria Fair Allocation

Martin Loebl[1], David Sychrovský[1], Cosimo Vinci[2]

[1]Charles University in Prague, Czechia
[2]University of Salento, Lecce, Italy
{loebl,sychrovsky}@kam.mff.cuni.cz, cosimo.vinci@unisalento.it

Abstract

Fair allocation of goods to agents has been extensively studied because of their applications in several socio-economic contexts, from division of inheritance to recommendation systems. In a common setting, we have n agents and m items, and each agent has an individual valuation for each of the goods. However, in many situations agents may have more than one valuation - for example when a recommendation platform must allocate ad slots in order to satisfy both visibility and marketing goals.

To deal with this and other general scenarios, in this paper we study a novel *bicriteria fair allocation* framework: a generalization of standard fair allocation settings where each agent has a common public valuation and an individual valuation for each good. The goal is to find an allocation that, for two integers $\gamma, \delta \geq 0$, is *envy-free-up-to-publicly-γ-and-privately-δ-goods* (EF-(γ, δ)): each agent becomes non-envious of each other w.r.t. the public (resp. her private) valuation, after deleting at most γ (resp. δ) goods from the bundle of each other.

We first provide a polynomial-time algorithm that is EF-(1,1) when private valuations are known to the system. Then, we focus on the realistic case in which agents can misreport their private valuations to the system, and we provide a randomized polynomial-time algorithm that returns EF-$(1, \delta)$ allocations with high probability, where $\delta = O(\alpha\sqrt{\log(n)m/n})$ and α is the maximum private valuation for any item.

Introduction

In recent years, the concept of fair allocation has gained significant attention in both theoretical and applied realms of economics and computer science, with the aim of addressing several real-life problems connected with fair goods distribution to people, e.g., division of inheritance, house allocation and dispute resolution and recommendation systems. The central idea is to distribute (or allocate) indivisible goods (or items, or resources) to agents in such a way that no individual feels envious of the resource bundle, that is, to guarantee the *envy-freeness* (EF) property (Gamow and Stern 1958). This ensures a certain level of fairness and satisfaction among participants, which, in turn, can enhance cooperation and system efficiency.

Differently from the case of divisible resources (Steinhaus 1948; Stromquist 1980), an outcome in which no agent is envious of each other does not always exist. Thus, several relaxation of envy-free allocations have been proposed, such as *envy-free-up-to-one-good* (EF1) (Budish 2011) and *envy-freeness-up-to-any-good* (EFX) (Caragiannis et al. 2019) allocations. In particular, an allocation is EF1 (resp. EFX) if no agent is envious after deleting opportunely a good (resp. any good) from each others' bundle. Several works have shown existence of EF1 allocations, by providing polynomial-time algorithms to compute them (e.g., the round-robin algorithm for additive valuations (Caragiannis et al. 2019) and the envy-cycle-elimination algorithm for more general monotone valuations (Lipton et al. 2004)). Instead, showing the existence of EFX allocations is a major open problem in fair allocation, and its existence has been addressed for restricted setting only (Amanatidis et al. 2021).

Traditional envy-free allocation systems have primarily considered a single valuation per-agent for each object being distributed. However, in many real-world scenarios, agents have complex perspectives towards goods that cannot be captured by a single valuation per-agent. This is the case of recommendation systems, where fairness constitutes a desirable aspect when the users and/or the interests of the platform must be satisfied under multiple criteria, without creating too many disparities between the agents involved (Li et al. 2023; Wang et al. 2023). As an example, assume that some companies (i.e., the users) vying for ad slots on a popular platform. The ad slots can have an objective value based on factors like visibility and reach (independently on the assignment), while companies might also have a personal valuation based on their specific marketing goals or target demographics. Allocating these slots in an envy-free manner, respecting both these valuations, can lead to enhanced advertiser satisfaction and platform trust. It can also pave the way for more efficient advertising ecosystems that align with both platform goals and individual advertiser objectives.

Our Contribution

To model the above scenarios (and many others), we introduce an ad-hoc fair allocation framework, called *bicriteria fair allocation*, where a set M of goods (e.g., the ad slots) must be allocated to some agents (e.g., the companies) in such a way that the resulting allocation is as fair as possible, according to two distinct additive valuations: (i) a *public (or common) valuation* w (e.g., modeling ad slot visibility and reach), that is common to all agents and is publicly known (e.g., if the recommendation system can well-estimate it according to its data), and (ii) a *private valuation* v_i for each agent i (e.g., modeling specific marketing goals of the companies) that may also be unknown or misreported (e.g., if some company tries to manipulate the choices of the recommendation system for their own interests).

In the same spirit of EF1 allocations, our goal is to compute allocations of goods to agents that are approximately fair according to both valuations. In particular, we aim at finding *envy-free-allocation-up-to-publicly-γ-and-privately-δ-goods* (EF-(γ, δ)), meaning that each agent is not envious with respect to the public (resp. private) valuation after deleting at most γ (resp. δ) goods from each others' bundle.

Assuming that the private valuations are known to the system, we provide a polynomial-time algorithm that returns an EF-(1,1) allocation (see Theorem 1). Such algorithm elegantly combines the round-robin (Caragiannis et al. 2019) and the envy-cycle-elimination (Lipton et al. 2004) approaches, and its fairness guarantee is the best possible, considering that, even for a single valuation, an agent might delete at least one good from the bundle of someone else to recover fairness.

