
SCIENTIFIC REPORTS

Algorithmics of Fair Division and Social Choice

25 Nov 2024–13 Dec 2024

Organizing Committee

Xiaohui Bei

Nanyang Technological University

Edith Elkind

University of Oxford

Warut Suksompong

National University of Singapore

CONTENTS PAGE

		Page
Georgios Amanatidis Athens University of Economics and Business, Greece	Pushing the Frontier on Approximate EFX Allocations	5
Vittorio Bilò University of Salento, Italy	Achieving Envy-Freeness through Items Sale	9
Martin Bullinger University of Oxford, UK	Stability in Random Hedonic Games	12
Diptarka Chakraborty National University of Singapore, Singapore	Fair Rank Aggregation	16
Jiehua Chen Technische Universität Wien, Austria	Fairness in Assignments with Congestion-Averse Agents: Concepts, Algorithms, and Complexity	18
Jiarui Gan University of Oxford, UK	Envy-free policy teaching to multiple agents	19
Jugal Garg University of Illinois Urbana-Champaign, USA	Fair Allocation of Indivisible Chores	21
Matthias Greger Technical University of Munich, Germany	Preference Aggregation on the Probability Simplex	22
Daniel Halpern Harvard University, USA	Aggregating Preferences with Limited Queries	24
Alexandros Hollender University of Oxford, UK	Envy-Free Cake-Cutting for Four Agents	27
Naoyuki Kamiyama Kyushu University, Japan	Some Recent Results on Super-stable Matchings	30
Yasushi Kawase University of Tokyo, Japan	The Random Assignment Problem Under Constraints	32
Alexander Lam The Hong Kong Polytechnic University, Hong Kong SAR	Facility Location Games with Scaling Effects	36
Jérôme Lang Université Paris Dauphine, France	How fair can strategy-proof fair division be?	40
Jérôme Lang Université Paris Dauphine, France	Irrelevant alternatives are relevant	43

		Page
Bo Li Hong Kong Polytechnic University, Hong Kong SAR	MMS Allocation of Indivisible Chores with Subadditive Valuations and the Fair Surveillance Assignment Problem	47
Xinhang Lu University of New South Wales, Australia	Best-of-Both-Worlds Fair Allocation of Indivisible and Mixed Goods	49
Ryoga Mahara The University of Tokyo, Japan	Proportionality up to the Least Valued Good on Average	53
David Manlove University of Glasgow, UK	Couples can be Tractable: New Algorithms and Hardness Results for the Hospitals / Residents Problem with Couples	54
Pasin Manurangsi Google Research, Thailand	Fair Division for Random Utilities	58
Ulrike Schmidt-Kraepelin TU Eindhoven, Netherlands	Monotone Randomized Apportionment	61
Erel Segal-Halevi Ariel University, Israel	Fairness in Real-estate Division	66
Erel Segal-Halevi Ariel University, Israel	Reducing Leximin Fairness to Utilitarian Optimization	69
Hanna Sumita Tokyo Institute of Technology, Japan	Fair Allocation with Binary Valuations for Mixed Divisible and Indivisible Goods via Hybrid Convex Optimization	72
Zhaohong Sun Kyushu University, Japan	Stable Matching in Practice: Daycare Matching Markets in Japan	75
Biaoshuai Tao Shanghai Jiao Tong University, China	Truthful and Almost Envy-Free Mechanism of Allocating Indivisible Goods: the Power of Randomness	79
Nicholas Teh University of Oxford, UK	Temporal Fair Division of Indivisible Items	83
Rohit Vaish Indian Institute of Technology Delhi, India	Fair Interval Scheduling of Indivisible Chores	87
Adrian Vetta McGill University, Canada	Six Candidates Suffice to Win a Voter Majority	89
Xiaowei Wu University of Macau, Macau	Fair Allocation of Chores with Subsidy	93
Yu Yokoi Tokyo Institute of Technology, Japan	Popular Matching under Matroid and Optimality Constraints	97

	Page
Sherman Yuen National University of Singapore, Singapore	Attaining EF1 Allocations by Exchanging Goods 99

PUSHING THE FRONTIER ON APPROXIMATE EFX ALLOCATIONS

GEORGIOS AMANATIDIS

Classification AMS 2020: 68W25, 68W40, 91B32

Keywords: fair division, envy-freeness up to any good, approximate EFX

This talk was based on joint work with Aris Filos-Ratsikas and Alkmini Sgouritsa [3].

General. The existence of EFX allocations is an important open problem in fair division with indivisible goods. In this setting, a set of agents have values over a set of indivisible goods, and the goal is to allocate the goods to the agents in a way that is perceived as *fair* by everyone. One of the most well-established notions of fairness is *envy-freeness*, introduced by Gamow and Stern [13] in the context of *divisible* resources; this notion stipulates that no agent would prefer another agent's allocation to her own. For indivisible goods, it is not hard to see that envy-free allocations may not be possible. Motivated by this impossibility, the literature has defined *relaxed* fairness notions, appropriate for indivisible goods allocation. Budish [5] defined *envy-freeness up to one good (EF1)*, which deems an allocation fair if the envy of an agent is eliminated after the removal of *some* good from the bundle of another envied agent. EF1 was implicitly introduced by Lipton et al. [15] who showed that EF1 allocations for monotone valuation functions always exist and can be found efficiently. However, in certain applications EF1 might be unsatisfactory, as it might require the removal of a very valuable good to restore envy-freeness. To address this, Gourvès et al. [14] and Caragiannis et al. [7] introduced *envy-freeness up to any good (EFX)*, which stipulates that the envy is eliminated even if the least valuable good, from the envious agent's perspective, is removed from the envied agent's bundle.

Contrary to EF1 allocations, the existence of EFX allocations is much more intricate. As we mentioned, this is a major open problem, carrying great momentum and being met with intensive efforts from the research community. The associated research has adopted a systematic approach to tackling this challenging question, by first obtaining existence results for special cases of the problem, developing a deeper understanding of its intricacies, and ultimately aiming to synthesize these ideas into an answer to the main problem. Three notable such results are that EFX allocations exist when there are at most three agents [10], or each agent's value for each good can be one of *two* numbers a or b [2], or the agents' valuation functions can be represented by a *graph*, with edges corresponding to goods and nodes corresponding to agents [11]. Other interesting restrictions that have been studied include agents with values that induce the same ordering over goods [17] and leaving some goods unallocated [6, 9, 4].

A related line of work has studied *approximations* to the EFX notion. An allocation is α -EFX if after the removal of any good from the envied agent's bundle, the envy is bounded by a factor of α . The state of the art for approximate EFX allocations is a $\phi - 1 = 0.618$ approximation due to Amanatidis et al. [1]. Markakis and Santorinaios [16] were able to produce $2/3$ -EFX allocations when all of the n agents agree on which n

goods are the most valuable. Whether $2/3$ -EFX allocations can be achieved in general is still open. If sufficiently many goods are left unallocated, near optimal approximations to EFX can be achieved [8].

The known results can thus be seen as lying on a certain kind of *frontier*: indeed, a certain set of parameters (e.g., number of agents, type of values, approximation ratio), can be seen as a point on a search space, with those points for which we have obtained existence results constituting the frontier of our current understanding of the problem. The ultimate goal is to move towards the point corresponding to (exact) EFX for any number of agents and without restrictions on the values or the structure.

Our Results and Techniques. In this work, we prove that $2/3$ -EFX allocations exist and can be efficiently computed for agents with additive valuation functions in three important cases, namely when:

- There are at most *seven* agents.
- Each agent’s value for each good can be one of *three* non-negative numbers a , b , or c .
- The agents’ values can be represented by a *multigraph*, with edges corresponding to goods and nodes corresponding to agents. Here an agent has nonzero value for a good only if this good is incident to her. This setting generalizes the setting studied recently by Christodoulou et al. [11].

We make progress in all three settings through the same algorithmic framework, although each one requires its own modifications. While all of these are nontrivial, the most intricate case is, somewhat surprisingly, the case of three values. We next present an overview of our techniques.

Property-Preserving Partial Allocations. Our approach is based on the following general principle: obtain a *partial* allocation \mathbf{X} of goods to agents that satisfies a certain set of properties. Then this allocation can be transformed into a complete allocation that is $2/3$ -EFX. To be more precise, all of the known algorithms for α -EFX in the literature [1, 16, 12] start by producing a partial allocation \mathbf{X} using only a subset of the goods. How this partial allocation is obtained may differ between different algorithms, but they all serve the same purpose: once \mathbf{X} is obtained, then one can run the Envy Cycle Elimination algorithm of Lipton et al. [15], with initial input \mathbf{X} , to produce an allocation that is approximately EFX. For this to be possible, \mathbf{X} has to satisfy certain properties, mainly that (a) it is α -EFX, for the approximation factor α that we are aiming to prove, and (b) none of the agents consider any of the goods that are left unallocated *too valuable*. Formally, an unallocated good is “too valuable” for an agent i if her value for that good is at least a factor β of her value for her allocated bundle in the partial allocation \mathbf{X} . These goods lie at the heart of our approach and we refer to them as *critical goods*. This is captured by a lemma of Markakis and Santorinaios [16] stating that *if a partial allocation \mathbf{X} is α -EFX and does not induce any critical goods (defined via a parameter β), then it can be transformed into a complete $\min\{\alpha, \frac{1}{\beta+1}\}$ -EFX allocation.*

Therefore, the value of β that makes a good critical depends on the approximation factor α that we are aiming to prove existence for. For $\alpha = 0.618$ (achieved by Amanatidis et al. [1]), it is also the case that $\beta = 0.618$. From a technical perspective, the “balance” of these terms makes the construction of a partial allocation \mathbf{X} that does not induce any critical goods achievable via relatively simple algorithms. In fact, said algorithms also guarantee that the cardinality of each bundle in \mathbf{X} is at most 2; working with bundles

of such size is much more manageable. For larger α however, $\alpha \neq \beta$ and we will have a natural imbalance. For $\alpha = 2/3$ in particular, we have $\beta = 1/2$. In this case, it can be shown that even for the cases that we consider, it is not possible for \mathbf{X} to be both $2/3$ -EFX and induce no critical goods, unless agents receive bundles of cardinality 3 or more.

