Near-Optimal Online Resource Allocation in the Random-Order Model

Extended Abstract

Saar Cohen
Department of Computer Science
Bar Ilan University, Israel
saar30@gmail.com

Noa Agmon
Department of Computer Science
Bar Ilan University, Israel
agmon@cs.biu.ac.il

ABSTRACT

We study the problem of allocating either divisible or indivisible items (goods or chores) among a set of agents, where the items arrive *online*, one at a time. Each agent's non-negative value for an item is set by an *adversary* upon the item's arrival. Our focus is on a unifying algorithmic framework for finding online allocations that treats both fairness and economic efficiency. For this sake, we aim to optimize the *generalized means* of agents' received values, covering a spectrum of welfare functions including average utilitarian welfare and egalitarian welfare. In the traditional adversarial model, where items arrive in an *arbitrary* order, no algorithm can give a decent approximation to welfare in the worst case. To escape from this strong lower bound, we consider the *random-order model*, where items arrive in a *uniformly random* order. This model provides us with a major breakthrough: we devise algorithms that guarantee a ***nearly-optimal*** competitive ratio for certain welfare functions, if the welfare obtained by the optimal allocation is sufficiently large. We prove that our results are *almost tight*: if the optimal solution's welfare is strictly below a certain threshold, then no nearly-optimal algorithm exists, even in the random-order model.

KEYWORDS

Online Fair Division; Resource Allocation; Random-Order Model; Generalized Means

ACM Reference Format:
Saar Cohen and Noa Agmon. 2024. Near-Optimal Online Resource Allocation in the Random-Order Model: Extended Abstract. 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.

1 INTRODUCTION

Fair and efficient allocation of resources (or items) to heterogeneous agents is a vital concern in many real-life scenarios [13, 25, 34, 38], ranging from course allocation [11], peer reviewing [5, 33, 36] and food donations [1], to distributing medical equipment and vaccines [3, 32]. Some items, such as a cake [38] or a land [17], may be *divisible* and can be split *fractionally* among several agents. Yet, items like rare artworks cannot be divided without losing their value, and are thus *indivisible* [2]: Each item must be allocated to a single agent. Prior work also distinguishes between the allocation of *goods* that bring happiness [12], such as birthday cakes shared among friends, and the assignment of undesirable *chores* that induce regret [4, 24], like household cleaning duties. From a welfarist perspective, fair division revolves around maximizing overall happiness from goods and minimizing overall regret from chores.

In this paper, we consider numerous practical scenarios where the items arrive *online*, one at a time [1, 21, 28, 39, 40, 42]. Example applications include food banks [1, 35], online advertising [10, 23] and so on. Our work explores online fair division from a welfarist viewpoint by employing a unified algorithmic framework that addresses both fairness and economic efficiency. Unlike prior studies that mainly focused on online allocation of *goods*, our analysis also considers *chores*. Depending on whether all the items are goods or chores, our goal is devising allocation algorithms that either maximize or minimize the *generalized means* of agents' received values [8, 16], encompassing various welfare functions such as average utilitarian welfare and egalitarian welfare.

2 PRELIMINARIES

We consider the problem of allocating a set of T items (goods or chores) among a set of $n \geq 2$ agents in an online manner. The items arrive one by one in T rounds overall. For each agent i, her value $v_i^t \in [0,1]$ for item t is revealed only when the item arrives (i.e., in round t). Existing literature often considers the *(worst-case) adversarial model* [22, 42], where an adversary controls agents' values and the items' order of arrival to harm algorithm performance. Yet, even in *more restricted* scenarios, this model has been proven to be too pessimistic and no algorithm can surpass trivial ones [6, 8, 37]. To avoid such negative results and enable the design of more effective algorithms in practice, there has been a shift to *beyond worst-case* models. In our work, we explore one notable such alternative, the *random-order model*, where the adversary first selects the agents' values for the items, but the items are then presented in a *uniformly random order*. In online fair division, optimizing the generalized means under this model is a major open problem [7, 8].

