Approximately Fair Allocation of Indivisible Items with Random Valuations

Alessandro Aloisio
Università degli Studi Internazionali di Roma
Rome, Italy
alessandro.aloisio@unint.eu

Vittorio Bilò
Università del Salento
Lecce, Italy
vittorio.bilo@unisalento.it

Antonio Mario Caruso
Università del Salento
Lecce, Italy
antonio.caruso@unisalento.it

Michele Flammini
Gran Sasso Science Institute
L'Aquila, Italy
University of Calabria
Rende, Italy
michele.flammini@gssi.it

Cosimo Vinci
Università del Salento
Lecce, Italy
cosimo.vinci@unisalento.it

ABSTRACT

In this work, we consider the problem of fairly allocating a set of indivisible items to agents, who have additive and random valuations for the bundles of items they receive. The valuations that each agent has for all items are independent and bounded, and their realizations are only revealed after allocating the items. The goal is to determine an allocation that minimizes, in expectation, the maximum envy that an agent has for the bundle assigned to each other, without knowing in advance the realization of the random valuations.

We first show how to compute in polynomial time and deterministically an allocation that guarantees an expected maximum envy of at most $O(w\sqrt{\ln(n)m/n})$, where n is the number of agents, m is the number of items and w is the maximum valuation for each item. Furthermore, we show that the above bound cannot be improved, that is, there is an instance for which the expected maximum envy of any allocation is at least $\Omega(w\sqrt{\ln(n)m/n})$. Finally, we resort to randomized algorithms that return (random) allocations satisfying further efficiency guarantees, such as ex-ante envy-freeness and ex-ante Pareto optimality. If we relax the constraint of ex-ante Pareto optimality, we provide an algorithm that still works without knowing the probability distributions of agent valuations.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory**.

KEYWORDS

Envy-free allocations, Pareto optimality, Randomness

ACM Reference Format:

Alessandro Aloisio, Vittorio Bilò, Antonio Mario Caruso, Michele Flammini, and Cosimo Vinci. 2024. Approximately Fair Allocation of Indivisible Items with Random Valuations. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

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

1 INTRODUCTION

Fair allocation problems have garnered significant attention from a multitude of researchers in recent decades. An early work in this domain is attributed to [25]. These types of problems have implications across diverse academic fields, encompassing mathematics, economics, and social science. They can be primarily classified into two categories: problems that pertain to *divisible items* and those that concern *indivisible items* (see [2, 14] for comprehensive surveys). Our study falls under the latter category, where we address a broadly applicable framework in which m indivisible items need to be allocated among n *agents* having a *valuation* for each *bundle* of items, and the aim is to allocate all items in a way that is as fair as possible. An item allocation is called *envy-free* if each agent prefers her own bundle to the one assigned to any other agent. For divisible items, envy-free allocations always exist [26, 27] and several protocols to find one have been developed (e.g., [6]), but, for the setting with indivisible items, it may not always be possible to allocate items fairly. Thus, a feasible goal becomes that of finding *approximately fair allocations*, in which the maximum envy among agents is somehow bounded.

In most of the prior literature, the problem of finding approximately fair allocations has been studied under valuations which are deterministic and known to the agents, and several existential results and polynomial time algorithms have been provided [15, 17, 21]. In particular, [21] considered the general case of monotone valuations over bundles, and provide a polynomial time algorithm to compute an allocation whose maximum envy is bounded by the maximum marginal increment. Their algorithm also satisfies a stronger fairness notion introduced by [15], *envy-freeness-up-to-one-good* (EF1), which requires that each agent can recover envy-freeness by deleting at most one item from the bundle assigned to any other agent. [17] showed that a simple round-robin algorithm guarantees EF1 for the restricted case of additive valuations, where the valuation for a bundle is given by the sum of the values of its elements. Several works [1, 3, 9, 10, 17, 18] also focused on the existence and computation of EF1 allocations satisfying further desirable efficiency properties, such as *Pareto optimality*, under which it is not possible to improve the utility of an agent without reducing that of somebody else.

More recently, several frameworks in which allocations and agent valuations are not deterministic, but *random*, have received considerable attention from the scientific community. Random allocations, generally called *lotteries*, have been successfully used to obtain approximately fair allocations that are also *ex-ante envy-free* (i.e., the expected value of each agent is at least as large as that she would have for the bundles assigned to other agents) [4, 5, 20] and/or *ex-ante (fractional) Pareto optimality* [4, 16]. Regarding random valuations, most of the works have focused on settings in which the agent valuations are picked from probability distributions verifying some mild assumptions, and their realizations are observed before allocating the items [7, 8, 19, 22–24]. In this work, we consider a fair allocation model where each agent has an independent random valuation for each item, that can be observed after allocating the items. This setting is motivated by real-life scenarios in which agents may not know the exact valuation for an item, but rather discovers it after receiving the item or observing how another agent enjoys it.