Our goal will be to obtain a $2/3$ -EFX partial allocation \mathbf{X} without critical goods *in two stages*. First, we devise a general algorithm called PROPERTY-PRESERVING PARTIAL ALLOCATION algorithm (3PA), which obtains a partial allocation \mathbf{X}^1 that satisfies a certain set of properties. One such property is that it is $2/3$ -EFX. This partial allocation \mathbf{X}^1 still has critical goods, but it limits their number to at most *one per agent*, and only for agents that have singleton bundles in \mathbf{X}^1 . This, together with the other properties of \mathbf{X}^1 will prepare the ground for allocating the critical goods to the agents in a subsequent stage, resulting in a new partial allocation \mathbf{X}^2 . This \mathbf{X}^2 will now satisfy the aforementioned lemma, and, thus, can be transformed into a complete $2/3$ -EFX allocation.

Swap steps and various envy graphs. The 3PA algorithm is based on a series of steps which are executed in sequence according to a certain priority structure. Most of these are *swap steps*, i.e., steps that enable certain agents to exchange (parts of) their bundles in the partial allocation \mathbf{X} with certain unallocated goods. The priority is determined by the cardinality of the bundles in \mathbf{X} , as well as the value of the agents for those unallocated goods. Besides the swap steps, the algorithm also includes steps that are performed on several different types of envy graphs, associated with allocation \mathbf{X} . A (standard) envy graph for a partial allocation \mathbf{X} is a graph in which the nodes corresponds to agents, and an edge (i, j) signifies that agent i envies agent j . Envy graphs are very common in fair division, starting with the Envy Cycle Elimination algorithm [15]. We also consider different types of graphs, named *reduced* graphs and *enhanced* graphs. In the former, any envy towards agents with singleton bundles is disregarded, unless it is high enough. In the latter, we add edges of *near* envy, as long as the value of the target bundle is above a certain threshold. By exchanging bundles along cycles and paths in these graphs, we deal with the inherently most challenging case of “moving value” from and to agents with singleton bundles. In the case of three values we further refine our graphs, including edges that indicate ties in the values. This allows for more options on how the partial allocations evolve, but adds an extra layer of complexity in the analysis.

Throughout the execution of the 3PA algorithm, the value of an agent for her bundle may decrease several times, but this is done in a controlled way that allows us to allocate the critical goods in subsequent steps. We remark that since the steps of the 3PA algorithm and its variants repeatedly allocate and de-allocate goods, their polynomial running time, or even their termination is far from obvious.

REFERENCES

- [1] Georgios Amanatidis, Evangelos Markakis, and Apostolos Ntokos. Multiple birds with one stone: Beating $1/2$ for EFX and GMMS via envy cycle elimination. *Theoretical Computer Science*, 841:94–109, 2020.
- [2] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863:69–85, 2021.

- [3] Georgios Amanatidis, Aris Filos-Ratsikas, and Alkmini Sgouritsa. Pushing the frontier on approximate EFX allocations. *CoRR*, abs/2406.12413, 2024. doi: 10.48550/ARXIV.2406.12413. URL <https://doi.org/10.48550/arXiv.2406.12413>.
- [4] Ben Berger, Avi Cohen, Michal Feldman, and Amos Fiat. Almost full EFX exists for four agents. In *Proceedings of the 36th AAI Conference on Artificial Intelligence (AAAI)*, pages 4826–4833, 2022.
- [5] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. of Political Economy*, 119(6):1061–1103, 2011.
- [6] Ioannis Caragiannis, Nick Gravin, and Xin Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pages 527–545, 2019.
- [7] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3):1–32, 2019.
- [8] Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, Ruta Mehta, and Pranabendu Misra. Improving EFX guarantees through rainbow cycle number. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 310–311, 2021.
- [9] Bhaskar Ray Chaudhury, Telikepalli Kavitha, Kurt Mehlhorn, and Alkmini Sgouritsa. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4): 1336–1358, 2021.
- [10] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. EFX Exists for Three Agents. *Journal of the ACM*, accepted, 2023.
- [11] George Christodoulou, Amos Fiat, Elias Koutsoupias, and Alkmini Sgouritsa. Fair allocation in graphs. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 2023.
- [12] Alireza Farhadi, MohammadTaghi Hajiaghayi, Mohamad Latifian, Masoud Seddighin, and Hadi Yami. Almost envy-freeness, envy-rank, and Nash social welfare matchings. In *Proceedings of the AAI Conference on Artificial Intelligence*, volume 35, pages 5355–5362, 2021.
- [13] George Gamow and Marvin Stern. *Puzzle-math*. Viking Press, 1958.
- [14] Laurent Gourvès, Jérôme Monnot, and Lydia Tlilane. Near fairness in matroids. In *21st European Conference on Artificial Intelligence (ECAI 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 393–398. IOS Press, 2014.
- [15] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.
- [16] Evangelos Markakis and Christodoulos Santorinaios. Improved EFX approximation guarantees under ordinal-based assumptions. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 591–599, 2023.
- [17] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2):1039–1068, 2020.

ACHIEVING ENVY-FREENESS THROUGH ITEMS SALE

VITTORIO BILO`

Classification AMS 2020: 68Q17, 68Q25, 91B14.

Keywords: Fair division, Envy-freeness, Efficient algorithms, Hardness of approximation.

Fair division refers to the algorithmic question, dating back to the origins of the civil society, of allocating resources or tasks to a set of agents according to some justice criteria. It is by now a prominent area within Computational Social Choice, [3, Part II]. One of the most natural and well studied notions of fairness is *envy-freeness* [5]: a division is envy-free if everyone thinks that her share is at least as valuable as the share of any other agent. When items are indivisible, obtaining an envy-free allocation is very challenging [4], and it is well known that, in the majority of cases, envy-free divisions do not exist.

An approach that has been followed by several works, in order to recover some existential guarantees, is to focus on relaxations of envy-freeness. Another natural direction that comes into mind is to insist on envy-freeness, but provide some compensation (e.g., monetary) to the agents who may feel unhappy by a proposed division. Such models have been considered in the literature, where money is either coming as an external subsidy from a third party or is already part of the initial endowment. Under this setting, [6] investigated the question of determining the minimum amount of money needed to obtain an envy-free division.

In this work, we also allow for monetary rewards, but we choose a different approach, as already initiated in [7]: we require that the money used to compensate the envious agents has to be raised from the set of available items, by selling some of them. This is what happens, for instance, in inheritance division. To provide some examples, as stated in Article n. 9 of the New York Laws - Real Property Actions and Article n. 720 of the Italian Civil Code, whenever an agreement is not possible, part of the inheritance can be sold through an auction. The same practice is also used in divorce settlements. Clearly, envy-freeness is then always feasible by selling, if needed, the whole inheritance, and equally sharing the proceeds. However, the amount of money raised by this process can be fairly below the real value of the sold items for at least two reasons. First, the bidders who participate in this type of auctions usually aim at winning items at very low prices; secondly, running an auction bears organizational costs which need to be subtracted from the proceeds. Thus, it is in the interest of the heirs to determine an envy-free division by selling assets with as little value loss as possible. This gives rise to an interesting optimization problem of determining which items to sell, so as to arrive at an envy-free allocation with optimal social welfare. Algorithmically, this question has been largely unexplored, with the exception of a particular case handled in [7].

Assuming that we are given the *market value* of each item as input, i.e., the money that can be raised by selling it, we embark on a thorough investigation of algorithmic and complexity questions for our problem and provide an almost tight set of results.

We start with the case where all agents have the same value for each item. After establishing NP-hardness, which can be easily shown even for 2 agents, our main results exhibit a sharp separation on the approximability between the cases of $n = 2$ and $n \geq 3$ agents. In particular, we prove that, with at least three agents, no polynomial time algorithm can obtain a solution that performs better than the one which sells all items, unless $P = NP$. On the other hand, for two agents, we are able to design a polynomial time approximation scheme (PTAS), under the assumption that the market value of each item is not smaller than half of the common agents' value. The idea behind the PTAS is to enumerate all partial allocations of the most valuable items, whose number is a constant depending on the desired approximation guarantee. Each such partial allocation, which consists of the two bundles assigned to the agents together with the bundle of sold items, is then completed processing the remaining items by non-increasing value. At every step, the next item is allocated to the agent having the lower valued bundle, until we reach a situation where it is possible to equalize the two bundles by using the money raised from the already sold items and from selling a subset of the not-yet-processed ones. The main technical effort is needed to show that, if this condition occurs, then the final allocation can be made envy-free at the expense of a negligible loss of social welfare, while, if the condition never occurs, then it is not possible to obtain an envy-free solution from the starting partial allocation. Finally, we design a dynamic-programming algorithm which runs in polynomial time when the number of distinct item values is constant; this assumption is in line with several other recent works on fair-division, e.g., [1, 2].

We then move to the case in which agents can have heterogeneous values for each item. While all computational barriers from the case of equal valuations carry over to this case as well, we are able to obtain two additional positive results. First, we focus on the setting where the value that an agent i has for any item lies in an interval of the form $[x_i, \beta x_i]$, where β is common across all agents. This means, essentially, that each agent attributes the same value to all items, up to a factor of β . For a constant number of agents, and for a constant value of β , we are able to design again a PTAS. This is very different from the PTAS for identical valuations and is based on an appropriate combination of two main ideas. First, by using a linear programming formulation, we compute a fractional solution with a bounded number of fractionally assigned items. Then, we apply a "reverse" version of the envy cycle elimination algorithm [8], so as to decide which items to sell, in addition to the fractional ones. We believe that this could be of independent interest for other allocation problems as well. Finally, even if we drop the assumption on β being constant, we can still provide a pseudo-polynomial time algorithm.

REFERENCES

- [1] H. Akrami, B. R. Chaudhury, M. Hoefler, K. Mehlhorn, M. Schmalhofer, G. Shahkarami, G. Varricchio, Q. Vermande and E. van Wijland. Maximizing Nash social welfare in 2-value instances. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence, (AAAI)*, pp. 4760–4767, 2022.
- [2] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, A. Hollender and A. A. Voudouris. Maximum Nash welfare and other stories about EFX. *Theor. Comput. Sci.*, 863:69–85, 2021.
- [3] F. Brandt, V. Conitzer, U. Endriss, J. Lang and A. D. Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.
- [4] D. Foley. Resource allocation and the public sector. *Yale Econ. Essays*, 7(1):45–98, 1967.
- [5] G. Gamow and M. Stern. *Puzzle-math*. Viking Press, 1958.

