Allocating Resources with Imperfect Information

Doctoral Consortium

Shiji Xing
The Hong Kong Polytechnic University
Hong Kong, China
shi-ji.xing@connect.polyu.hk

ABSTRACT

The distribution of resources is a critical issue that impacts all aspects of the Internet and society, with fairness playing a key role. Current standard algorithms for distribution typically measure fairness through methods based on envy or proportionality, requiring precise numerical values. However, there is a clear discrepancy between how these algorithms are intended to work in theory and their application in real-world situations. This is because users often do not have exact information about the resources and struggle to assign a numerical value to them. Our goal is to explore settings where agents do not take exact numeric values as input. In this framework, we do not assume that individuals can specify exact numerical values for resources. Instead, we assume that each agent has an ordinal preference for the items. That is, given two items, an agent can identify which is better, without assigning cardinal values to them. We consider new criteria for fairness in this setting, and discuss about their achievability in this article. Besides, we investigate the Probabilistic Serial mechanism where agents also only provides ordinal ranking over items, and particularly research on the incentive ratio of the mechanism.

KEYWORDS

Fair Division; Ordinal Preference; Possible Fairness; Probabilistic Serial; Incentive Ratio

ACM Reference Format:

Shiji Xing. 2024. Allocating Resources with Imperfect Information: Doctoral Consortium. 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 BACKGROUND

Allocating resources fairly has been a challenge from the beginning of human society and is still of great significance in modern life. It occurs at many levels within both the digital world and social structures [8]. Traditional examples of fair division include allocating food and water among a community or evenly distributing inheritance among offsprings, and similar concepts are still applied to modern society in areas like cloud computing [4, 10].

The modern theory of fair allocation arguably originated in the work of Steinhaus [9] and has been the focus of economics, mathematics, and computer science for most of the last century [7]. To allocate resources fairly, one needs first to define fairness. There are generally two types of well-studied fairness, namely *envy freeness* (EF) and *proportionality* (PROP). Assume that in an allocation, each agent gets a bundle of resources. Each agent has an evaluation of all bundles. An allocation is envy-free if, for each agent, they think that their bundle is the best among all bundles. In addition, an allocation is considered proportional when each agent receives a bundle perceived to be at least as valuable as the total value of all resources divided by the number of agents. Put simply, an allocation is proportional if each agent feels they have obtained a fair share of the resources. In this work, we only consider indivisible items.

As indicated by the above definitions, in traditional research, problems are often defined with mathematically precise conditions, referred to as cardinal fairness; agents are assumed to be able to assign a concrete numeric value on any subset of the resources. However, this method often overlooks the practicality of implementing such algorithms. It misses the fact that the computational burden placed on users to express and communicate their precise preferences regarding resources could render the algorithms unusable; for instance, an algorithm might be rejected if users find it too complex to operate.

Perhaps more crucially, it is often challenging for users to articulate their preferences using precise cardinal values. For example, while individuals might recognize that a car is generally more valuable than a cellphone, quantifying the exact value difference is nearly impossible.

Therefore, we argue that a gradual relaxation of the perfect information assumption is necessary for theory to more closely align with reality and for theoretical insights to be deployed effectively in practice. In this context, our project aims to investigate the problem of fair allocation with imperfect information. We assume that agents have an ordinal ranking preference of all items so that they can compare any pair of items without assigning value to them.

2 ORDINAL FAIRNESS

When agents only have ordinal preferences on items, it is trivial to compare two items. However, there is no natural way to compare two bundles. In some cases, it is easy for an agent to compare two bundles: when one is strictly better than the other. If there exists a matching between items from the two bundles in which, for each pair of items, the item in a certain bundle is better than the other, then it is safe to say that certain bundle is preferred over the other. This is also called *stochastic dominance* (SD). However, the SD condition is strict, and not all bundles are comparable by SD. Therefore, other ways of comparison are needed.

Two ways are commonly used to extend preferences over items to preferences over bundles: *downward lexicographic* (DL) dominance and additivity.

