

Relaxations of Envy-Freeness Over Graphs

Extended Abstract

Justin Payan

University of Massachusetts Amherst
Amherst, United States
jpayan@umass.edu

Rik Sengupta

University of Massachusetts Amherst
Amherst, United States
rsengupta@umass.edu

Vignesh Viswanathan

University of Massachusetts Amherst
Amherst, United States
vviswanathan@umass.edu

ABSTRACT

In allocating a set of indivisible items among agents, the condition of *envy-freeness* cannot always be achieved. *Envy-freeness up to any good* (EFX) and *envy-freeness with k hidden items* (HEF- k) are two compelling relaxations of envy-freeness, which remain elusive in many settings. We study a natural relaxation of these two fairness constraints, where we place the agents on the vertices of a graph, and only require that our allocations satisfy the EFX (resp. HEF) constraint on the edges of the graph. We refer to these allocations as *graph-EFX* (resp. *graph-HEF*) or simply *G-EFX* (resp. *G-HEF*) allocations. We show that, for any graph G , there always exists a *G-HEF- k* allocation of goods, where k is the size of a minimum vertex cover of G , and this is tight. We show that *G-EFX* allocations of goods exist for three different classes of graphs — two of them generalizing the star $K_{1,n-1}$ and the third generalizing the three-edge path P_4 . Many of these results extend to allocations of *chores* as well. Overall, we show several natural settings in which the graph structure helps obtain strong fairness guarantees. Finally, we devise an algorithm tested using Spliddit data, to show that *G-EFX* allocations appear to exist for paths P_n , pointing the way towards generalizing our results to even broader families of graphs.

KEYWORDS

Fair Allocation; Envy; Local Fairness

ACM Reference Format:

Justin Payan, Rik Sengupta, and Vignesh Viswanathan. 2023. Relaxations of Envy-Freeness Over Graphs: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The problem of fairly allocating a set of indivisible goods among agents with preferences has been extensively studied by the multi-agent systems community [8].

Several notions of fairness have been proposed and analyzed in the last two decades; of all these notions, arguably the most compelling one is that of *envy-freeness*. In an envy-free allocation of goods, no agent prefers the set of goods allocated to any other agent over their own. Unfortunately, with indivisible goods, an envy-free allocation is not guaranteed to exist: consider an example with two agents and one indivisible good.

Several natural relaxations of envy-freeness have been proposed, most notably *envy-freeness up to any good* (EFX) [1], and *envy-freeness up to k hidden goods* (HEF- k) [3]. Both notions have many open questions, and in fact, the existence of EFX or uHEF- k allocations are not known beyond some special cases. The existence of EFX remains one of the biggest open questions in this subfield.

An allocation is EFX if whenever an agent envies another agent, the envy can be eliminated by removing *any* item from the other agent’s allocated bundle. An allocation is (uniformly) HEF- k (or uHEF- k) if there are k or fewer agents who can each “hide” a single good from their bundle (so that they themselves can see it but all other agents are unaware of it) and the resulting allocation is envy-free. Note that these two notions are incomparable; there are allocations that are EFX but not uHEF- $(n - 1)$, and ones that are uHEF-1 but not EFX.

We introduce a relaxation of these fairness criteria, where agents are represented by vertices on a fixed undirected graph, and allocations only need to satisfy the relaxed envy constraint for all neighboring pairs of agents in the graph. For hidden envy-freeness, this amounts to an agent needing to hide a good in order to eliminate the envy from its neighbors. For EFX, this amounts to only needing to satisfy the EFX criterion along the graph edges. These reduce to the usual notions of uHEF- k or EFX allocations when the underlying graph is complete.

In addition to being a generalization of both these fairness constraints, this model is also quite natural, as it captures envy under partial information. In the real world, agents typically do not envy other agents whose allocated bundles they are unaware of. In these cases, it suffices to only consider pairs of agents who are aware of each other and therefore know only each other’s allocated bundles.