If we assume that the private valuations are initially unknown and are reported by the agents to the system (i.e, the recommendation platform), the high fairness considered above is guaranteed to each player that truthfully declares her goods valuations (that is, if an agent reports truthfully, she will be non-envious up to one good w.r.t. both valuations). However, it could be the case that some agents strategically misreport their valuations to influence and manipulate the goods allocation, in order to maximize her public and/or private valuation, but possibly, at the expense of the system performance (e.g., measured in terms of overall satisfaction with the recommendation). Such possibility of strategising may negatively influence the trust of agents. To address the above issue, we provide a randomized algorithm that is *strategy-proof* (i.e., agents have no incentive to disclose or misreport their private valuations) and returns an EF-$(1, \delta)$ allocation with high probability, where $\delta = O\left(\alpha\sqrt{\log(n)m/n}\right)$ and α is the maximum goods value. The considered algorithm combines the round-robin approach with further randomness features to recover strategy-proofness, at the expense of the constant fairness guarantee on the private valuation achieved by the first proposed algorithm. Anyway, despite the obtained bound on the fairness under private valuations is non-constant, it is still sublinear in m/n and is quite small under several cases (e.g., it is logarithmic if the maximum item value α and the ratio m/n are kept constant); furthermore, the envy-freeness-up-to-one-good property is still preserved under public valuations. As a result of independent interest, we also have that the fairness guarantee of our randomized algorithm outperforms that achieved by the existing strategy-proof approaches (Caragiannis et al. 2009), even for the case of a single valuation per-agent.

Further Related Work

Further connections of our model and results with the existing works are provided below.

Group-fair allocation. Our bicriteria fair allocation framework can be seen as a particular case of a general goods allocation framework, known as *group-fair allocation* (Kyropoulou, Suksompong, and Voudouris 2020), that generalizes the standard fair allocation setting as follows: r players[1] are organized in several fixed groups, and we must distribute goods to groups in such a way that the resulting allocations are as fair as possible for each player in the group. Several works (e.g., (Segal-Halevi and Suksompong 2019; Kyropoulou, Suksompong, and Voudouris 2020; Manurangsi and Suksompong 2022)) have studied existence and computation of allocations which are envy-free-up-to-δ-goods (EF-δ) under known valuations, for some integer δ. Among these, the work of (Manurangsi and Suksompong 2022) applies some tools from Discrepancy theory (Alon and Spencer 2008) to show asymptotically tight bounds on the lowest integer δ guaranteeing the existence of EF-δ allocations. In particular, they show that EF-$O(\sqrt{r})$ always exist and can be computed in polynomial time.

We observe that our model of bicriteria fair allocation can be instantiated in the group-fair allocation model. Indeed, each agent of our model can be represented as a distinct group, and each valuation of our model can be represented as a player having the same valuation, so that each group has two distinct players, that is, $r = 2n$. Thus, by applying the results of (Manurangsi and Suksompong 2022) to our model, we can compute an allocation where each agent can recover fairness by removing at most $O(\sqrt{r}) = O(\sqrt{n})$ items, that is worse than the fairness guarantee of envy-freeness-up-to-one-good that our algorithm for known valuations achieved. We conclude that, despite the existing results on group-fair allocation (e.g., that of (Manurangsi and Suksompong 2022)) can be directly applied to our framework, the resulting fairness guarantee is not as good as that provided by our algorithm. This is not surprising, considering that our algorithm has been developed to work with a proper sub-case of group-fair allocation problems, where each group has exactly two valuations and one of them is common.

Strategy-proofness. Several works addressed strategy-proofness issues for the allocation of indivisible goods (e.g., (Lipton et al. 2004; Caragiannis et al. 2009; Amanatidis et al. 2017; Bouveret and Lang 2011; Padala and Gujar 2022; Halpern et al. 2020; Arbiv and Aumann 2022; Arnosti and Bonet 2022)). In particular, (Lipton et al. 2004) provides a randomized strategy-proof algorithm for standard fair allocation with a single valuation per-agent, which guarantees a sub-linear maximum envy with high probability. The analysis of such algorithm has been subsequently improved by

[1]We use term "player" to avoid ambiguity with the agents considered in our bicriteria fair allocation framework.

(Caragiannis et al. 2009), who showed that the maximum envy among agents is at most $\alpha\sqrt{\log(n)m}$ with high probability, where α denotes the maximum item value; such bound is proved by the $O(\alpha\sqrt{\log(n)m/n})$ bound achieved by our randomized strategy-proof algorithm. (Amanatidis et al. 2017) provided some useful characterization of truthful mechanisms for fair allocation, and showed that there is no strategy-proof algorithm that returns EF1 allocations (so as for other fairness notions). In light of this result, our EF-$(1,1)$ algorithm cannot be strategy-proof.

(Arbiv and Aumann 2022; Arnosti and Bonet 2022) studied strategy-proofness in the context of group-fair allocations, but with partially different objectives than ours.

Model and Definitions

Let $[k] := \{1,\dots,k\}$ denote the set of the first k integers.

Let $N := \{1,\dots,n\}$ denote a set of n *agents* and M a set of m *goods*. Each agent $i \in N$ has a *private* valuation v_i that assigns a value $v_i(g) \geq 0$ to each good $g \in M$, and all agents have a *public* (or *common*) valuation w assigning a value $w(g) \geq 0$, that is common to all agents. Valuations can be extended in an additive way to *bundles* $S \subseteq M$ of goods, that is, $v_i(S) = \sum_{g\in S} v_i(g)$ for any $i \in N$ and $w(S) = \sum_{g\in S} w(g)$. The tuple $I = (N, M, w, (v_i)_{i\in N})$ denotes a generic input instance of *bicriteria fair allocation*.