We consider the cases of *divisible* items (where each item can be matched to multiple agents fractionally) and *indivisible* items (where each item must be matched to a single agent integrally). In the *divisible* case, when item t arrives, an online algorithm must *irrevocably* and *immediately* decide on the fraction $x_i^t \in [0,1]$ of item t that should be assigned to each agent i such that $\Sigma_{i=1}^n x_i^t \leq 1$. The *indivisible* case is obtained when $x_i^t \in \{0,1\}$ for each agent i and item t. Let $\mathbf{x}_i := (x_i^t)_{t\in[T]} \in [0,1]^T$ be the *bundle* of each agent i and $\mathbf{x} := (\mathbf{x}_i)_{i\in[n]} \in [0,1]^{n\times T}$ be the fractional allocation of the

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. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

items among all the agents. We also denote the allocation at time t as $\mathbf{x}^t := (x_i^t)_{i\in[n]} \in [0,1]^n$. We focus on agents with *additive* valuations. In *goods-only* problems, an agent's valuation is referred to as her *utility*, and, in *chores-only* problems, as her *disutility*. Let $\mathbf{v}(\mathbf{x}) := (v_i(\mathbf{x}_i))_{i=1}^n$ be the agents' *joint* valuation for an allocation $\mathbf{x}$ and $\mathbf{v}(\mathbf{x}^t) := (v_i^t x_i^t)_{i=1}^n$ be the agents' joint valuation at time t.

The agents' valuations under an allocation $\mathbf{x}$ can be aggregated into a single non-negative value using a welfare function, which measures the fairness and efficiency of the allocation. In our work, we focus on the *generalized means* of agents' valuations. Formally, for $p \in \mathbb{R}$, the *p-mean welfare* of an allocation $\mathbf{x}$ is given by:

$$M_p(\mathbf{x}) := M_p(\mathbf{v}(\mathbf{x})) = \left(\frac{1}{n}\sum_{i=1}^{n} v_i(\mathbf{x}_i)^p\right)^{1/p} \quad (1)$$

This family of objectives captures multiple well-known fairness and efficiency measures: $p = 1$ yields the common *average utilitarian welfare*, the limit $p \to \infty$ provides the *maximax welfare*, the limit $p \to -\infty$ induces the *egalitarian welfare*, and the limit $p \to 0$ gives the *Nash welfare*. In goods-only problems, we aim to *maximize* the generalized means of agents' utilities, while in chores-only problems, our goal is to *minimize* it. As standard in online algorithms (See, e.g., [20]), we measure an algorithm's performance in terms of its *competitive ratio*, i.e., the worst-case ratio between the welfare of the algorithm's allocation to that of the optimal (offline) solution.

3 OUR CONTRIBUTIONS AND TECHNIQUES

Near-Optimal Algorithms for Divisible Items. In instances with only *divisible* goods or *divisible* chores, the random-order model offers a significant breakthrough: we develop algorithms that ensure a ***nearly-optimal*** competitive ratio for certain welfare functions, if the optimal allocation yields a sufficiently large welfare. Before describing our results more formally, we provide the intuition behind our algorithms. Note that the naive greedy approach of allocating each arriving item in a greedy manner incurs a high additive regret as the change in welfare can vary rapidly from one round to another. Inspired by the scheme of Molinaro for online scheduling [31], we avoid this problem by allocating items greedily with respect to a *smoother* version of the generalized-means.

Further, a key challenge in the random-order model is that there are correlations among the items arriving in different time instants as they are being sampled *without replacement* from the underlying set of items. Hence, our algorithms combine the use of the smoothed function with the *restart strategy*: the arriving items are bisected into two sequences, which reduces the correlations among them so as to obtain a *near-optimal* competitive ratio under mild conditions. This holds as each item's allocation only depends on at most $T/2 - 1$ other allocations. Formally, for the case of *divisible* items, our algorithms have the following guarantees:

Theorem 3.1. *In the random-order model with* divisible **goods***, for any $\varepsilon \in [0,1]$, the following are satisfied:*

(1) *For each $p \in (1,\infty)$, if the* maximum *p-mean welfare is at least $\Omega(p(1-n^{-1/p})/\varepsilon^2)$, then there is an algorithm with a* nearly-optimal *competitive ratio of at least $1-\varepsilon$ for* maximizing *the p-mean welfare.*
(2) *For each $p \in (-\infty,-1)$, if $-3p(1-n^{1/p})/\varepsilon$ is little-o of the* minimum *$(-p)$-mean welfare, then, under mild assumptions, there is an algorithm with a* nearly-optimal *competitive ratio of at least $1-\varepsilon$ for* maximizing *the p-mean welfare.*

Theorem 3.2. *In the random-order model with* divisible **chores***, for any $\varepsilon \in [0,1]$, the following are satisfied:*