Throughout this extended abstract, we will assume there is a set of agents $N = \{1, 2, \dots, n\}$ and a set of goods $M = \{g_1, g_2, \dots, g_m\}$. We assume that the agents N are placed on the vertices of a graph $G = (N, E)$. We use $i \in N$ interchangeably to denote an agent or the corresponding vertex of G . Each agent i has an additive valuation function $v_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ such that for every bundle of goods $S \subseteq M$, we have $v_i(S) = \sum_{g \in S} v_i(\{g\})$. Some of our results extend beyond additive valuations, but for this abstract we restrict ourselves to additivity. An allocation $\mathbf{X} = (X_1, X_2, \dots, X_n)$ is an n -partition of M ; X_i denotes the set of goods (“bundle”) allocated to agent i .

For detailed proofs of all our results, we refer the reader to [6].

2 HIDDEN ENVY-FREENESS OVER GRAPHS

An allocation \mathbf{X} is said to be *G-uHEF- k* if there exists a set of goods $S \subseteq M$ with $|S| \leq k$, such that (a) for every $(i, j) \in E$, we have $v_i(X_i) \geq v_i(X_j \setminus S)$, and (b) for every $i \in N$, $|X_i \cap S| \leq 1$. Intuitively,

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

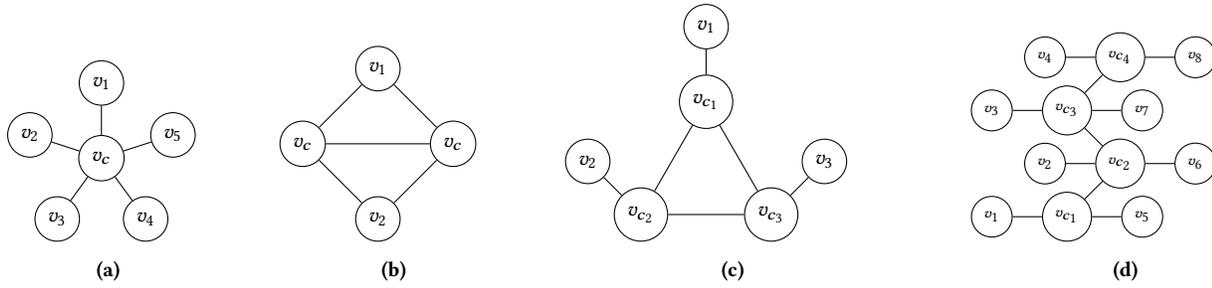


Figure 1: Generalizations of the star. The valuation functions $v_{c_1}, v_{c_2}, v_{c_3}$ and v_{c_4} are identically ordered.

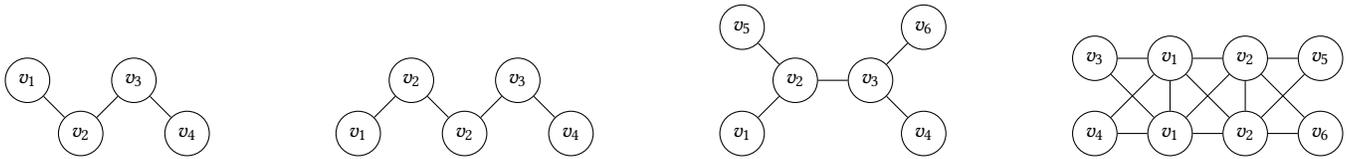


Figure 2: Generalizations of the three-edge path.

an allocation is G -uHEF- k if at most k agents can each hide a good so that no agent envies their neighbors in G .

We devise a simple round-robin based algorithm that computes a G -uHEF- $|C|$ allocation for any vertex cover C of the graph G . The algorithm starts with all goods unallocated; in every round, each agent chooses a single unallocated good, with agents in C choosing before agents outside C . The procedure terminates only when all goods are allocated. We prove that this simple protocol yields a G -uHEF- $|C|$ allocation.

We show further that this is, in fact, optimal. Formally, we prove that every graph G with a minimum vertex cover of size k admits an instance with no G -uHEF- k' allocation for any $k' < k$.

3 EFX OVER GRAPHS

An allocation X is said to be G -EFX if for every edge $(i, j) \in E$, and for every good $g \in X_j$, we have $v_i(X_i) \geq v_i(X_j - g)$. Intuitively, an allocation is G -EFX if the usual EFX criterion is maintained over every edge of the graph G .