A *full allocation* $A = (A_1,\dots,A_n)$ (simply denoted as *allocation*) is a partition of M in n disjoint bundles of goods, so that A_i is the bundle assigned to agent $i \in N$. A *partial allocation* is a general collection $A = (A_1,\dots,A_n)$ of disjoint subsets of M, so that A_i is still the bundle assigned to agent $i \in N$, but some goods are unallocated.

Bicriteria Fairness. Given two integers $\gamma, \delta \geq 0$, a (full or partial) allocation $A = (A_1,\dots,A_n)$ is called *bicriteria-envy-free-up-to-γ-and-δ-goods* (BEF-(γ,δ)) if, for any agent $i, j \in N$, there exist subsets $S_j, Q_{i,j} \subseteq A_j$ with $|S_j| \leq \gamma$ and $|Q_{i,j}| \leq \delta$ such that $w(A_i) \geq w(A_j \setminus S_j)$ and $v_i(A_i) \geq v_i(A_j \setminus Q_{i,j})$ (where $Q_{i,j}$ may depend on the choice of both i and j, while S_j depends on j only). Informally, an allocation is BEF-(γ,δ) if no agent i is envious of the bundle assigned to each other agent j with respect to her private (resp. public) valuation, up to at most γ (resp. δ) goods.

In the remainder of the paper, we will assume that the public valuation is known to the decision maker who will be responsible in assigning goods to agents, while private valuations, depending on the cases, may be either reported truthfully, or misreported/disclosed.

The Case of Known Valuations

In this section, we assume that agents truthfully report their private valuations to the designer and we show that, in such case, BEF-$(1,1)$ allocations always exist and can be computed in polynomial time, thus matching the best possible bound achieved for standard fair allocation problems with a single valuation per agent (Lipton et al. 2004). To show this result, we propose a novel algorithm called *Round-robin+Envy-Cycle-elimination* (REC), that elegantly combines two known algorithmic approaches, namely the round-robin (Caragiannis et al. 2019) and the envy-cycle-elimination (Lipton et al. 2004) algorithms, already used to find almost fair allocations in standard goods allocation problems.

Before describing the algorithm, we first give some preliminary definitions.

Envy-cycle-elimination procedure. Given a partial allocation $A = (A_1,\dots,A_n)$, the *envy-graph* of A is a graph $G_A = (V, E)$ where nodes are the agents in N, and there exists an edge $(i,j) \in E$ iff $v_i(A_i) < v_i(A_j)$, i.e., if agent i is envious of the bundle assigned to agent j, w.r.t. the private valuation. (Lipton et al. 2004) defined a procedure, called *envy-cycle-elimination* that, given a partial allocation A, rearranges the bundles of A in such a way that the envy-graph of the obtained partial allocation becomes a direct-acyclic-graph (DAG). The envy-cycle-elimination procedure works as follows: (i) if there is a cycle $A_{i_1},\dots,A_{i_r} = A_{i_1}$ in G, then A_{i_1} becomes the bundle assigned to i_0, A_{i_2} becomes the bundle assigned to i_1, and in general, A_{i_s} becomes the bundle assigned to i_{s-1} for any $s \in [r]$; (ii) we iterate the above procedure until G becomes a DAG. We observe that, after each application of (i), the cardinality $|E|$ of the edges decreases by at least 1, thus, as the number of edges in the initial envy-graph is at most $O(n^2)$, after $O(n^2)$ steps the envy graph necessarily-becomes a DAG.

REC algorithm. The REC algorithm works in $T := \lceil m/n \rceil$ rounds, and is defined as follows:

1. We first reorder all goods in non-increasing order w.r.t. the public valuation, that is, $w(g_1) \geq w(g_2) \geq \ldots \geq w(g_m)$, where $g_1,\dots,g_m$ is the sequence of the goods in M according to the new order. Then, the algorithm splits all goods into T disjoint subsets $M_1,\dots,M_T$, where $M_t = \{g_{(t-1)n+1},\dots,g_{tn}\}$ for any $t \in [T-1]$, and $M_T := \{g_{(T-1)n+1},\dots,g_m\}$.
2. Let $A := (A_1,\dots,A_n)$ be the partial allocation initially made of n empty sets. At round $t = 1$, for each $i = 1,\dots,n$, agent i picks the good $g^* \in M_1$ that maximizes the private valuation $v_i(g^*)$ among the goods in M_1, removes g^* from M_1 and includes it in her bundle A_i. We observe that, after this first round, the envy-graph G_A associated with A is already a DAG without applying envy-cycle-elimination (indeed, there are no edges (i,j) with $i < j$ because of the greedy choice of the agents, and a graph with such a property is a DAG).
3. At each subsequent round $t = 2,\dots,T$, we do the following sub-steps:
   (a) As G_A is a DAG, we can compute a topological ordering $\sigma_1 \prec \sigma_2 \prec \ldots \prec \sigma_n$ of G_A, that is, a total ordering of nodes/agents such that, if there is an edge from agent σ_i to agent σ_j in G_A, then $\sigma_i \prec \sigma_j$.
   (b) We assign goods of M_t to agents as in step 2, but following the above topological ordering, that is, for each $i = 1,\dots,n$, agent σ_i picks the good $g^* \in M_t$ that maximizes the private valuation $v_{\sigma_i}(g^*)$ among the goods in M_t, removes g^* from M_t and includes it

in her bundle A_{σ_i}; if $t = T$, M_t could become empty before all agents pick a good, thus such agents will not receive their T-th good.