(1) *For each $p \in (1,\infty)$, if $3p(1-n^{-1/p})/\varepsilon$ is little-o of the* minimum *p-mean welfare, then there is an algorithm with a* nearly-optimal *competitive ratio of at most $1+\varepsilon$ for* minimizing *the p-mean welfare.*
(2) *For each $p \in (-\infty,-1)$, if the* maximum *$(-p)$-mean welfare is at least $\Omega(-p(1-n^{1/p})/\varepsilon^2)$, then, under mild assumptions, there is an algorithm with a* nearly-optimal *competitive ratio of at most $1+\varepsilon$ for* minimizing *the p-mean welfare.*

We prove that the conditions on the generalized means in Theorems 3.1 and 3.2 are *necessary*, i.e., our results are *almost tight*. However, our algorithms still have theoretically good guarantees in this case and are near-optimal within constants. Formally:

Theorem 3.3. *For both divisible* goods *and* chores*, if the* maximum *generalized means is strictly below a certain threshold, then* no *nearly-optimal algorithm exists, even in the random-order model.*

Randomized Rounding for Indivisible Items. For *indivisible* items, we show how standard randomized rounding can convert the *fractional* allocations produced by our algorithms into *integral* allocations while maintaining their guarantees. Formally:

Theorem 3.4. *In the random-order model with* indivisible items*, if the number of agents n is sufficiently large, then, under the conditions of Theorems 3.1-3.2, standard randomized rounding combined with each of our algorithms has a* nearly-optimal *competitive ratio.*

4 DISCUSSION AND FUTURE WORK

Note that the random-order model is *stronger* than the *i.i.d. model* [14, 42], where items are i.i.d. drawn from an *unknown* distribution picked by an adversary. Indeed, one can view i.i.d. draws as first sampling from an underlying distribution and then randomly permuting them. Thus, our results also apply to this model.

Our research opens the way for many future works. An immediate direction is examining the random-order model for $p \in [-1,1]$, which we believe to be inherently harder as it includes the Nash social welfare, known to be APX-hard even in *offline* setting with indivisible items [27]. Future studies can also develop algorithms that are effective on both purely stochastic and purely adversarial inputs [30], or even on inputs that are a *mix* of both [18, 30]. Another interesting direction is considering scenarios where input items may follow a Poisson arrival process [29, 41] or more general arrival distributions [15, 19]. Future research may also consider other objectives, such as proportional fairness [7] and envy-freeness [9, 42]. While our work focuses on allocating either goods or chores, future works should consider *mixtures* of them, which has proven to be inherently challenging even in *offline* settings and under domain restrictions [24, 26]. Finally, investigating domains *beyond* agents with additive valuations may lead to additional profound insights regarding online fair division.

ACKNOWLEDGMENTS

This research was funded in part by ISF grant #1563/22.

REFERENCES

[1] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online fair division: analysing a food bank problem. In *Proceedings of the 24th International Conference on Artificial Intelligence*. 2540–2546.

[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. *Artificial Intelligence* (2023), 103965.

[3] Haris Aziz and Florian Brandl. 2021. Efficient, fair, and incentive-compatible healthcare rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 103–104.

[4] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. 2022. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems* 36 (2022), 1–21.

[5] Haris Aziz, Evi Micha, and Nisarg Shah. 2023. Group Fairness in Peer Review. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2889–2891.

[6] Siddhartha Banerjee, Vasilis Gkatzelis, Artur Gorokh, and Billy Jin. 2022. Online nash social welfare maximization with predictions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1–19.

[7] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. 2023. Proportionally Fair Online Allocation of Public Goods with Predictions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. 20–28.

[8] Siddharth Barman, Arindam Khan, and Arnab Maiti. 2022. Universal and tight online algorithms for generalized-mean welfare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4793–4800.

[9] Gerdus Benade, Aleksandr M Kazachkov, Ariel D Procaccia, and Christos-Alexandros Psomas. 2018. How to make envy vanish over time. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 593–610.

[10] Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. 2006. Online algorithms for market clearing. *Journal of the ACM (JACM)* 53, 5 (2006), 845–879.

[11] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.

[12] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–32.

[13] Yann Chevaleyre, Paul E Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A Rodríguez-Aguilar, and Paulo Sousa. 2006. Issues in Multiagent Resource Allocation. *Informatica* 30 (2006), 3–31.

[14] John Dickerson, Jonathan Goldman, Jeremy Karp, Ariel Procaccia, and Tuomas Sandholm. 2014. The computational rise and fall of fairness. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 28.