In this setting, we start by proving that a G -EFX allocation always exists on any instance when the underlying graph G is a star $K_{1,n-1}$. The proof for this is a straightforward application of a result from [7], which states that an EFX allocation always exists when agents have identical valuations. We can use this to compute an allocation X where all n agents have the same valuation function as the center of the star. We can now reassign these n bundles to the star $K_{1,n-1}$ as follows: we first let the outer $n - 1$ agents take their preferred bundle from X in any order, and allocate the final bundle to the center. The outer agents do not envy the center since they picked their bundles first, and the center does not envy any of the spokes beyond EFX due to our choice of the allocation X .

This technique and result can be generalized to a larger class of graphs which contain a central *core* set of agents with identical valuations and a set of unrestricted *outer* agents with no edges between themselves but any number of neighbors among the core

agents. We also present a second generalization where the core agents only need to have *identically ordered* valuations. Examples of these graphs are shown in Figure 1.

We next show that a G -EFX allocation always exists when $G = P_4$. The proof is similar, but uses the result in [5] that an EFX allocation always exists when the agents only have two different types of valuations. We can again generalize this result to graphs with a core set of agents with two types of valuations, and outer agents with no edges between themselves but arbitrary neighbors among core agents of any one type. Examples are shown in Figure 2.

We also investigate settings with both goods (positively-valued items) and chores (negatively-valued items). In this setting, it is known that EFX allocations do not exist even when agents have *lexicographic* valuations [4]. We show, however, that G -EFX allocations exist under arbitrary lexicographic valuations for all graphs with diameter at least 4.

Empirically, we evaluate a simple algorithm on real fair allocation data (from Spliddit [2]) and show that G -EFX allocations are likely to exist when G corresponds to the path P_n . We use observations from our experiments to discuss several possible potential functions that can be used to prove the existence of EFX on more general classes of graphs. We refer to [6] for further details.

4 CONCLUSION

The graph-based relaxation of our two fairness notions is natural, as many real-life agents only care about the agents with whom they interact. In many cases, this graph setting enables us to obtain results that have not been possible in general. Our hope is that obtaining positive results on natural classes of graphs may help prove the existence of EFX and HEF allocations more broadly. Our empirical results show that G -EFX allocations likely exist for quite general classes of graphs such as paths. Several other fairness notions like local proportionality and local max-min share can be naturally defined on graphs; these offer scope for future research.

ACKNOWLEDGMENTS

We would like to thank Raghavendra Addanki for participating in early discussions of this work. We would also like to thank Rohit Vaish for pointing us to several interesting concepts in fair allocations, and suggesting approaches towards some of the proof techniques we used. Finally, we would like to thank Yair Zick, Hadi Hosseini, and Andrew McGregor for incredibly helpful comments and support. RS acknowledges support from NSF grant no. CCF1908849.

REFERENCES

- [1] 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 (2019), 1–32.
- [2] Jonathan Goldman and Ariel D. Procaccia. 2015. Spliddit: Unleashing Fair Division Algorithms. *SIGecom Exchanges* 13 (2015), 41–46.
- [3] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Hejun Wang, and Lirong Xia. 2020. Fair Division Through Information Withholding. 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, 2014–2021. <https://ojs.aaai.org/index.php/AAAI/article/view/5573>
- [4] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. 2022. Fairly Dividing Mixtures of Goods and Chores under Lexicographic Preferences. <https://doi.org/10.48550/ARXIV.2203.07279>
- [5] Ryoga Mahara. 2021. Extension of Additive Valuations to General Valuations on the Existence of EFX. In *Proceedings of the 29th Annual European Symposium on Algorithms, (ESA)*, Vol. 204. 1–15.
- [6] Justin Payan, Rik Sengupta, and Vignesh Viswanathan. 2022. Relaxations of Envy-Freeness Over Graphs. <https://doi.org/10.48550/ARXIV.2202.10946>
- [7] Benjamin Plaut and Tim Roughgarden. 2020. Almost Envy-Freeness with General Valuations. *SIAM Journal of Discrete Mathematics* 34 (2020), 1039–1068.
- [8] Toby Walsh. 2021. Fair Division: The Computer Scientist’s Perspective. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*.