(c) We apply the envy-cycle-elimination procedure to A, so that the envy-graph of the new partial allocation is again a DAG.

4. We repeat step 3 for all T rounds, that is, until all goods are allocated; finally, we return the resulting (full) allocation A.

See the supplementary material for the pseudo-code of the REC algorithm.

Theorem 1. *For any input instance I of bicriteria fair allocation, the* REC *algorithm returns in polynomial time an EF-$(1,1)$ allocation.*

To show that REC returns an EF-(1,1) allocation, we first show that each agent is envy-free-up-to-one-good w.r.t. the public valuation, and this is done by exploiting the definition of groups $M_1, \ldots, M_T$. Then, we show that each agent envy-free-up-to-one-good w.r.t. her private valuation, by exploiting the topological ordering used to assign items at each round, and the subsequent envy-cycle elimination procedure. The full proof on the fairness guarantee of REC and its polynomial-time complexity is deferred to the supplementary material.

Strategy-proofness via Randomization

In this section, we assume that agents can misreport their private valuations with the aim of manipulating the assignment and reaching a certain individual goal (e.g., improving their public or private valuation). However, misreporting can affect public or private values of other agents in the final outcome, possibly creating mistrust among agents. To overcome the problem of misreporting, we design and analyze a randomized variant of the round-robin algorithm, called *probabilistic-round-robin* (PRR), that only uses the knowledge of the public valuation, without using any information on the private valuations, and returns an EF-$(1,\delta)$ allocation with high probability, where $\delta = O(\alpha\sqrt{\log(n)m/n})$ and $\alpha \geq 0$ is the maximum goods valuation.

In the remainder of the section, we will also assume w.l.o.g. that the minimum non-zero public or private valuation is 1.[2]

PRR Algorithm. The PRR algorithm works as follows:

1. As in REC, we reorder all goods in non-increasing order w.r.t. the public valuation, that is, $w(g_1) \geq w(g_2) \geq \ldots \geq v_i(g_m)$, where $g_1, \ldots, g_m$ is the sequence of the goods in M after the ordering. Then, the algorithm splits all goods into T groups $M_1, \ldots, M_T$, where $M_t = \{g_{(t-1)n+1}, \ldots, g_{tn}\}$ for any $t \in [T-1]$, and $M_T := \{g_{(T-1)n+1}, \ldots, g_m\}$. Let $A := (A_1, \ldots, A_n)$ be the partial allocation initially made of n empty sets.
2. At each round $t \in [T]$, we randomly choose an ordering $\sigma_1 \prec \sigma_2 \ldots \prec \sigma_n$ of all agents. Then, for each $i = 1, \ldots, n$, good $g_{(t-1)n+i}$ is assigned to agent σ_i and is removed from M_t.
3. After executing step 2 over all T rounds, we return the resulting (full) allocation A.

PRR could be interpreted as a multiple-round variant of the *random serial dictatorship algorithm* (RSD) (Bogomolnaia and Moulin 2001), a strategy-proof algorithm applied to match n houses to n agents, where the agents' ordering is picked uniformly at random, and each agents chooses the best house for her. PRR randomly picks a distinct ordering at each round and, differently from RSD, does not assign items based on the agents' choices, but following a fixed items ordering. See the supplementary material for the pseudo-code of the PRR algorithm.

Theorem 2. *For any input instance $I = (N, M, w, (v_i)_{i \in N})$ of bicriteria fair allocation with n agents and m goods, and any constant $\beta > 0$, the* PRR *algorithm is strategy-proof and returns in polynomial time an EF-$(1,\delta)$ allocation with probability at least $1 - 1/n^\beta$, where $\delta = \left\lceil \alpha\sqrt{2(\beta+2)\log(n)\left\lceil \frac{m}{n}\right\rceil} \right\rceil = O\left(\alpha\sqrt{\log(n)\frac{m}{n}}\right)$ and $\alpha = \max_{i \in N, g \in M} v_i(g)$.*

The proof of Theorem 2 is based on the probabilistic analysis via the Hoeffding's inequality (Hoeffding 1963), and it is deferred to the supplementary material.

Future Works

Our work left several interesting research directions. First of all, it would be interesting to study more general agents valuations, e.g., submodular. In particular, it would be nice to provide new analysis for our algorithms under more general valuations, or to design new algorithms to deal with them.

Furthermore, it would be good to study more general multicriteria fair allocation settings than that studied in this work (e.g., with more than 2 public and/or private valuations per agent), and seeing if one can obtain better fairness guarantees than those already provided for the general framework of group-fair allocation (Kyropoulou, Suksompong, and Voudouris 2020). To this aim, a direction we are considering is that of applying a variant of the REC algorithm to a bicriteria fair allocation setting in which each agent i has two private valuations v_i^1 and v_i^2, where v_i^1 is binary, agent specific, and used in place of the public valuation w considered in this work. Relatively to this setting, we are exploiting interesting connections between the EF-(1,δ) guarantee (with δ as small as possible) and the *Gap* of the *value hypergraph*[3] $H = (V, E)$ associated with v^1 (Alon 1998), where V is the set of items and E contains, for each agent i, the set of items that i values positively under binary valuation v_i^1.

Finally, it would be nice to see how to get better strategy-proof algorithms than ours, e.g., if we assume weaker notions of strategy-proofness.

[2]Indeed, if it is not the case, it is sufficient to divide the goods valuation of each agent i by $\min\{g \in M : v_i(g) > 0\}$ to match this requirement and to preserve the ordering relation among valuations over bundles (i.e., we obtain an equivalent instance of bicriteria fair allocation).