2 MODEL AND DEFINITIONS

Given a non-negative integer k, let $[k] := \{1, \ldots, k\}$ denote the set of the first k positive integers. We denote random variables by bold letters, while we use italic letters to denote their realizations. Let $N := [n]$ be a set of n *agents* and $M := \{g_1, \ldots, g_m\}$ be a set of m *items*. Each agent $i \in N$ has a *random item valuation* $\boldsymbol{v}_{i,j}$ for each item g_j, distributed according to a probability distribution $\mathcal{P}_{i,j}$ that takes values in $[0, w]$, for some $w > 0$; we assume that, for any $i \in [n]$, all the distributions in $\{\mathcal{P}_{i,j} : g_j \in M\}$ are independent. Let $p_{i,j} := \mathbb{E}[\boldsymbol{v}_{i,j}]$ denote the *expected valuation* of agent i for item g_j; in the following, we implicitly assume that $p_{i,j}$ can be computed via a polynomial-time oracle; the assumption on the polynomial-time oracle can be relaxed, by assuming that we are able to sample from each distribution $\mathcal{P}_{i,j}$. If $\boldsymbol{v}_{i,j}$ is constant for any $i \in N$ and $g_j \in M$ (that is, $\boldsymbol{v}_{i,j} = p_{i,j}$), we say that valuations are *deterministic*. As we assume *additive valuations*, for each agent i, a random valuation over *bundles* $\boldsymbol{v}_i$ such that $\boldsymbol{v}_i(A) = \sum_{g_j \in A} \boldsymbol{v}_{i,j}$ for each bundle $A \subseteq M$ is induced. The triple $(N, M, (\mathcal{P}_{i,j})_{i \in N, g_j \in M})$ constitutes the input instance of the problem described below. An *allocation* $\vec{A} = (A_1, \ldots, A_n)$ is a partition of M in n bundles, where A_i is the bundle assigned to agent i. Given an allocation $\vec{A}$ and a realization $(v_{i,j})_{i,j}$ of the random item valuations $(\boldsymbol{v}_{i,j})_{i,j}$, the *maximum envy* under allocation $\vec{A}$ and realization $(v_{i,j})$ is defined as $Envy(\vec{A}, (v_{i,j})_{i,j}) := \max_{i,h \in N} v_i(A_h) - v_i(A_i) = \max_{i,h \in N} \left(\sum_{g_j \in A_h} v_{i,j} - \sum_{g_j \in A_i} v_{i,j} \right)$, that is, the maximum difference between the valuation that any agent i has for the bundle assigned to any other agents and the valuation of the bundle that she receives. Analogously, the *expected maximum envy* under allocation $\vec{A}$ is defined as $Envy(\vec{A}) := \mathbb{E}\left[Envy(\vec{A}, (\boldsymbol{v}_{i,j})_{i,j}) \right] = \mathbb{E}\left[\max_{i,h \in N} \left(\sum_{g_j \in A_h} \boldsymbol{v}_{i,j} - \sum_{g_j \in A_i} \boldsymbol{v}_{i,j} \right) \right]$, i.e., the expected value of the maximum envy.

3 OUR RESULTS

Considering that agent valuations are random, we aim at finding an allocation with a bounded expected maximum envy. We point out that the problem of minimizing the maximum expected envy has been also considered in [11], where randomized algorithms for computing assignments in an online setting are designed.

We show how to compute in polynomial time an allocation that guarantees an expected maximum envy $Envy(\vec{A})$ of at most $w\left(1 + \frac{1}{\sqrt{2}}\right)\sqrt{\ln(n)\left\lceil \frac{m}{n} \right\rceil} + w \in O\left(w\sqrt{\ln(n)\frac{m}{n}}\right)$ (Theorem 1), that is sublinear in both n and m. We also show that the above guarantee cannot be improved, as we provide an instance for which any allocation with a sufficiently large number of items has an expected maximum envy of at least $\Omega(w\sqrt{\ln(n)m/n})$ (Theorem 2). This result introduces a separation with the classic framework of deterministic valuations, where the envy is bounded by the maximum marginal increment w [21].

Then, we consider the possibility of allocating items in a random way, and this allows to obtain other desirable efficiency properties, such as ex-ante envy freeness and ex-ante Pareto optimality. In particular, we first provide a polynomial time randomized algorithm that returns a random allocation satisfying ex-ante envy-freeness and matching the above sublinear bound on the expected maximum envy (Theorem 3). Furthermore, our randomized algorithm works even in the online setting considered by [11], and asymptotically improves the expected maximum envy achieved by their randomized algorithm. Finally, we show how to compute in polynomial time a random allocation that is both ex-ante envy-free and ex-ante Pareto optimal and has an expected maximum envy of at most $O(w\sqrt{\ln(n)m})$ (Theorem 4). By exploiting the lower bound of Theorem 2, we show that the upper bound of Theorem 3 cannot be asymptotically improved, and the one provided in Theorem 4 cannot be improved up to a poly-logarithmic factor.