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.

DL dominance assumes that the value differences between items are huge, such that, for an item o and an agent i, i will consider o more valuable than the bundle containing all items worse than o. Formally, let A and B be two bundles. If agent i DL-prefers bundle A over B, then there exists an integer k, such that i SD-prefers A' over B', where A' (*resp.* B') is the set of the k most preferred items in A (*resp.* B).

Additivity means that agents have additive cardinal valuations of the items, although the allocator does not know the exact function. Additive valuation means that each agent assigns a numeric value to each item and that their value to a bundle is the sum of the value of all the items in the bundle.

2.1 DL Dominance

Using DL dominance, we can compare two bundles in the ordinal setting and thus give a concrete definition of fairness. We mainly considered the DL version of EFX (DL-EFX) and MMS (DL-MMS), which are the relaxed versions of EF and PROP, respectively.

THEOREM 2.1. *Given any instance with arbitrary ordinal preferences, an allocation that is simultaneously DL-EFX, DL-MMS exists and can be computed in polynomial time [1].*

The theorem is proven by constructing an algorithm that can always find such an allocation given any instance, and the main technique behind the algorithm is the *exchange graph*. Similarly to the *envy cycle procedure* [6], the exchange graph aims to cease the envy relationships between agents by exchanging. However, instead of exchange bundles, the exchange graph tries to exchange single items among a chain of agents. This procedure allows the algorithm to handle envy relationships without affecting the structure of agents' bundles and is crucial to maintaining the MMS property.

2.2 Possible and Necessary Fairness

Another way to extend the ordinal preference over bundles is by additivity. We assume that each agent has an underlying cardinal additive valuation function that is consistent with their ordinal preference. Since we do not actually know the exact valuation function, we prompt two variations of fairness notions in this setting: possible and necessary fairness. We say that an allocation satisfies a certain possible fairness notion if there exists at least one set of cardinal valuations consistent with the provided ordinal rankings that fulfills the fairness criterion when those cardinal valuations are considered the true valuations of the agents. Similarly, necessary fairness demands that, for every unique cardinal valuation that aligns with the provided ordinal preference, the allocation must guarantee the fairness criteria under the cardinal realization.

Following the above discussion, we can define *possible EFX* (p-EFX), *possible MMS* (p-MMS), *necessary EFX* (n-EFX) and *necessary MMS* (n-MMS).

Since DL dominance can be expressed by additive cardinal valuations, possible versions of fair notions are weaker than their DL counterparts. Hence, we have a stronger result for possible fairness.

THEOREM 2.2. *For any instance with ordinal preferences, an allocation that is simultaneously p-EFX, p-MMS, and balanced exists and can be computed in polynomial time [1].*

Necessary fairness, on the other hand, is much more difficult to achieve. In fact, necessary EFX and necessary MMS do not always exist. Given an ordinal preference profile, one can have a DL cardinal valuation where the best item is much better than the rest; however, there exists another valid valuation where each item is of similar value with minor difference. Although both are consistent with ordinal preference, the two cardinal valuations require largely different allocation structures to satisfy fairness, rendering it impossible to consistently find an allocation that satisfies the criteria for necessary fairness.

3 INCENTIVE RATIO OF THE PROBABILISTIC SERIAL RULE

In addition to the algorithms discussed above, ordinal preference is also used as input to other allocation mechanisms, among which a well-studied one is the *Probabilistic Serial* (PS) Rule [2]. The PS Rule is a randomized allocation rule for indivisible items. In this mechanism, each agent will have an ordinal ranking over of items, and agents will try to 'consume' all the items. Agents have the same consumption rate and start at the same time. We assume that each item also has the same size. Agents will consume items according to their ordinal preferences, and the mechanism ends after all items are consumed. An agent's consumed share of a certain item is considered the probability that they can get it.