- [6] D. Halpern and N. Shah. Fair division with subsidy. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, LNCS 11801, Springer, pp. 374–389, 2019.
- [7] J. Karp, A. M. Kazachkov and A. D. Procaccia. Envy-free division of sellable goods. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, (AAAI)*, pp. 728–734, 2014.
- [8] R. J. Lipton, E. Markakis, E. Mossel and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pp. 125–131, 2004.

DEPARTMENT OF MATHEMATICS AND PHYSICS, UNIVERSITY OF SALENTO, ITALY
Email address: vittorio.bilo@unisalento.it

STABILITY IN RANDOM HEDONIC GAMES

MARTIN BULLINGER

Classification AMS 2020: 91A68, 68W40, 68Q87

Keywords: Algorithmic game theory, hedonic games, stability, random games

In this talk, I demonstrate that stable outcomes are highly likely to exist in random additively separable hedonic games and can be efficiently identified using appropriate algorithms. The talk is based on a joint paper with Sonja Kraiczy [4].

1. HEDONIC GAMES

A frequent challenge is to partition a set of agents into meaningful groups. For instance, this might involve assigning employees to company departments or organizing students into project teams. However, even more abstract applications with nonhuman agents are possible, e.g., when classifying a large body of webpages. Such settings can be conceptualized as coalition formation games and they are studied across many fields in computer science, mathematics, and economics [5, 7].

In many scenarios, especially when agents are assumed to be humans, preferences are based on which group an agent is part of, whereas the composition of other groups does not matter for an agent's own utility. This assumption has led to the introduction of *hedonic games* as a model of coalition formation over four decades ago [6]. The algorithmic task of coalition formation is then to produce a partition of a set of agents, given their preferences over possible coalitions. A first elementary question, therefore, is how to elicit agents' preferences. In principle, one could ask agents to provide a ranking of all possible coalitions they could be part of. However, this would require exponentially large space with respect to the number of agents, and is thus infeasible, even for moderately small instances. A lot of the research on hedonic games has, therefore, addressed the question of how to succinctly express preferences, see, e.g., [1, 2]. A common approach is to assume that preferences can be encoded as a complete and weighted directed graph, where the vertices represent the agents and the weight of an edge from agent a to agent b represents the valuation of a for b . These valuations can then be aggregated to utilities over coalitions. A natural and widely studied model is to assume that the utility of a for a coalition C is the sum of edges from a to other agents in C . Games that can be represented like this are called *additively separable hedonic games* [2].

2. STABILITY

We now assume that we have given a hedonic game (in some representation) and want to specify desirable outcomes. A key concept in this context are *stable* partitions, in which no agent would prefer joining another coalition (or forming a coalition on their own) over staying in their assigned coalition.

Let us consider a simple scenario that is illustrated in Figure 1. Assume that there are two agents, a mouse and a cat. For two agents, there are only two possible coalitions for each agent: the singleton coalition, where they are on their own, and the grand coalition, where they form a coalition together with the other agent. In our game, the preferences might be as follows: The mouse is scared of the cat and prefers to be in the singleton coalition, whereas the cat is hungry and prefers the grand coalition. These preferences can be captured by an additively separable hedonic game. A possible representation is provided in Figure 1.




FIGURE 1. Additively separable hedonic game representing the run-and-chase scenario.

What would constitute a stable outcome in this scenario? In the singleton partition, where both agents form singleton coalitions, the cat would prefer joining the mouse. Indeed, this would increase her (additively separable) utility from 0 (the utility for the singleton coalition) to 10. If, however, the grand coalition was proposed, then the mouse would prefer to leave and be on her own. Again, this would lead to an increase in utility, this time from -10 to 0. Intuitively, there is no stable outcome. We are caught in an endless run-and-chase loop.

The stability concept considered in the example requires that no agent can unilaterally decide to abandon their coalition and join another coalition. An outcome in which no such deviation leads to an improvement is called *Nash-stable* [2]. As we have seen, Nash-stable outcomes are not guaranteed to exist, even in simple settings. Arguably, the unilateral nature of Nash stability is demanding and can even be inappropriate in a coalition formation setting.

In some cases, deviating may require the consent of other agents. For example, coalition membership might be governed by contracts, allowing an agent to leave only if no member of the coalition exercises a veto. This concept is known as *contractual Nash stability* [8]. In the previous example, the grand coalition remains contractually Nash-stable because the mouse would be prevented from making its beneficial deviation.

Instead of requiring consent to abandon a coalition, an agent might also need permission to join one. This leads to the concept of *individual stability* [2]. In our example, the singleton partition is individually stable. Indeed, the mouse would veto the cat's attempt to join her.

3. COMPUTATIONAL COMPLEXITY

The discussion of the example in Figure 1 looks like it is easier to construct contractually Nash-stable or individually stable partitions. And indeed, both are weakenings of Nash stability that may exist when Nash stability is not satisfiable. However, their existence is not guaranteed in general either. There are simple examples with a small number of agents, in which no contractually Nash-stable (see [8],

Example 2) or individually stable (see [2], Example 5) partition exists. This gives rise to the following decision problem: given an additively separable hedonic game, does there exist a stable partition? For all three of our stability concepts, this problem is NP-complete [3, 8]. In the case of Nash stability, this hardness result holds even for games where valuations are restricted to one positive and one negative value [3].

4. RESULTS

How significant are the hardness results discussed in the last section? As our example along Figure 1 suggests, Nash stability is a volatile concept and there are many instances without a Nash-stable outcome. However, the additively separable hedonic games constructed in all three hardness results are combinatorially involved and rely on a careful choice of valuations. By contrast, assume that games are based on randomly selected valuation graphs. Our central question is whether this helps to achieve stability.

How likely is it for stable partitions to exist in random games?

It is easy to see that the consideration of random games can lead to stability: Assume that valuations are drawn independently from some distribution. If the probability of a positive valuation is 0, then all valuations are nonpositive with probability 1. In this case, the singleton partition is Nash-stable (and, therefore, also contractually Nash-stable). If, however, the probability of a positive valuation is positive, then every agent will be liked by some other agent with high probability for a large number of agents. Hence, the grand coalition is contractually Nash-stable with probability 1 in the large agent limit. We capture this in the following proposition.

Proposition 1 ([4]). *In the large agent limit, a contractually Nash-stable partition exists with probability 1 in random additively separable hedonic games.*

It seems that contractual Nash stability is easy to achieve. However, the partitions that lead to Theorem 1 are the “trivial” ones. Clearly, if there are no positive valuations, then the singleton partition is the only sensible choice: it is most preferred by every agent, and there is no incentive to form nonsingleton coalitions at all. However, even in the presence of positive weights, the grand coalition can be highly undesirable: first, all agents might still obtain a negative utility, and the coalition is only rendered stable due to the vetoes of unhappy agents. In other words, a partition is not required to be *individually rational*, defined as yielding a utility of at least 0 for each agent. Second, the grand coalition seems unreasonable in a context with many agents. By contrast, smaller coalitions might be more reasonable or desirable.

Interestingly, since they do not require consent to abandon a coalition, individual stability and Nash stability entail individual rationality. Hence, they seem to be the harder objectives to satisfy. With high probability, they are neither satisfied by the singleton partition nor the grand coalition.

The first main contribution of our paper is an algorithm that efficiently finds partitions with the following properties with high probability:

- (1) Individually stable,
- (2) Contractually Nash-stable, and
- (3) Coalition sizes of $\Theta(\log n)$ where n is the number of agents.

This algorithm holds for the natural special case, where the random valuations are drawn from a uniform distribution over the interval $[-1, 1]$. We refer to such games as *uniformly random* additively separable hedonic games.

Theorem 2 ([4]). *There exists a polynomial-time algorithm that, in the large agent limit, produces individually stable and contractually Nash-stable partitions consisting of coalitions of size $\mathcal{O}(\log n)$ with probability 1 for uniformly random additively separable hedonic games.*

Notably, the produced partitions simultaneously satisfy both of the discussed weakenings of Nash stability. This raises the question of whether Nash-stable partitions also exist with high probability. Interestingly, this is not the case, as we prove in our second main contribution.

Theorem 3 ([4]). *In the large agent limit, uniformly random additively separable hedonic games do not admit Nash-stable partitions with probability 1.*

Theorems 2 and 3 present a strict separation of Nash stability and weaker stability concepts. We interpret them as follows: In a coalition formation context, where the agents affected by a deviation are involved in determining its permissibility, stability is likely to exist. By contrast, agents acting single-handedly are usually to blame for instabilities.

REFERENCES

- [1] Coralio Ballester. NP-completeness in hedonic games. *Games and Economic Behavior*, 49(1):1–30, 2004.
- [2] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- [3] Felix Brandt, Martin Bullinger, and Leo Tappe. Stability based on single-agent deviations in additively separable hedonic games. *Artificial Intelligence*, 334:104160, 2024.
- [4] Martin Bullinger and Sonja Kraiczky. Stability in random hedonic games. In *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*, page 212, 2024.
- [5] Martin Bullinger, Edith Elkind, and Jörg Rothe. Cooperative game theory. In Jörg Rothe, editor, *Economics and Computation: An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, chapter 3, pages 139–229. Springer, 2024.
- [6] Jacques H. Drèze and Joseph Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica*, 48(4):987–1003, 1980.
- [7] Debraj Ray. *A Game-Theoretic Perspective on Coalition Formation*. Oxford University Press, 2007.
- [8] Shao C. Sung and Dinko Dimitrov. On myopic stability concepts for hedonic games. *Theory and Decision*, 62(1):31–45, 2007.