[3]Given a hypergraph $H = (V, E)$, the Gap of H is defined as $Gap(H) = \min_{o \in L(V)} Gap(o)$, where $L(V)$ is the set of all linear orders of V, $Gap(o)$ is defined as $\max_{e \in E} Gap(o, e)$ and $Gap(o, e)$ is the number of times o leaves e in H.

References

Alon, N. 1998. Piercing d -Intervals. *Discret. Comput. Geom.*, 19(3): 333–334.

Alon, N.; and Spencer, J. H. 2008. *The Probabilistic Method, Third Edition*. Wiley-Interscience series in discrete mathematics and optimization. Wiley.

Amanatidis, G.; Birmpas, G.; Christodoulou, G.; and Markakis, E. 2017. Truthful Allocation Mechanisms Without Payments: Characterization and Implications on Fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC*, 545–562. ACM.

Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.*, 863: 69–85.

Arbiv, T.; and Aumann, Y. 2022. Fair and Truthful Giveaway Lotteries. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, 4785–4792.

Arnosti, N.; and Bonet, C. 2022. Lotteries for Shared Experiences. In *The 23rd ACM Conference on Economics and Computation, EC*, 1179–1180. ACM.

Bogomolnaia, A.; and Moulin, H. 2001. A New Solution to the Random Assignment Problem. *J. Econ. Theory*, 100(2): 295–328.

Bouveret, S.; and Lang, J. 2011. A General Elicitation-Free Protocol for Allocating Indivisible Goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence IJCAI*, 73–78.

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

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. On Low-Envy Truthful Allocations. In Rossi, F.; and Tsoukiàs, A., eds., *Algorithmic Decision Theory, First International Conference, ADT. Proceedings*.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Trans. Economics and Comput. (EC 2019)*, 7(3): 12:1–12:32.

Gamow, G.; and Stern, M. 1958. *Puzzle-math*. Viking Press.

Halpern, D.; Procaccia, A. D.; Psomas, A.; and Shah, N. 2020. Fair Division with Binary Valuations: One Rule to Rule Them All. In *Web and Internet Economics - 16th International Conference, WINE, Proceedings*, volume 12495 of *Lecture Notes in Computer Science*, 370–383.

Hoeffding, W. 1963. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301): 13–30.

Kyropoulou, M.; Suksompong, W.; and Voudouris, A. A. 2020. Almost envy-freeness in group resource allocation. *Theor. Comput. Sci.*, 841: 110–123.

Li, Y.; Chen, H.; Xu, S.; Ge, Y.; Tan, J.; Liu, S.; and Zhang, Y. 2023. Fairness in Recommendation: Foundations, Methods, and Applications. *ACM Trans. Intell. Syst. Technol.*, 14(5).

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *EC 2004*, 125–131.

Manurangsi, P.; and Suksompong, W. 2022. Almost envy-freeness for groups: Improved bounds via discrepancy theory. *Theor. Comput. Sci.*, 930: 179–195.

Padala, M.; and Gujar, S. 2022. Mechanism Design without Money for Fair Allocations. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, WI-IAT '21, 382–389.

Segal-Halevi, E.; and Suksompong, W. 2019. Democratic fair allocation of indivisible goods. *Artif. Intell.*, 277.

Steinhaus, H. 1948. The problem of fair division. 16: 101–104.

Stromquist, W. 1980. How to Cut a Cake Fairly. *The American Mathematical Monthly*, 87(8): 640–644.

Wang, Y.; Ma, W.; Zhang, M.; Liu, Y.; and Ma, S. 2023. A Survey on the Fairness of Recommender Systems. *ACM Trans. Inf. Syst.*, 41(3).

Supplementary Material

REC Algorithm

Algorithm 1: REC Algorithm

Require: An instance $I = (N, M, (v_i)_{i \in N}, w)$ of bicriteria fair allocation.
Ensure: An allocation $A = \mathsf{REC}(I)$.

Let $T := \lceil m/n \rceil$.
Compute an ordering $g_1, g_2, \ldots, g_m$ of all goods in such a way that $w(g_1) \geq w(g_2) \geq \ldots \geq w(g_m)$.
Let $M_t := \{g_{(t-1)n+1}, \ldots, g_{tn}\}$ for any $t \in [T-1]$, and $M_T := \{g_{(T-1)n+1}, \ldots, g_m\}$.
Let $A := (A_1, \ldots, A_n)$ be the empty allocation.
for $t = 1, \ldots, T$ **do**
  Let $\sigma_1 \prec \sigma_2 \prec \ldots \prec \sigma_n$ be a topological ordering of the envy-graph G_A of allocation A, that is, if (σ_i, σ_j) is an edge in G_A, then $\sigma_i \prec \sigma_j$.
  $i \leftarrow 1$.
  while $M_t \neq \emptyset$ **do**
    Let $g^* \in \arg\max_{g \in M_t} v_{\sigma_i}(g)$.
    $A_{\sigma_i} \leftarrow A_{\sigma_i} \cup \{g^*\}$.
    $M_t \leftarrow M_t \setminus \{g^*\}$.
    $i \leftarrow i + 1$.
  end while
  $A \leftarrow \mathsf{EnvyCycleElimination}(A)$.
end for
return $A := (A_1, \ldots, A_n)$.

Algorithm 2: EnvyCycleElimination

Require: An allocation $A = (A_1, \ldots, A_n)$.
Ensure: An allocation $\mathsf{EnvyCycleElimination}(A)$.