4 FUTURE WORKS

Our work leaves several research directions. An open problem left by our work is that of closing the gap between $O(w\sqrt{\ln(n)m})$ and $\Omega(w\sqrt{m})$ on the expected maximum envy of allocations which satisfy both ex-ante EF and PO. Similarly as done in [15] for the notion of EF1, it would be interesting to know how many items, in expectation, each agent should remove from any bundle to recover the envy-freeness. It would be also nice to embed the stochastic aspects introduced in this work in other fair allocation settings (e.g., fair allocation with graph connectivity constraints [12, 13, 28]) and other fairness criteria (e.g., proportionality, maximin share [2, 14]).

ACKNOWLEDGMENTS

We acknowledge: the PNRR MIUR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI; the PNRR MIUR project VITALITY (ECS00000041), Spoke 2 ASTRA - Advanced Space Technologies and Research Alliance; the PNRR MIUR project HaMMon - Hazard Mapping and vulnerability Monitoring, Innovation Grant; the PNRR MIUR project SERICS - Security and Rights in Cyber Space (PE00000014), Spoke 2 - Misinformation and Fakes; the Italian MIUR PRIN 2017 project ALGADIMAR - Algorithms, Games, and Digital Markets (2017R9FHSR_002); GNCS-INdAM.

REFERENCES

[1] Hannaneh Akrami, Bhaskar Ray Chaudhury, Martin Hoefer, Kurt Mehlhorn, Marco Schmalhofer, Golnoosh Shahkarami, Giovanna Varricchio, Quentin Vermande, and Ernest van Wijland. 2022. Maximizing Nash Social Welfare in 2-Value Instances. In *Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2022)*. 4760–4767.

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

[3] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. 2021. Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.* 863 (2021), 69–85.

[4] Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2023. Best of Both Worlds: Ex Ante and Ex Post Fairness in Resource Allocation. *Operations Research* (2023).

[5] Haris Aziz, Aditya Ganguly, and Evi Micha. 2023. Best of Both Worlds Fairness under Entitlements. In *The 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*. 941–948.

[6] Haris Aziz and Simon Mackenzie. 2016. A Discrete and Bounded Envy-Free Cake Cutting Protocol for Four Agents. In *The Forty-Eighth Annual ACM Symposium on Theory of Computing (STOC 2016)*. 454–464.

[7] Yushi Bai, Uriel Feige, Paul Gölz, and Ariel D. Procaccia. 2022. Fair Allocations for Smoothed Utilities. In *The 23rd ACM Conference on Economics and Computation (EC 2022)*. 436–465.

[8] Yushi Bai and Paul Gölz. 2022. Envy-Free and Pareto-Optimal Allocations for Agents with Asymmetric Random Valuations. In *The Thirty-First International Joint Conference on Artificial Intelligence (IJCAI 2022)*. 53–59.

[9] Siddharth Barman and Sanath Kumar Krishnamurthy. 2019. On the Proximity of Markets with Integral Equilibria. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*. 1748–1755.

[10] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. 2018. Finding Fair and Efficient Allocations. In *The 2018 ACM Conference on Economics and Computation (EC 2018)*. 557–574.

[11] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. 2018. How to Make Envy Vanish Over Time. In *EC 2018*. 593–610.

[12] Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S. Zwicker. 2022. Almost envy-free allocations with connected bundles. *Games Econ. Behav.* 131 (2022), 197–221.

[13] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. 2017. Fair Division of a Graph. In *The Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI 2017)*. 135–141.

[14] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair Allocation of Indivisible Goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 284–310.

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

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

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

[18] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. 2019. Group Fairness for the Allocation of Indivisible Goods. In *The Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*. 1853–1860.

[19] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. 2014. The Computational Rise and Fall of Fairness. In *The Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2014)*. 1405–1411.

[20] Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. 2023. Best of Both Worlds: Agents with Entitlements. In *The 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*. 564–572.

[21] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *EC 2004*. 125–131.

[22] Pasin Manurangsi and Warut Suksompong. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences* 89 (2017), 100–108.

[23] Pasin Manurangsi and Warut Suksompong. 2020. When Do Envy-Free Allocations Exist? *SIAM Journal on Discrete Mathematics* 34, 3 (2020), 1505–1521.

[24] Pasin Manurangsi and Warut Suksompong. 2021. Closing Gaps in Asymptotic Fair Division. *SIAM Journal on Discrete Mathematics* 35, 2 (2021), 668–706.

[25] Hugo Steinhaus. 1948. The problem of fair division. 16 (1948), 101–104.

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

[27] Francis Edward Su. 1999. Rental Harmony: Sperner's Lemma in Fair Division. *Amer. Math. Monthly* 106 (1999), 930–942.

[28] Warut Suksompong. 2019. Fairly allocating contiguous blocks of indivisible items. *Discret. Appl. Math.* 260 (2019), 227–236.