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OXFORD, UK
Email address: martin.bullinger@cs.ox.ac.uk

FAIR RANK AGGREGATION

DIPTARKA CHAKRABORTY

Classification AMS 2020: 68W25 Approximation algorithms

Keywords: Ranking, Median, Fairness, Kendall-tau distance, Ulam distance

Aggregating multiple input rankings over a group of candidates to derive a consensus ranking is a vital challenge, with numerous applications in areas such as social choice theory, hiring processes, college admissions, web search, and databases. One of the popular objectives is to find a *median* ranking – compute a ranking that minimizes the sum of distances to all the input rankings. Different distance metrics have been considered to study this problem. For example, with the Kendall-tau metric, a polynomial-time approximation scheme (PTAS) is known [2]. For the Spearman-footrule metric, there is an exact polynomial-time algorithm [1]. When considering the Ulam metric, we have developed an algorithm offering a 1.999-approximation [5, 6].

However, achieving the optimal consensus ranking can sometimes result in biases against individual candidates or groups, particularly those from marginalized communities. This concern has prompted investigations into rank aggregation with a focus on fairness. The goal, in addition to generating a consensus ranking, is to ensure equitable representation of each group in the top tiers of the final aggregated ranking.

Proportionate fairness in rank aggregation was first explored by [3, 4] among others. They proposed a basic algorithm that identifies the closest fair ranking for each input ranking and selects the one with the lowest total distance. Using the triangle inequality, it can be shown that this approach achieves a 3-approximation, regardless of the distance function used. However, for metrics like Ulam and Kendall-tau, this 3-factor can be improved with more advanced methods. It's important to note that the approximation factor is influenced by the specific fairness criterion applied. For certain fairness variations, it is possible to achieve an approximation factor approaching 2.

The challenge becomes significantly more complex when it comes to determining a *center* ranking – to compute a ranking that minimizes the maximum distance to input rankings. For all three distance measures, identifying a center ranking (without any fairness constraint) is recognized as an NP-hard problem [7, 8]. Currently, the only widely known approach is a basic 2-approximation algorithm that simply selects one of the input rankings as the output. Exploring the potential for achieving a better approximation guarantee presents an intriguing avenue for further research.

REFERENCES

- [1] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the Tenth International World Wide Web Conference, WWW 10, Hong Kong, China, May 1-5, 2001*, pages 613–622, 2001.
- [2] Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 95–103, 2007.

- [3] Dong Wei, Md Mouinul Islam, Baruch Schieber, and Senjuti Basu Roy. Rank aggregation with proportionate fairness. In *Proceedings of the 2022 international conference on management of data*, pages 262–275, 2022.
- [4] Diptarka Chakraborty, Syamantak Das, Arindam Khan, and Aditya Subramanian. Fair rank aggregation. *Advances in Neural Information Processing Systems*, 35:23965–23978, 2022.
- [5] Diptarka Chakraborty, Debarati Das, and Robert Krauthgamer. Approximating the median under the ulam metric. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 761–775. SIAM, 2021.
- [6] Diptarka Chakraborty, Debarati Das, and Robert Krauthgamer. Clustering permutations: New techniques with streaming applications. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2023.
- [7] Christian Bachmaier, Franz J. Brandenburg, Andreas Gleißner, and Andreas Hofmeier. On the hardness of maximum rank aggregation problems. *Journal of Discrete Algorithms*, 31:2–13, 2015. 24th International Workshop on Combinatorial Algorithms (IWOCA 2013).
- [8] Therese Biedl, Franz J Brandenburg, and Xiaotie Deng. On the complexity of crossings in permutations. *Discrete Mathematics*, 309(7):1813–1823, 2009.

SCHOOL OF COMPUTING,, NATIONAL UNIVERSITY OF SINGAPORE
Email address: diptarka@comp.nus.edu.sg

FAIR CONGESTED ASSIGNMENTS: CONCEPTS, ALGORITHMS, AND COMPLEXITY

JIEHUA CHEN

Classification AMS 2020:

Keywords: Many-to-one matching with sided preferences; congestion games

The congested assignment problem is concerned with assigning agents to posts where agents care about both the posts and their congestion levels. Here, agents are averse to congestion, consistently preferring lower over higher congestion for the same resource. Such scenarios are prevalent across many domains, including traffic management and school choice, where fair resource allocation is crucial. Congested assignment can be considered as a restricted variant of the Group Activity Selection problem, introduced by Darmann et al. [Darmann et al.(2012)]. Additionally, it is related to many-to-one matching in matching under preferences.

In this talk, I will explore one ex-ante fairness concept, top-fairness, and two ex-post fairness concepts, envy-freeness and competitiveness. The top-fairness and competitiveness concepts were recently introduced by Bogomolnaia and Moulin [Bogomolnaia and Moulin(2023)]. While a top-fair or envy-free assignment always exists and can be found easily, competitive assignments do not always exist. The talk will cover the following key points:

- (1) An efficient method to determine the existence of competitive or maximally competitive assignments for a given congestion profile.
- (2) Two optimization variants of congested assignments and their computational complexity: a) Finding a top-fair assignment that is envy-free b) Finding a top-fair assignment that is maximally competitive. Both variants are NP-hard, unfortunately.
- (3) Parameterized algorithms for these NP-hard problems.

REFERENCES

- [Darmann et al.(2012)] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, Gerhard J. Woeginger. 2012 Group Activity Selection Problem. In Proceedings of the 8th Workshop on Internet and Network Economics volume 7695. pages: 156–169.
- [Bogomolnaia and Moulin(2023)] Anna Bogomolnaia and Herve Moulin. 2023. The congested assignment problem. Technical report, arXiv:2301.12163v3.

TU WIEN, AUSTRIA

Email address: jiehua.chen@tuwien.ac.at

ENVY-FREE POLICY TEACHING TO MULTIPLE AGENTS

JARUI GAN

Classification AMS 2020: 91B32, 91A15, 91A65

Keywords: Policy teaching, envyfreeness, Markov decision process

Incentive design is a key approach to influencing the behavior of rational agents. In reinforcement learning (RL), an agent's incentives are defined by its reward function [1]. Thus, one can guide an agent toward a desired policy by modifying its reward function so that the target policy becomes optimal with respect to the adjusted rewards. For example, in safe RL, penalties can be assigned to hazardous actions to discourage an agent from taking them [2]. In many cases, personalized teaching programs are beneficial for heterogeneous agents, who may have vastly different innate reward functions or apply different discount rates. As a result, agents may perceive the same action in the same situation as being rewarded or penalized differently. This raises concerns about fairness, leading to the question: how can we design fair personalized teaching programs that ensure agents perceive their treatment as fair?

Objectives and Results. Our first goal is to define fairness in the context of policy teaching. We adopt the well-established notion of envy-freeness (EF) from fair division theory, which has been applied to resolving disputes over, e.g., property division and rent splitting [3, 4]. In the context of policy teaching, our aim is to design a set of personalized teaching programs such that no agent prefers another agent's program over their own. At the same time, each program must incentivize its corresponding agent to follow the target policy—a fundamental requirement of policy teaching. Beyond the standard EF condition, we consider two stronger variants: one allows agents to evaluate alternative teaching programs by considering deviations from the target policy, while the other mandates identical teaching programs for all agents, ensuring absolute fairness.

We investigate several fundamental questions about EF policy teaching.

- *Existence of an EF Solution.* The first question is about the existence of an EF solution under the EF notions of interest. We show that an EF solution always exists and can be obtained by penalizing undesired actions by a sufficiently large value. Nevertheless, the reverse does not hold true: one cannot hope to find an EF solution only by rewarding actions desired by the target policy. We demonstrate instances that do not admit any EF solution when penalties are not allowed even with the weakest EF notion; we also prove that this non-existence issue is resolved if the agents have the same discount factor.
- *Cost Minimization.* Since reward modification can be costly, we are interested in least-cost EF solutions. We show that computing a cost-minimizing EF solution can be formulated as convex optimization and hence can be solved efficiently.
- *Price of Fairness.* Finally, we analyze the *price of fairness* (PoF), a quantity that measures the (multiplicative) increase of the cost due to consideration of fairness, in the spirit of the *price of anarchy* (PoA) in game theory [5]. We

present tight asymptotic bounds on the PoF. The PoF increases at most quadratically with the geometric sum of the discount factor and linearly with the size of the MDP in general. Additionally, it may also grow linearly with the number of agents involved depending on the specific EF notion considered.

In summary, our results suggest that incorporating fairness into policy teaching can lead to the non-existence of feasible solutions in some cases, but existence is ensured in a broad range of important settings. Fairness does not appear to increase the computational complexity of policy teaching, and the additional cost it incurs grows only moderately with problem size. Overall, our findings indicate that fairness can be integrated into multi-agent teaching without significant computational overhead or a high Price of Fairness (PoF).

REFERENCES

- [1] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. In *International Conference on Learning Representations (ICLR '18)*, 2018.
- [3] Francis Edward Su. Rental harmony: Sperner’s lemma in fair division. *The American mathematical monthly*, 106(10):930–942, 1999.
- [4] Jiarui Gan, Warut Suksompong, and Alexandros A Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.
- [5] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science (STACS'99)*, pages 404–413. Springer, 1999.

UNIVERSITY OF OXFORD, OXFORD, UK
Email address: jiarui.gan@cs.ox.ac.uk

FAIR ALLOCATION OF INDIVISIBLE CHORES

JUGAL GARG

Classification AMS 2020: 68W25 Approximation algorithms

Keywords: EFX, Earning-Restricted Equilibrium

We study the problem of *fair* allocation of chores among agents with additive preferences. In the discrete setting, envy-freeness up to any chore (EFX) has emerged as a compelling fairness criterion. However, establishing its (non-)existence or achieving a meaningful approximation remains a major open question in fair division. The current best guarantee is the existence of $O(n^2)$ -EFX allocations, where n denotes the number of agents, obtained through a sophisticated algorithm [Zhou and Wu (2022)]. In this paper, we show the existence of 4-EFX allocations, providing the first constant-factor approximation of EFX.

We further investigate the existence of allocations that are both fair and *efficient*, using Pareto optimality (PO) as our efficiency criterion. For the special case of bivalued instances, we establish the existence of allocations that are both 3-EFX and PO, thereby improving upon the current best factor of $O(n)$ -EFX without any efficiency guarantees. For general additive instances, the existence of allocations that are α -EF k and PO has remained open for any constant values of α and k , where EF k denotes envy-freeness up to k chores. We provide the first positive result in this direction by showing the existence of allocations that are 2-EF2 and PO.