The envy-graph $G_A = (V, E)$ associated with A is defined in such a way that $V := [n]$, and there exists an edge $(i, j) \in E$ iff $v_i(A_i) < v_i(A_j)$.
while $\exists$ a cycle $A_{i_1}, \ldots, A_{i_r} = A_{i_1}$ in the envy-graph G_A **do**
  $A_{temp} \leftarrow A_{i_1}$.
  for $h = 1, \ldots, r-1$ **do**
    $A_{i_h} \leftarrow A_{i_{h+1}}$.
    Update A and G_A.
  end for
  $A_{i_r} \leftarrow A_{temp}$.
  Update A and G_A.
end while
return A

Proof of Theorem 1

Let $I = (N, M, w, (v_i)_{i \in N})$ denote the input instance of bicriteria fair allocation, with $|N| = n$ and $|M| = m$. We first show that REC can be executed in polynomial time.

Lemma 1. REC *can be executed in $O(m \log(m) + n^3 m)$ time.*

Proof of Lemma 1. REC first requires to order the goods according to the public valuation and to split them in $T := \lceil m/n \rceil$ disjoint subsets, and this can be done in $O(m \log(m))$ time. If $m \leq n$, we have that $T = 1$, thus the time complexity is uniquely determined by the previous ordering, and the claim follows. If $m > n$, we have $T > 1$, and for each round $t \in [T]$, the algorithm computes the topological ordering of the envy-graph $G_A = (V, E)$; this can be done in time $O(|V| + |E|) \leq O(n^2)$. Then, REC let each agent chooses in set M_t her best good according to the private valuation, and this can be done in $O(n^2)$ time. Finally, REC applies the envy-cycle elimination procedure at each round, and this can be done in $O(n^4)$ time. Indeed, each cycle detection can be performed in $O(|V| + |E|) \leq O(n^2)$ time, and each cycle-elimination is applied at most $O(n^2)$ times at each round, that is, the overall time complexity at each round is $O(n^4)$. We conclude that, the time-complexity of REC over all rounds is $O(m \log(m) + T(n^2 + n^4)) = O(m \log(m) + (m/n)(n^2 + n^4)) = O(m \log(m) + n^3 m)$. □

In the remainder of the proof, we only focus on the fairness guarantee achieved by REC. For the sake of simplicity, we assume w.l.o.g. that n divides m. Indeed, if it is not the case, it is sufficient to add some dummy goods having zero value that will be assigned at the last round only, to the agents who would not have receive any good in the initial instance. Such dummy goods having zero value will not affect the value of the bundles, so as the fairness guarantees.

Let $A^t = (A_1^t, \ldots, A_n^t)$ be the partial allocation obtained by REC at the end of round t, for any $t \in [T]$; furthermore, let $g_{i,r}^t$ denote the r-th good included in bundle A_i^t for any $t \in [T]$, $i \in N$, $r \in [t]$. In the following lemmas, we show that, under partial allocation A^t, each agent is not envious of each other up to one good, with respect to both valuations, and this will be sufficient to show the claim of the theorem.

Lemma 2. *For any $i, j \in N$ and $t \in [T]$, we have that $w(A_i^t) \geq w(A_j^t \setminus \{g_{j,1}^t\})$, that is, each agent is not envious of each other up to one good, with respect to the public valuation and under partial allocation A^t.*

Proof of Lemma 2. Let $i, j \in N$ and $t \in [T]$. REC assigns goods in such a way that agent i evaluates any good assigned at each round $r \in [t-1]$ at least as any good assigned at round $r+1$, under the public valuation (that is common to all agents). This implies that

$$w(\{g_{i,r}^t\}) \geq w(\{g_{j,r+1}^t\}), \quad \forall r \in [t-1]. \tag{1}$$

Then, we have that

$$\begin{aligned} w(A_i^t) &= \sum_{r=1}^{t} w(\{g_{i,r}^t\}) \\ &\geq \sum_{r=1}^{t-1} w(\{g_{i,r}^t\}) \\ &\geq \sum_{r=2}^{t} w(\{g_{j,r}^t\}) \qquad (2) \\ &= w(A_j^t \setminus \{g_{j,1}^t\}), \end{aligned}$$

where (2) follows from (1). □

Lemma 3. *For any $i, j \in N$ and $t \in [T]$, there exists $g \in A_j^t$ such that $v_i(A_i^t) \geq v_i(A_j^t \setminus \{g\})$, that is, each agent is not envious of each other up to one good with respect to the private valuation, under partial allocation A^t.*

Proof of Lemma 3. We show the lemma by induction on t. For $t = 1$, the lemma is obviously satisfied, as each bundle has one good only.

Assume that, for some $t \geq 2$, the claim is true for $t-1$, and let us show it for t. To obtain partial allocation A^t from A^{t-1}, the REC algorithm first assigns the most valuable good (among the remaining ones) to each agent, but following the topological ordering defined by the envy-graph $G_{A^{t-1}}$, and let $\tilde{A}^t = (\tilde{A}_1^t, \ldots, \tilde{A}_n^t)$ denote the partial allocation obtained after this procedure. We will show that no agent is envious up to one good, w.r.t. the private valuation and partial allocation $\tilde{A}^t$, that is, for any $i, j \in N$ there exists $g^* \in \tilde{A}_j^t$ such that $v_i(\tilde{A}_i^t) \geq v_i(\tilde{A}_j^t \setminus \{g^*\})$. Given $i, j \in N$, we have two possible cases:

Case $v_i(A_i^{t-1}) \geq v_i(A_j^{t-1})$**:** Let $g^* = \tilde{A}_j^t \setminus A_j^{t-1}$ be the good assigned to agent j at round t. Then, we have that

$$v_i(\tilde{A}_i^t) \geq v_i^2(A_i^{t-1}) \geq v_i(A_j^{t-1}) = v_i(\tilde{A}_j^t \setminus \{g^*\}). \quad (3)$$

Case $v_i(A_i^{t-1}) < v_i(A_j^{t-1})$**:** By the considered topological ordering, at round t agent i will choose her most valuable good before agent j. Thus, denoting as g_i^* and g_j^* the goods assigned to i and j, respectively, we have that

$$v_i(\{g_i^*\}) \geq v_i(\{g_j^*\}). \quad (4)$$

Furthermore, by the inductive hypothesis, we have that there exists $g^* \in A_j^{t-1}$ such that $v_i(A_i^{t-1}) \geq v_i(A_j^{t-1} \setminus \{g^*\})$. Then, we obtain

$$\begin{aligned} v_i(\tilde{A}_i^t) &= v_i(A_i^{t-1}) + v_i(\{g_i^*\}) \\ &\geq v_i(A_i^{t-1}) + v_i(\{g_j^*\}) & (5) \\ &\geq v_i(A_j^{t-1} \setminus \{g^*\}) + v_i(\{g_j^*\}) & (6) \\ &= v_i(\tilde{A}_j^t \setminus \{g^*\}), \end{aligned}$$

where (5) holds by (4), and (6) holds by the inductive hypothesis.

We conclude that, in both cases, there exists a good $g^* \in \tilde{A}_j^t$ such that $v_i(\tilde{A}_i^t) \geq v_i(\tilde{A}_i^t \setminus \{g^*\})$.

After obtaining the partial allocation $\tilde{A}^t$, the REC algorithm applies the envy-cycle-elimination procedure to $\tilde{A}^t$, and the resulting allocation is A^t. As each agent is not envious up to one good under the allocation $\tilde{A}^t$ and w.r.t. the private valuation, if we apply the envy-cycle-elimination procedure to $\tilde{A}^t$ the above fairness guarantee continues to hold under the new allocation A^t. Indeed, the envy-cycle-elimination procedure, when assigning a new bundle to some agent, the valuation of such agent does not decrease if compared with that she had for the previously assigned bundle. Thus, we have that, for any i, j, there exists a good $g^* \in \tilde{A}_j^t$ such that $v_i(A_i^t) \geq v_i(A_i^t \setminus \{g^*\})$. This shows the inductive step, and concludes the proof of the lemma. □

By using Lemma 2 and 3 with $t = T$, we obtain that no agent is envious up to one good, with respect to both public and private valuations, thus the claim of the theorem follows.

Remark 1. We observe that the REC algorithm would return an EF-$(1,1)$ even if the public valuation is not common to all agents, but agents assign the same ranking to the goods under their individual public valuation. Indeed, under such partial generalization, the REC algorithm would order the goods according to the common ranking, and the proof of Lemma 2 (stating that the returned allocation is envy-free-up-to-one good w.r.t. the public valuation) would continue to hold, as it only uses the fact the goods ranking is common.

PRR Algorithm

Algorithm 3: PRR Algorithm

Require: An instance $I = (N, M, w, (v_i)_{i \in N})$ of bicriteria fair allocation.
Ensure: An allocation $A = \mathsf{PRR}(I)$.
Let $T := \lceil m/n \rceil$.
Compute an ordering $g_1, g_2, \ldots, g_m$ of all goods in such a way that $w(g_1) \geq w(g_2) \geq \ldots \geq w(g_m)$.
Let $M_t := \{g_t, g_{t+1}, \ldots, g_{t+n-1}\}$ for any $t \in [T-1]$ and $M_T := \{g_T, \ldots, g_m\}$.
Let $A := (A_1, \ldots, A_n)$ be the empty allocation.
for $t = 1, \ldots, T := \lceil m/n \rceil$ **do**
  Compute a random ordering $\sigma_1 \prec \sigma_2 \prec \ldots \prec \sigma_n$ of all agents.
  $i \leftarrow 1$.
  while $M_t \neq \emptyset$ **do**
    $g^* \leftarrow g_{(t-1)n+i}$
    $A_{\sigma_i} \leftarrow A_{\sigma_i} \cup \{g^*\}$.
    $M_t \leftarrow M_t \setminus \{g^*\}$.
    $i \leftarrow i + 1$.
  end while
end for
return $A := (A_1, \ldots, A_n)$.

Proof of Theorem 2

Strategy-proofness holds since PRR does not use any information about the private valuations of the agents. Thus, agents have no benefit in misreporting their valuations (that, indeed, are not even used by the algorithm).

We observe that the complexity of the algorithm depends on the ordering of all goods according to the public valuation, that requires $O(m \log(m))$ steps. Thus, we have shown that PRR runs in polynomial time and, in the remainder of the proof, we can only focus on the fairness guarantee of PRR.

We assume w.l.o.g. that $m \geq n$, otherwise PRR will trivially returns an EF-$(1,1)$ allocation, and the claim follows. By exploiting the same arguments as in Lemma 2, we can show that the resulting allocation is envy-free up to one good w.r.t. to the public valuation.