The PS Rule satisfied many desirable properties, such as envy-freeness, proportionality, and Pareto efficiency. However, the mechanism is not truthful since an agent can misreport their ordinal preference in order to gain more profit than reporting their real preference. A benchmark to measure this ability to manipulate is the incentive ratio [3]. The incentive ratio denotes the ratio between the highest value one agent can gain by manipulating their reported preference and the value they get when reporting according to their real preference. Note that no cardinal valuations of agents are required, since the ratio accounts for the worst-case scenario among all additive valuation functions. We have researched the incentive ratio for the PS Rule and extended an existing result [11] to a more general setting.

THEOREM 3.1. *The incentive ratio of Probabilistic Serial is at most* $2 - \frac{1}{2^{n-1}}$ *[5].*

4 FUTURE DIRECTIONS

Our findings indicate that, while possible fairness is achievable, necessary fairness presents significant challenges. It remains to be seen whether necessary fairness can always exist in some restricted settings or whether there are any approximations. Regarding the incentive ratio of the PS Rule, though the incentive ratio of the case where items are more than agents are tightly bounded, we believe that there exists a more accurate bound when agents are no fewer than items, which remains to be discovered.

ACKNOWLEDGMENTS

The author would expresses appreciation for his supervisor, Bo Li, for his significant support and advice. The author would also thank Haris Aziz, Ankang Sun, and Yu Zhou for the help and collaboration of this work.

REFERENCES

[1] Haris Aziz, Bo Li, Shiji Xing, and Yu Zhou. 2023. Possible Fairness for Allocating Indivisible Resources. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh (Eds.). ACM, 197–205. https://doi.org/10.5555/3545946.3598637

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

[3] Ning Chen, Xiaotie Deng, and Jie Zhang. 2011. How Profitable Are Strategic Behaviors in a Market?. In *ESA 2011 (Lecture Notes in Computer Science, Vol. 6942)*, Camil Demetrescu and Magnús M. Halldórsson (Eds.). Springer, 106–118.

[4] Robert Grandl, Ganesh Ananthanarayanan, Srikanth Kandula, Sriram Rao, and Aditya Akella. 2014. Multi-resource packing for cluster schedulers. In *ACM SIGCOMM 2014 Conference, SIGCOMM'14, Chicago, IL, USA, August 17-22, 2014*, Fabián E. Bustamante, Y. Charlie Hu, Arvind Krishnamurthy, and Sylvia Ratnasamy (Eds.). ACM, 455–466. https://doi.org/10.1145/2619239.2626334

[5] Bo Li, Ankang Sun, and Shiji Xing. 2024. Bounding the Incentive Ratio of the Probabilistic Serial Rule. In *Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024, Auckland, New Zealand, 6-10 May 2024*. ACM.

[6] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings 5th ACM Conference on Electronic Commerce (EC-2004), New York, NY, USA, May 17-20, 2004*, Jack S. Breese, Joan Feigenbaum, and Margo I. Seltzer (Eds.). ACM, 125–131. https://doi.org/10.1145/988772.988792

[7] Hervé Moulin. 2003. *Fair division and collective welfare.* MIT Press.

[8] Roman Seidl. 2018. Handbook of Computational Social Choice *by Brandt Felix, Vincent Conitzer, Ulle Endriss, Jerome Lang, Ariel Procaccia. J. Artif. Soc. Soc. Simul.* 21, 2 (2018). http://jasss.soc.surrey.ac.uk/21/2/reviews/4.html

[9] Hugo Steinhaus. 1949. Sur la division pragmatique. *Econometrica: Journal of the Econometric Society* (1949), 315–319.

[10] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. 2015. Large-scale cluster management at Google with Borg. In *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, Laurent Réveillère, Tim Harris, and Maurice Herlihy (Eds.). ACM, 18:1–18:17. https://doi.org/10.1145/2741948.2741964

[11] Zihe Wang, Zhide Wei, and Jie Zhang. 2020. Bounded Incentives in Manipulating the Probabilistic Serial Rule. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 2276–2283. https://ojs.aaai.org/index.php/AAAI/article/view/5605