Our results are obtained via a novel economic framework called *earning restricted (ER) competitive equilibrium* for fractional allocations, which imposes limits on the earnings of agents from each chore. We show the existence of ER equilibria by carefully formulating a linear complementarity problem (LCP) that captures all ER equilibria, and then prove that the classic complementary pivot algorithm applied to this LCP terminates at an ER equilibrium. By carefully setting earning limits and leveraging the properties of ER equilibria, we design algorithms that involve rounding the fractional solutions and then performing swaps and merges of bundles to meet the desired fairness and efficiency criteria. We expect that the concept of ER equilibrium will play a crucial role in deriving further results on related problems.

REFERENCES

- [Zhou and Wu (2022)] Shengwei Zhou and Xiaowei Wu. Approximately EFX allocations for indivisible chores. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, (IJCAI)*, pages 783–789, 2022.
- [Garg et al. (2025)] Jugal Garg and Aniket Murhekar and John Qin. Constant-Factor EFX Exists for Chores. In *Proceedings of the 57th Symposium on Theory of Computing (STOC)* (to appear), 2025.

UNIVERSITY OF ILLINOIS URBANA-CHAMPAIGN
Email address: jugal@illinois.edu

PREFERENCE AGGREGATION ON THE PROBABILITY SIMPLEX

MATTHIAS GREGER

Classification AMS 2020: 91A35, 91A80, 91B03, 91B14, 91B18

Keywords: mechanism design, collective decision-making, participatory budgeting, portioning, single-peaked preferences

1. INTRODUCTION

Participatory budgeting (see, e.g., [2]) defines a framework that allows citizens to actively participate in deciding how to distribute an exogenously given budget among a set of public projects. For that, the preferences of the citizens (or agents) are collected and aggregated into a collective decision.

The standard model assumes each project has a fixed cost, meaning that it is either fully funded or not funded. However, some instances require us to handle projects and preferences in a more “continuous” way, e.g., when allocating money to whole areas of public interest like education, nature conservation, and public transport. There, preferences depend on the distribution of the budget rather than the set of implemented projects. This class of problems is called portioning, fair mixing, or budget aggregation.

2. RELATED WORK

Having a convex set of outcomes, one has to impose some structure on the agents’ preferences in order to efficiently elicit them. This can be done by asking each agent to report her set of approved projects ([3]) or her favorite distribution. For the latter case, ℓ_1 and Leontief preferences were introduced to the portioning setting by [8] and [5], respectively.

We want to find sensible outcomes and more general mechanisms with good axiomatic properties based on the favorite distributions reported by the agents. Strategyproof mechanisms have initially been investigated by [9] for single-peaked preferences and by [7] for ℓ_1 preferences. We consider two fairness axioms, namely *proportionality* ([7]) and the (α) -core due to [1].

3. MODEL

Given a finite set A of m projects, the set of all possible outcomes consists of all distributions over A and is denoted by $\Delta(1)$. Furthermore, we have a set N of n agents where an agent i ’s preferences over the outcomes can be represented by a continuous and quasi-concave utility function $u_i : \Delta(1) \rightarrow \mathbb{R}$.

Utility functions are assumed to be *star-shaped* [4], i.e., for any distribution $\mathbf{q} \neq \mathbf{p}_i$ and $\lambda \in (0, 1)$, $u_i(\mathbf{p}_i) > u_i(\lambda\mathbf{p}_i + (1 - \lambda)\mathbf{q}) > u_i(\mathbf{q})$, where $\mathbf{p}_i = (p_{i,x})_{x \in A}$ denotes agent i ’s unique peak which corresponds to her favorite outcome.

In my talk, I considered two specific star-shaped utility models, namely ℓ_1 preferences ($u_i(\mathbf{q}) = -\sum_{x \in A} |p_{i,x} - q_x|$) and Leontief preferences ($u_i(\mathbf{q}) = \min_{x \in A: p_{i,x} > 0} q_x/p_{i,x}$).

4. MAIN RESULTS

For $m = 2$, star-shaped preferences coincide with single-peaked preferences, and a unique fair and strategyproof mechanism exists, the uniform phantom mechanism ([6, 7]).

Theorem 4.1 ([5]). *For $m = 2$, the only continuous mechanism that satisfies strategyproofness and proportionality is the uniform phantom mechanism (independent of the underlying utility model).*

In addition, this mechanism always returns Pareto optimal outcomes.

For larger m , the underlying utility model becomes essential. First, I covered an impossibility result for ℓ_1 preferences.

Theorem 4.2 ([5]). *With ℓ_1 preferences, no mechanism satisfies Pareto optimality, strategyproofness, and proportionality when $m \geq 3$ and $n \geq 3$.*

Second, I presented a characterization of the Nash product rule (which returns the outcome that maximizes the product of the agents' utilities) for Leontief preferences.

Theorem 4.3 ([5]). *With Leontief preferences, the Nash product rule is the only continuous mechanism that satisfies group-strategyproofness and always returns core outcomes.*

This theorem implies that Pareto optimality, strategyproofness, and fairness can be achieved simultaneously for Leontief preferences.

REFERENCES

- [1] Robert J. Aumann. The core of a cooperative game without side payments. *Transactions of the American Mathematical Society*, 98, pages 539–552, 1961.
- [2] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236, Springer International Publishing, 2021.
- [3] Anna Bogomolnaia and Hervé Moulin. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2), pages 165–184, 2005.
- [4] Kim C. Border and James S. Jordan. Straightforward elections, unanimity and phantom voters. *Review of Economic Studies*, 50(1), pages 153–170, 1983.
- [5] Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. Optimal budget aggregation with single-peaked preferences. *Proceedings of the 25th ACM Conference on Economics and Computation (ACM-EC)*, page 49, 2024.
- [6] Ioannis Caragiannis, Ariel D. Procaccia, and Nisarg Shah. Truthful univariate estimators. *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, pages 127–135, 2016.
- [7] Rupert Freeman, David M. Pennock, Dominik Peters, and Jennifer Wortman Vaughan. Truthful aggregation of budget proposals. *Journal of Economic Theory*, 193(105234), 2021.
- [8] Tobias Lindner, Klaus Nehring, and Clemens Puppe. Allocating public goods via the midpoint rule. *Proceedings of the 9th International Meeting of the Society of Social Choice and Welfare*, 2008.
- [9] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4), pages 437–455, 1980.

SCHOOL OF COMPUTATION, INFORMATION AND TECHNOLOGY, TECHNICAL UNIVERSITY OF MUNICH, GERMANY

Email address: matthias.greger@tum.de

SOCIAL CHOICE WITH LIMITED QUERIES

DANIEL HALPERN

Classification AMS 2020: 91A68

Keywords: Computational Social Choice, Partial Information, Query Complexity.

1. INTRODUCTION

Social choice theory studies how to aggregate individual preferences into a single collective decision. Traditionally, this assumes complete access to each individual's complete preferences. However, in a variety of settings, this assumption does not hold. For example, modern online platforms promoting civic participation, such as <https://Pol.is>, aggregate complex preferences over a vast space of alternatives, rendering it infeasible to learn any individual's preferences completely. Instead, preferences are elicited by asking each user a query about a small subset of their preferences.

In this talk, I present a simple framework for analyzing what social choice guarantees are possible in these setups, where preferences can only be elicited via small queries. It covers two papers: one on approval-based preferences [1] and one on ranking-based preferences [2].

Contributions include (i) positive algorithmic results: efficient algorithms that produce representative outcomes with limited queries. (ii) information-theoretic impossibilities: fundamental limits on what can be learned, regardless of the number of queries, and (iii) query-complexity lower bounds: situations where, even if it is possible in theory to achieve a desired outcome, an exponential number of queries may be required, making it practically infeasible.

2. MODEL

Preferences. A population of voters have *preferences* over a set of m candidates, C . In the approval-based model, each voter has an *approval set* $A \subseteq C$ of candidates that they approve of. In the ranking-based model, each voter has a ranking σ over C , where σ is a bijection from $C \mapsto \{1, \dots, m\}$, where $\sigma(c)$ indicates the position of candidate c in the ranking. The entire population is represented by a distribution over preferences π , called a profile.

Output Goals. Given a profile π , we would like to output a candidate $c \in C$, or set of candidates $W \subseteq C$ of a fixed size k (also called a committee), that satisfies some desirable property. In the approval-based model, the goal will be finding committees that satisfy some representation axiom. Two of interest are the following. These definitions are adapted from [3] to fit our distributional model.

Definition 2.1 (Justified Representation). A committee $W \subseteq C$ of size $|W| = k$ satisfies Justified Representation (JR) if there is no group of voters making up $1/k$ of the population that all agree on a candidate, and do not like any of the members of W . More formally, there is no candidate $c \in C$ such that $\Pr_{A \sim \pi}[c \in A \text{ and } W \cap A = \emptyset] \geq 1/k$.

Definition 2.2 (Extended Justified Representation). A committee $W \subseteq C$ of size $|W| = k$ satisfies Extended Justified Representation (EJR) if there is no group of voters making up ℓ/k of the population that all agree on ℓ candidates, and do not like any of the members of W . More formally, for all subsets of candidates $S \subseteq C$, $\Pr_{A \sim \pi}[S \subseteq A \text{ and } |W \cap A| < \ell] \geq \ell/k$ where $\ell = |S|$.

For ranking-based preferences we are interested in maximizing scoring rules. A scoring-rule is given by a vector $\mathbf{s} \in \mathbb{R}^m$. Intuitively, each voter gives s_j points to the candidate in their j 'th position. More formally, the score for a candidate $c \in C$ is $\mathbb{E}_{\sigma \sim \pi}[s_{\sigma(c)}]$. Our goal is to output a candidate with maximal score.

Query Model. Instead of us learning the entire population distribution, we interact with it only by querying arriving voters on a subset of candidates. At each time step, our algorithms may choose a subset $Q \subseteq C$ of size t . A random voter arrives from the population, and the algorithm learns their preferences restricted to Q . For approval-based preferences, $A \sim \pi$ and we learn $A \cap Q$. For ranking based preferences, $\sigma \sim \pi$, and we learn the ranking restricted to Q , denoted σ^Q , i.e., σ^Q maps $Q \mapsto \{1, \dots, t\}$, where $\sigma^Q(c)$ is the ranking of c out of the candidates in Q . We are interested in whether our algorithms can find one of the desirable outputs making a reasonable number of queries.