[15] John P Dickerson, Karthik A Sankararaman, Aravind Srinivasan, and Pan Xu. 2021. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. *ACM Transactions on Economics and Computation (TEAC)* 9, 3 (2021), 1–17.

[16] Soroush Ebadian, Rupert Freeman, and Nisarg Shah. 2022. Efficient Resource Allocation with Secretive Agents. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. 272–278.

[17] Edith Elkind, Erel Segal-Halevi, and Warut Suksompong. 2023. Keep your distance: Land division with separation. *Computational Geometry* 113 (2023), 102006.

[18] Hossein Esfandiari, Nitish Korula, and Vahab Mirrokni. 2015. Online allocation with traffic spikes: Mixing adversarial and stochastic models. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 169–186.

[19] Seyed Esmaeili, Sharmila Duppala, Davidson Cheng, Vedant Nanda, Aravind Srinivasan, and John P Dickerson. 2023. Rawlsian fairness in online bipartite matching: Two-sided, group, and individual. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 5624–5632.

[20] Amos Fiat and Gerhard J Woeginger. 1998. *Online algorithms: The state of the art*. Vol. 1442. Springer.

[21] Eric Friedman, Christos-Alexandros Psomas, and Shai Vardi. 2015. Dynamic fair division with minimal disruptions. In *Proceedings of the sixteenth ACM conference on Economics and Computation*. 697–713.

[22] Vasilis Gkatzelis, Alexandros Psomas, and Xizhi Tan. 2021. Fair and efficient online allocations with normalized valuations. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 5440–5447.

[23] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. 2007. Automated online mechanism design and prophet inequalities. In *AAAI*, Vol. 7. 58–65.

[24] Hadi Hosseini, Aghaheybat Mammadov, and Tomasz Wąs. 2023. Fairly Allocating Goods and (Terrible) Chores. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. International Joint Conferences on Artificial Intelligence Organization, 2738–2746.

[25] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2021. Fair and efficient allocations under lexicographic preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5472–5480.

[26] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2023. Fairly Dividing Mixtures of Goods and Chores under Lexicographic Preferences. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 152–160.

[27] Euiwoong Lee. 2017. APX-hardness of maximizing Nash social welfare with indivisible items. *Inform. Process. Lett.* 122 (2017), 17–20.

[28] Bo Li, Wenyang Li, and Yingkai Li. 2018. Dynamic fair division problem with general valuations. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 375–381.

[29] Mathieu Mari, Michał Pawłowski, Runtian Ren, and Piotr Sankowski. 2023. Online Matching with Delays and Stochastic Arrival Times. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 976–984.

[30] Vahab S Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. 2012. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 1690–1701.

[31] Marco Molinaro. 2017. Online and random-order load balancing simultaneously. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1638–1650.

[32] Parag A Pathak, Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez. 2021. Fair allocation of vaccines, ventilators and antiviral treatments: leaving no ethical value behind in health care rationing. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 785–786.

[33] Justin Payan and Yair Zick. 2022. I Will Have Order! Optimizing Orders for Fair Reviewer Assignment. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. International Joint Conferences on Artificial Intelligence Organization, 440–446.

[34] Benjamin Plaut and Tim Roughgarden. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics* 34, 2 (2020), 1039–1068.

[35] Canice Prendergast. 2017. How food banks use markets to feed the poor. *Journal of Economic Perspectives* 31, 4 (2017), 145–162.

[36] Nihar B Shah. 2022. Challenges, experiments, and computational solutions in peer review. *Commun. ACM* 65, 6 (2022), 76–87.

[37] Max Springer, MohammadTaghi Hajiaghayi, Debmalya Panigrahi, and Mohammad Khani. 2022. Online Algorithms for the Santa Claus Problem. In *Advances in Neural Information Processing Systems*, Vol. 35. Curran Associates, Inc., 30732–30743.

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

[39] Shai Vardi, Alexandros Psomas, and Eric Friedman. 2022. Dynamic fair resource division. *Mathematics of Operations Research* 47, 2 (2022), 945–968.

[40] Toby Walsh. 2011. Online cake cutting. In *Algorithmic Decision Theory: Second International Conference, ADT 2011, Piscataway, NJ, USA, October 26-28, 2011. Proceedings 2*. Springer, 292–305.

[41] Pan Xu and Yifan Xu. 2022. Equity Promotion in Online Resource Allocation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 9962–9970.

[42] David Zeng and Alexandros Psomas. 2020. Fairness-efficiency tradeoffs in dynamic fair division. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 911–912.