It remains to show the probabilistic fairness guarantee on the private valuation, and to do this, we resort the Hoeffding's concentration bounds (Hoeffding 1963). For any

$i, j \in N$ and $t \in [T]$, let (i) $X_{i,j,t}$ be the random variable equal to the private valuation of agent i for the good that agent j receives at round t, (ii) $Y_{i,j} := \sum_{t \in T} X_{i,j,t}$, i.e., $Y_{i,j}$ is the random valuation of agent i for the bundle assigned by PRR to agent j, (iii) $W_{i,j,t} := X_{i,j,t} - X_{i,i,t}$, (iv) $Z_{i,j} := Y_{i,j} - Y_{i,i} = \sum_{t \in T} W_{i,j,t}$, i.e., $Z_{i,j}$ is the difference between the valuation of agent i for the bundle assigned to agent j and her own bundle, and (v) $Q := \max_{i,j \in N} Z_{i,j}$ is the maximum additive envy among all agents.

In the remainder of the proof, we will show that the following inequality holds:

$$\mathbb{P}(Q < \delta) \geq 1 - \frac{1}{n^\beta}, \tag{7}$$

where $\delta := \left\lceil \alpha \sqrt{2(\beta+2)\log(n) \left\lceil \frac{m}{n} \right\rceil} \right\rceil$. As α is defined as the maximum goods valuation and the minimum non-zero valuation is 1 (w.l.o.g., by the assumptions on the model), if $Q < \delta$ holds, then each agent can delete at most δ goods from the bundle of each other to be non-envious according to her private valuation. Thus, by showing that $Q < \delta$ holds with probability at least $1 - \frac{1}{n^\beta}$ (i.e., inequality (7)), we get the fairness guarantee on the private valuation stated in the theorem, and this would conclude the proof.

We have that

$$\begin{aligned}
\mathbb{P}(Q < \delta) &= \mathbb{P}\left(\max_{i,j \in N} Z_{i,j} < \delta\right) \\
&= \mathbb{P}(Z_{i,j} < \delta,\ \forall i, j \in N) \\
&= 1 - \mathbb{P}(\exists i, j \in N :\ Z_{i,j} \geq \delta) \\
&\geq 1 - \sum_{i,j \in N} \mathbb{P}(Z_{i,j} \geq \delta) \\
&= 1 - \sum_{i,j \in N} \mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} \geq \delta\right). \qquad (8)
\end{aligned}$$

The following lemma resorts to the Hoeffding's concentration bounds to provide an upper bound on each term $\mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} \geq \delta\right)$.

Lemma 4. *For any $i, j \in N$, we have*

$$\mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} \geq \delta\right) \leq \frac{1}{n^{\beta+2}}. \tag{9}$$

Proof of Lemma 4. As $Y_{i,j}$ is the random valuation of agent i for the bundle assigned to agent j, and the agents ordering are picked uniformly at random by PRR, by symmetry arguments we necessarily have that, for any fixed $i \in N$, $\mathbb{E}[Y_{i,j}]$ does not depend on j. Thus, we have

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t \in [T]} W_{i,j,t}\right] = \mathbb{E}[Z_{i,j}] \\
&= \mathbb{E}[Y_{i,j} - Y_{i,i}] = \mathbb{E}[Y_{i,j}] - \mathbb{E}[Y_{i,i}] = 0, \qquad (10)
\end{aligned}$$

where the last equality holds by the above observations. We observe that the random variables $(W_{i,j,t})_{t \in [T]}$ are independent, as the agents ordering selected by PRR at each round $t \in [T]$ are picked independently. Furthermore, as each variable $W_{i,j,t}$ represents the difference $X_{i,j,t} - X_{i,i,t}$ between two goods valuations bounded by α from above, we necessarily have that $W_{i,j,t}$ belongs to interval $[a, b] := [-\alpha, \alpha]$. Then, by applying the Hoeffding's inequality, we get

$$\begin{aligned}
&\mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} \geq \delta\right) \\
&= \mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} - \mathbb{E}\left[\sum_{t \in [T]} W_{i,j,t}\right] \geq \delta\right) \qquad (11) \\
&\leq e^{-\frac{2\delta^2}{\sum_{t \in [T]}(b-a)^2}} = e^{-\frac{\delta^2}{\lceil m/n \rceil 2\alpha^2}} \qquad (12) \\
&\leq \frac{1}{n^{\beta+2}}, \qquad (13)
\end{aligned}$$

where (11) holds by (10), (12) holds by Hoeffding's inequality and (13) holds by definition of δ. By (13), the claim follows. □

By using the upper bound provided by Lemma 4 in (8), we get

$$\begin{aligned}
\mathbb{P}(Q < \delta) &\geq 1 - \sum_{i,j \in N} \mathbb{P}\left(\sum_{t \in [T]} W_{i,j,t} \geq \delta\right) \\
&\geq 1 - \sum_{i,j \in N} \frac{1}{n^{\beta+2}} \\
&= 1 - n^2 \cdot \frac{1}{n^{\beta+2}} \\
&= 1 - \frac{1}{n^\beta},
\end{aligned}$$

that is, we showed inequality (7). Thus, each agent can delete at most δ goods from the bundle of each other to be non-envious according to her private valuation, with probability at least $1 - \frac{1}{n^\alpha}$, and this concludes the proof.

Remark 2. As observed in Remark 1 for the REC algorithm, we have that the PRR algorithm would return an allocation that is envy-free up to one good w.r.t. to the public valuation, even if the public valuation is not common to all agents, but agents assign the same ranking to the goods under their individual public valuation. Furthermore, the fairness guarantee on the public valuation always holds, and not only with high probability, even if the algorithm only knows the ranking of goods, and not the precise public valuation.