3. RESULTS

For approval-based preferences, we show the following negative result. We say an algorithm is *non-adaptive* if, at each step, it chooses a query Q drawn from a distribution, independent of responses seen so far.

Theorem 3.1. For any constants k and t with $k \geq 2$, and for any $\delta > 0$, any non-adaptive algorithm making fewer than $\Omega(m^2)$ queries outputs a committee of size k satisfying JR with probability at most δ . Furthermore, if $k \geq \Omega(\log(1/\delta))$, then this holds for any algorithm making fewer than $\Omega(m^{11})$ queries.

On the other hand, if we allow for adaptive algorithms, then it is possible to guarantee the stronger EJR.

Theorem 3.2. For any $m \geq t > k$ and any $\delta > 0$, there is an algorithm outputting an EJR committee of size k with probability $1 - \delta$ making at most $O(mk^6 \log k \log m)$ queries.

For ranking-based preferences, we have the following negative result. From [4], it can be inferred that for all $t > 0$, there is a t -dimensional subspace $S^t \subseteq \mathbb{R}^m$ such that using queries of size t , for all $\mathbf{s} \in S^t$, the score of each candidate can be estimated to arbitrary precision assuming enough voters arrive. On the other hand, we show the following.

Theorem 3.3. For all $\mathbf{s} \notin S^t$, there exist m profiles π^1, \dots, π^m such that each has a unique candidate with maximal \mathbf{s} -score, but all are indistinguishable to algorithms making queries of size t .

REFERENCES

- [1] Daniel Halpern, Gregory Kehne, Ariel D. Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. Representation with Incomplete Votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5657–5664, 2023.
- [2] Daniel Halpern, Safwan Hossain, and Jamie Tucker-Foltz. Computing Voting Rules with Elicited Incomplete Votes. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pages 941–963, 2024.
- [3] Harris Aziz, Markus Brill, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare*, 42(2): 461–485, 2017.
- [4] Matthias Bentert and Peter Skowron. Comparing election methods where each voter ranks only few candidates. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2218–2225, 2020.

150 WESTERN AVENUE, OFFICE 5.449, ALLSTON, MA 02134, USA

Email address: dhalpern@g.harvard.edu

ENVY-FREE CAKE-CUTTING FOR FOUR AGENTS

ALEXANDROS HOLLENDER

Classification AMS 2020: 68Q25, 91B32, 68Q17

Keywords: cake cutting, envy-free, query complexity

The aim of this talk is to present some recent results and open questions for the envy-free cake-cutting problem with connected pieces. This is based on a joint paper with Aviad Rubinfeld [7].

Cake-Cutting. The cake-cutting problem, first introduced by Steinhaus [10], is perhaps the most famous, and certainly the most well-studied problem in fair division. The cake serves as a metaphor for modeling settings where the resource is divisible – namely it can be divided arbitrarily – and heterogeneous – different parts of the resource have different attributes and can thus be valued differently by different agents. Examples where the resource can be modeled as a cake, include, among other things, division of land, time, and other natural resources. The problem has been extensively studied in mathematics and economics [2, 9] and more recently in computer science [8].

More formally, the cake is usually modeled as the interval $[0, 1]$, and there are n agents, each with their own preferences over different pieces of the cake. A piece of the cake is simply an interval of $[0, 1]$. The preferences of each agent i are represented by a valuation function v_i which assigns a value to each piece of the cake. Formally, $v_i : [0, 1]^2 \rightarrow [0, 1]$, where $v_i(a, b)$ represents the value that agent i has for interval $[a, b]$. For this to be well-defined, we require that $v_i(a, b) = 0$ whenever $b \leq a$. Furthermore, we always assume that the valuations v_i are continuous.

The goal is to divide the cake into n pieces A_1, \dots, A_n and assign one piece to each agent in a *fair* manner. Here we consider the fairness notion of *envy-freeness*. An allocation is envy-free, if no agent is envious of another agent's piece. Formally, we require that $v_i(A_i) \geq v_i(A_j)$ for all agents i, j . The problem of envy-free cake-cutting was popularized by Gamow and Stern [6] and has since been studied extensively. Stromquist [11], Woodall [14], and more recently Su [13] have shown that an envy-free allocation always exists for continuous valuations.

Computation. The existence proofs all rely on tools such as Brouwer's fixed point theorem, or Sperner's lemma. These powerful tools prove the existence of envy-free allocations, but they do not yield an efficient algorithm for finding one. Mathematicians and economists have tried to address this by proposing so-called *moving-knife* protocols to solve the problem. Unfortunately, the definition of a moving-knife protocol is somewhat informal and thus not well-suited for theoretical investigation (see [3] for a discussion). In more recent years, research in theoretical computer science has started studying these questions in formal models of computation. However, despite extensive efforts on the envy-free cake-cutting problem, its complexity is still poorly understood.

We study the query complexity of the problem. A (value) query consists of the endpoints of an interval $[x, y]$, and the agent responds with its value for that interval, i.e., $v_i(x, y)$. The running time of an algorithm consists of the number of queries. Given that we insist on *connected pieces*, it is known that an exact envy-free allocation cannot be found with a finite number of queries [12]. Thus, we consider *approximate* envy-free allocations instead.¹ An allocation is ε -envy-free, if $v_i(A_i) \geq v_i(A_j) - \varepsilon$ for all agents i, j . In this setting, an ε -envy-free allocation can be found by brute force using $\text{poly}(1/\varepsilon)$ queries when the number of agents is constant [4]. We say that an algorithm is *efficient* if it uses $\text{poly}(\log(1/\varepsilon))$ queries.

A simple algorithm is known for two agents: the *cut-and-choose* protocol. The first agent cuts the cake in two (approximately) equal parts, according to its own valuation, and then the second agent picks its favorite piece, and the first agent receives the remaining piece. This simple algorithm can be implemented using $O(\log(1/\varepsilon))$ queries. For three agents, the problem can be solved using $O(\log^2(1/\varepsilon))$ queries by using an algorithm by Deng, Qi, and Saberi [5]. As shown by Brânzei and Nisan [4], moving-knife algorithms due to Barbanel and Brams [1] and Stromquist [11] can also be simulated in the query model to obtain the same bound. The positive results for three agents require that the valuations be *monotone*, a standard assumption in many works on the topic. A valuation function v_i is monotone if $v_i(A) \geq v_i(B)$, whenever A is a superset of B .

For four agents, the problem has remained open, and Brânzei and Nisan [4] conjectured that there is no efficient algorithm. In our work, we disprove this conjecture.

Theorem 1. *For four agents with monotone Lipschitz valuations, we can compute an ε -envy-free connected allocation using $O(\log^3(1/\varepsilon))$ value queries.*

In the second part of our work, we investigate whether monotonicity is necessary for obtaining efficient algorithms. We prove that this is indeed the case, in a very strong sense. Namely, we show that the *communication complexity* of finding an ε -envy-free allocation with four *non-monotone* agents is $\Omega(\text{poly}(1/\varepsilon))$. To the best of our knowledge, this is the first intractability result for any version of the cake-cutting problem in the communication model. For the case of agents with identical non-monotone valuations, our reduction also yields an $\Omega(\text{poly}(1/\varepsilon))$ query lower bound, as well as a PPAD-hardness result in the standard Turing machine model.

Open problems and future directions. In Tables 1 and 2 we summarize the current state-of-the-art results for connected ε -envy-free cake-cutting, including our results, in two natural models of computation for this problem.

The following two questions are particularly interesting:

- What is the complexity of the problem for five agents with monotone valuations? Can the problem be solved efficiently, or can we show a lower bound? Can we at least show a lower bound for some larger number of players?
- What is the complexity of the problem for three agents with general valuations?

¹We also have to assume that the valuations are L -Lipschitz-continuous for some constant L .

valuations	$n = 2$	$n = 3$	$n = 4$	$n \geq 5$
monotone	$\Theta(\log(1/\varepsilon))$	$O(\log^2(1/\varepsilon))$	$O(\log^3(1/\varepsilon))$?
general		?	$\Theta(\text{poly}(1/\varepsilon))$	$\Theta(\text{poly}(1/\varepsilon))$

TABLE 1. **Query complexity bounds for ε -envy-free cake-cutting.** Here “ $\Theta(\text{poly}(1/\varepsilon))$ ” denotes that there is a polynomial upper bound and a (possibly different) polynomial lower bound. All the lower bounds also hold for agents with identical valuations.

valuations	$n = 2$	$n = 3$	$n = 4$	$n \geq 5$
monotone	$O(\log(1/\varepsilon))$	$O(\log(1/\varepsilon))$	$O(\log^2(1/\varepsilon))$?
general		?	$\Theta(\text{poly}(1/\varepsilon))$	$\Theta(\text{poly}(1/\varepsilon))$

TABLE 2. **Communication complexity bounds for ε -envy-free cake-cutting.**

REFERENCES

- [1] Julius B. Barbanel and Steven J. Brams. Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. *Mathematical Social Sciences*, 48(3):251–269, 2004.
- [2] Steven J. Brams and Alan D. Taylor. *Fair Division: From cake-cutting to dispute resolution*. Cambridge University Press, 1996.
- [3] Steven J. Brams, Alan D. Taylor, and William S. Zwicker. A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of the American Mathematical Society*, 125(2):547–554, 1997.
- [4] Simina Brânzei and Noam Nisan. The query complexity of cake cutting. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pages 37905–37919, 2022.
- [5] Xiaotie Deng, Qi Qi, and Amin Saberi. Algorithmic solutions for envy-free cake cutting. *Oper. Res.*, 60(6):1461–1476, 2012.
- [6] George Gamow and Marvin Stern. *Puzzle-math*. Viking Press, 1958.
- [7] Alexandros Hollender and Aviad Rubinfeld. Envy-free cake-cutting for four agents. In *Proceedings of the IEEE 64th Symposium on Foundations of Computer Science (FOCS)*, pages 113–122, 2023.
- [8] Ariel D. Procaccia. Cake cutting: Not just child’s play. *Communications of the ACM*, 56(7):78–87, 2013.
- [9] Jack Robertson and William Webb. *Cake-cutting algorithms: Be fair if you can*. CRC Press, 1998.
- [10] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.
- [11] Walter Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.
- [12] Walter Stromquist. Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15(1):R11, 2008.
- [13] Francis Edward Su. Rental harmony: Sperner’s lemma in fair division. *The American Mathematical Monthly*, 106(10):930–942, 1999.
- [14] Douglas R. Woodall. Dividing a cake fairly. *Journal of Mathematical Analysis and Applications*, 78(1):233–247, 1980.

ALL SOULS COLLEGE, UNIVERSITY OF OXFORD, UK
 Email address: alexandros.hollender@cs.ox.ac.uk

SOME RECENT RESULTS ON SUPER-STABLE MATCHINGS

NAOYUKI KAMIYAMA

Classification AMS 2020: 68R01, 68W01, 91A68

Keywords: stable matching, super-stability, matroid,

1. INTRODUCTION

The topic of this talk is the stable matching problem in a bipartite graph. This problem was introduced by Gale and Shapley [3], and it is one of the most famous mathematical models of matching. In the basic setting of the stable matching problem, it is assumed that each agent has a strict preference, i.e., the preference of an agent does not contain ties. It is known that, in this setting, a stable matching always exists [3]. In this talk, we consider the stable matching problem with ties. In particular, we focus on super-stability in the stable matching problem with ties. Super-stability is one of the stability concepts in the stable matching problem with ties. Roughly speaking, super-stability guarantees that there does not exist an unmatched pair of agents such that both agents weakly prefer the other agent in the pair to the current partner (i.e., the new partners are not worse than their current partners). It is known that there may not exist a super-stable matching, and the existence of a super-stable matching can be checked in polynomial time [5, 8].

2. MODIFYING AN INSTANCE OF THE SUPER-STABLE MATCHING PROBLEM

The aim of the first part of this talk is to consider how to cope with an instance of the stable matching problem with ties in which there does not exist a super-stable matching. In the first part, we consider the problem of modifying an instance of the stable matching problem with ties by deleting some bounded number of agents in such a way that there exists a super-stable matching in the modified instance. Our problem is motivated by the following question. How far is a given instance of the stable matching problem with ties from the set of instances where there exists a super-stable matching? Our problem gives one approach to this question. If we have to delete many agents from the given instance, then we could conclude that this instance is far from the set of instances where there exists a super-stable matching.

The contribution of this part is summarized as follows. First, we consider the setting where we are allowed to delete agents on only one side. We prove that, in this setting, our problem can be solved in polynomial time. Interestingly, this positive result is obtained by carefully observing the existing algorithm [5, 8] for checking the existence of a super-stable matching. Next, we consider the setting where we are given an upper bound on the number of deleted agents for each side, and we are allowed to delete agents on both sides. We prove that our problem is NP-complete in this setting. See [6] for the details of the results of this part.

3. SUPER-STABLE MATCHINGS WITH GENERALIZED MATROID CONSTRAINTS

In typical settings of variants of the stable matching problem, we are given an upper bound on the number of partners assigned to each agent. In some applications, we need to consider not only an upper bound but also a lower bound on the number of partners assigned to each agent. For example, we consider the problem of assigning residents to hospitals. In order to balance the numbers of residents assigned to hospitals, we could impose lower bounds on the numbers of residents assigned to hospitals.

In this talk, we focus on the classified stable matching problem, which was introduced by Huang [4]. Here we consider the problem of assigning residents to hospitals again. In the classified stable matching problem, a hospital has a laminar family of subsets of acceptable residents, and imposes an upper bound and a lower bound on each subset in this family. In the setting where the preferences of agents are strict, Huang [4] proved that we can check the existence of a stable matching in polynomial time. Furthermore, Fleiner and Kamiyama [2] proposed a matroid approach to the classified stable matching problem, and proved that its many-to-many variant can be solved in polynomial time. Yokoi [9] proposed an abstract generalization of the classified stable matching problem where the constraints are generalized by generalized matroids [1], and proved that, in this setting, we can check the existence of a stable matching in polynomial time.

In the second part, we consider a variant of the problem proposed by Yokoi [9] where the preferences of agents may contain ties. We prove that the problem of determining whether there exists a super-stable matching [5] can be solved in polynomial time in this setting. See [7] for the details of the results of this part.

Acknowledgments. The work was supported by JSPS KAKENHI Grant Number JP20H05795 and JST ERATO Grant Number JPMJER2301, Japan.

REFERENCES

- [1] Éva Tardos. Generalized matroids and supermodular colourings. In András Recski László Lovász, editor, *Colloquia Mathematica Societatis János Bolyai*, volume 40, pages 359–382, Amsterdam, the Netherlands, 1985. North-Holland.
- [2] Tamás Fleiner and Naoyuki Kamiyama. A matroid approach to stable matchings with lower quotas. *Mathematics of Operations Research*, 41(2):734–744, 2016.
- [3] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [4] Chien-Chung Huang. Classified stable matching. In Moses Charikar, editor, *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1235–1253, Philadelphia, PA, 2010. Society for Industrial and Applied Mathematics.
- [5] Robert W. Irving. Stable marriage and indifference. *Discrete Applied Mathematics*, 48(3):261–272, 1994.
- [6] Naoyuki Kamiyama. Modifying an instance of the super-stable matching problem. *Information Processing Letters*, 189:106549, 2025.
- [7] Naoyuki Kamiyama. Super-stable common independent sets of generalized matroids. *Operations Research Letters*, accepted.
- [8] David F. Manlove. Stable marriage with ties and unacceptable partners. Technical Report TR-1999-29, The University of Glasgow, Department of Computing Science, 1999.
- [9] Yu Yokoi. A generalized polymatroid approach to stable matchings with lower quotas. *Mathematics of Operations Research*, 42(1):238–255, 2017.

INSTITUTE OF MATHEMATICS FOR INDUSTRY, KYUSHU UNIVERSITY, FUKUOKA, JAPAN
Email address: kamiyama@imi.kyushu-u.ac.jp

THE RANDOM ASSIGNMENT PROBLEM UNDER CONSTRAINTS

YASUSHI KAWASE

Classification AMS 2020: 91B14, 91B32

Keywords: random assignment, probabilistic serial mechanism, sd-efficiency, sd-envy-freeness, matroid

This is a short report based on joint work with Hanna Sumita and Yu Yokoi [3].

Assigning indivisible items to agents with preferences is one of the most fundamental problems in computer science and economics [4, 5]. Examples of such problems include university housing assignments, student course placements, employee shift assignments, and professional sports drafts. In these kinds of problems, we are given a set of agents, a set of indivisible items, and preferences of the agents. The goal of the problem is to find an assignment that satisfies efficiency and fairness. This study deals with the case where only ordinal information on preferences is available. Such an assumption is common in the literature because eliciting precise cardinal preferences would be difficult in practice (see Bogomolnaia and Moulin [2] for more detailed justifications).

Model. For a nonnegative integer k , we write $[k]$ to denote $\{1, 2, \dots, k\}$. An instance of our problem is a tuple $(N, E, (\succ_i, \mathcal{F}_i)_{i \in N})$, where $N = [n]$ represents the set of agents and $E = \{e_1, e_2, \dots, e_m\}$ represents the set of indivisible items. Each agent $i \in N$ has a strict preference \succ_i over E and can consume a set of items in $\mathcal{F}_i \subseteq 2^E$, which is the feasible set family of agent i . We assume that \mathcal{F}_i is given by a membership oracle for each $i \in N$. The preferences over sets of items are additively separable across items, meaning that each agent i has a cardinal utility function $u_i: E \rightarrow \mathbb{R}_{++}$, and her utility for a bundle $E' \in \mathcal{F}_i$ is $\sum_{e \in E'} u_i(e)$. Here, \mathbb{R}_{++} is the set of positive real numbers. We assume that the preference of each agent i has no ties and that the central authority knows only the ordinal preferences \succ_i that are consistent with u_i . In other words, \succ_i is a strict order on E such that $e \succ_i e'$ if and only if $u_i(e) > u_i(e')$.

For each agent $i \in N$, we assume that the pair (E, \mathcal{F}_i) forms an *independence system*: the feasible set family $\mathcal{F}_i \subseteq 2^E$ is nonempty and satisfies the *hereditary property*, that is, $X \subseteq Y \in \mathcal{F}_i$ implies $X \in \mathcal{F}_i$. We will also consider a special case where each (E, \mathcal{F}_i) is a *matroid*, which is an independence system satisfying a property called the *augmentation axiom*: if $X, Y \in \mathcal{F}_i$ and $|X| < |Y|$ then there exists $e \in Y \setminus X$ such that $X \cup \{e\} \in \mathcal{F}_i$.

A *deterministic assignment* is a list $\mathbf{A} = (A_1, \dots, A_n)$ of subsets of E such that (i) $A_i \in \mathcal{F}_i$ for all $i \in N$ and (ii) $A_i \cap A_j = \emptyset$ for all distinct $i, j \in N$. Let \mathcal{A} be the set of all deterministic assignments. A *lottery assignment* is a probability distribution over \mathcal{A} . We denote the set of all lottery assignments by $\Delta(\mathcal{A})$.

A *fractional assignment* is a matrix $\pi = (\pi_{ie})_{i \in N, e \in E} \in \mathbb{R}^{N \times E}$ such that, for every item $e \in E$, $\sum_{i \in N} \pi_{ie} \leq 1$. We interpret π_{ie} as the probability that agent $i \in N$ receives item $e \in E$. For each $i \in N$, we denote the row in π corresponding to agent i by π_i , that is, $\pi_i = (\pi_{ie})_{e \in E} \in [0, 1]^E$. A lottery assignment $p \in \Delta(\mathcal{A})$ induces a fractional assignment

$\pi \in \mathbb{R}^{N \times E}$ such that $\pi_{ie} = \Pr_{A \sim p}[e \in A_i] = \sum_{A \in \mathcal{A}: e \in A_i} p_A$ for all $i \in N$ and $e \in E$. We will write π^p for the fractional assignment induced from p . A fractional assignment is called *feasible* if it is induced from some lottery assignment.

Desirable Properties. For a preference \succ_i , let $U(\succ_i, e) := \{e' \in E : e' \succeq_i e\}$ be the set of items that are not worse than e with respect to \succ_i . We say that $x \in \mathbb{R}_+^E$ *weakly stochastically dominates* $y \in \mathbb{R}_+^E$, denoted by $x \succeq_i^{\text{sd}} y$, if $\sum_{e' \in U(\succ_i, e)} x_{e'} \geq \sum_{e' \in U(\succ_i, e)} y_{e'}$ for all $e \in E$. If $x \succeq_i^{\text{sd}} y$ and $x \neq y$, we say that x *stochastically dominates* y and denote $x \succ_i^{\text{sd}} y$. Note that x stochastically dominates y if and only if the expected utility of x is greater than that of y for all possible cardinal utilities consistent with \succ_i .

Pareto-efficiency is a standard efficiency concept where no agents can be made better off without making at least one other agent worse off. A natural notion of efficiency for lottery assignments is defined as Pareto-efficiency with respect to the SD relation.

Definition 1 (sd-efficiency). *A lottery assignment $p \in \Delta(\mathcal{A})$ is called sd-efficient (also called ordinally efficient or necessarily Pareto-efficient) if there is no lottery assignment $q \in \Delta(\mathcal{A})$ that satisfies $\pi_i^q \succeq_i^{\text{sd}} \pi_i^p$ for all $i \in N$ and $\pi_j^q \succ_j^{\text{sd}} \pi_j^p$ for some $j \in N$.*

As a notion of fairness, we consider envy-freeness. For the unconstrained setting, a standard definition of sd-envy-freeness requires a fractional assignment to satisfy $\pi_i^p \succeq_i^{\text{sd}} \pi_j^p$ for any agents $i, j \in N$. This condition is equivalent to the expected utility of the fractional assignment of agent i being no worse than that of any other agent j with respect to any cardinal utility consistent to \succ_i [1]. In our setting, however, this equivalence does not hold due to the existence of constraints. Indeed, the bundle assigned to agent j is not feasible for agent i in general. Therefore, we have to take constraints into account when considering each agent's envy toward other agents. For a utility function u_i consistent to \succ_i , let $\tilde{u}_i(X)$ be i 's evaluation of a bundle $X \subseteq E$ (that may be infeasible for i to consume). That is, $\tilde{u}_i(X) = \max\{\sum_{e \in Y} u_i(e) \mid Y \subseteq X, Y \in \mathcal{F}_i\}$. Then, a natural generalization of sd-envy-freeness is to impose a lottery assignment $p \in \Delta(\mathcal{A})$ to satisfy

$$(1) \quad \mathbb{E}_{A \sim p}[\tilde{u}_i(A_i)] \geq \mathbb{E}_{A \sim p}[\tilde{u}_i(A_j)] \quad (\forall i, j \in N, \forall u_i \in \mathbb{R}_{++}^E \text{ consistent to } \succ_i).$$

It turns out that the condition (1) is equivalent to the condition (2) below. Since (2) does not use utility functions, we adopt (2) as the definition of sd-envy-freeness. The envy-freeness with respect to the SD relation is defined as follows.

Definition 2 (sd-envy-freeness). *A lottery assignment $p \in \Delta(\mathcal{A})$ is called sd-envy-free (also called necessary envy-free or not envy-possible) if*

$$(2) \quad \sum_{A \in \mathcal{A}} p_A |A_i \cap U(\succ_i, e)| \geq \sum_{A \in \mathcal{A}} p_A \max_{\substack{Y \subseteq A_j: \\ Y \in \mathcal{F}_i}} |Y \cap U(\succ_i, e)| \quad (\forall i, j \in N, \forall e \in E).$$

Our Results. We investigate the existence of sd-efficient and sd-envy-free assignments in 16 settings according to the following: (i) the number of agents is 2 or arbitrary n , (ii) the constraints are matroids or general hereditary constraints, (iii) the constraints of the agents are identical or heterogeneous, and (iv) the ordinal preferences of the agents are identical or heterogeneous.

We show that an sd-efficient and sd-envy-free assignment always exists in the following cases.

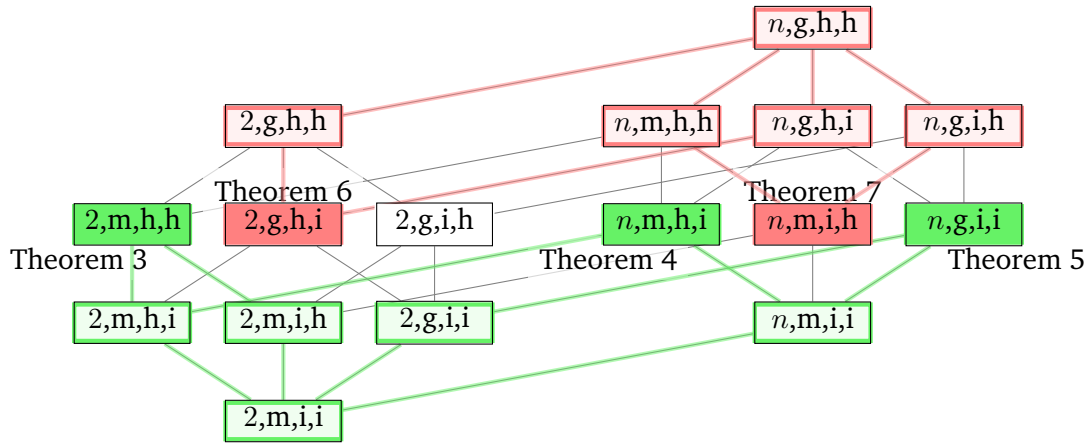


FIGURE 1. Summary of our results on the existence of an sd-efficient and sd-envy-free assignment. Each of the 16 cases is identified by four characters, such as “2,m,i,i.” The first, second, third, and fourth characters, respectively, indicate whether there are 2 or an arbitrary n number of agents, whether the constraints are matroids or general, whether the constraints are identical or heterogeneous, and whether the preferences are identical or heterogeneous. For each case, the box is painted green if such a lottery assignment always exists and red otherwise.

Theorem 3. *A lottery assignment that satisfies sd-efficiency and sd-envy-freeness always exists and can be computed in polynomial time if the number of agents is 2 and the constraints are matroids.*

Theorem 4. *An sd-efficient and sd-envy-free lottery assignment always exists and can be computed in polynomial time if the constraints are matroids, and the preferences are identical.*

Theorem 5. *An sd-efficient and sd-envy-free lottery assignment must exist for any instance with identical constraints and preferences.*

We also show that an sd-efficient and sd-envy-free assignment may not exist in the following cases.

Theorem 6. *An sd-efficient and sd-envy-free lottery assignment may not exist even with two agents, and the preferences are identical.*

Theorem 7. *An sd-efficient and sd-envy-free lottery assignment may not exist even with three agents and identical matroid constraints.*

Taking the inclusion relations into account, we obtain the results shown in Figure 1.

REFERENCES

- [1] H. Aziz, S. Gaspers, S. Mackenzie, and T. Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence*, 227:71–92, 2015.
- [2] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

- [3] Y. Kawase, H. Sumita, and Y. Yokoi. Random assignment of indivisible goods under constraints. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2792–2799, 2023.
- [4] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*. Cambridge University Press, New York, NY, USA, 2007.
- [5] J. Rothe. *Economics and computation*, volume 4. Springer, 2015.
- [6] A. Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998.

UNIVERSITY OF TOKYO, 7-3-1 HONGO, BUNKYO-KU, TOKYO 113-8656, JAPAN
Email address: kawase@mist.i.u-tokyo.ac.jp

FACILITY LOCATION GAMES WITH SCALING EFFECTS

ALEXANDER LAM

Classification AMS 2020: 90B85, 91A10, 91B03, 91B08, 91B14

Keywords: Facility Location, Mechanism Design, Scaling Effects, Strategyproofness

1. INTRODUCTION

In the classic variant of the facility location problem, agents are located on a unidimensional interval, and we are tasked with finding an ideal location to place a facility which serves these agents. The agents have single-peaked preferences for the facility location, as they incur a cost equal to their distance from the facility. This problem models many real-world single-peaked social choice and preference aggregation problems.

In our setting, we assume that the agents' locations are private information, and take a strategyproof mechanism design approach to the classic facility location problem. We also add an additional dimension of complexity: each agent's cost is scaled by an external factor corresponding to the facility's location on the domain. Specifically, there is a (continuous) *scaling function* which maps the facility location to a positive scaling factor which is multiplied by each agent's distance to calculate their individual costs. A local minimum of the scaling function implies that the facility is particularly effective when placed at this point.

The paper most similar to our work considers the facility location model with entrance fees [1], in which agents incur, in addition to their distance from the facility, an additive cost which depends on the facility placement. In this model, the authors give strategyproof mechanisms and compute their approximation ratios with respect to the total and maximum cost objectives. The key difference with our model is that our agents experience a multiplicative scaling factor to their cost, whilst their agents experience an additional additive cost. Our work and [1] are inspired by the line of research on approximate mechanism design without money, initially proposed for the facility location problem by [2]. In their paper, the authors compute the worst-case ratio between the performance of strategyproof mechanisms and welfare-optimal mechanisms.

2. MODEL

We have a set of agents $N = \{1, \dots, n\}$, where agent i has location x_i on the domain¹ $X := [0, 1]$. Although we consider the space of potentially non-anonymous mechanisms, we assume for simplicity that agent locations are ordered such that $x_1 \leq \dots \leq x_n$. This does not affect the nature of our results. A scaling function $q : X \rightarrow \mathbb{R}_{>0}$ gives the

¹Although we consider the unit interval domain, our results can be scaled and shifted to any compact interval on \mathbb{R} . We consider a bounded domain so that a single-peaked linear scaling function does not take negative values.

effectiveness of a facility.² We also assume that the scaling function is continuous. Note that continuity of q is required for the optimal facility location to be well-defined.

Denoting \mathcal{Q} as the space of all scaling functions, a *continuous*³ facility location mechanism $f : \mathcal{Q} \times X^n \rightarrow X$ maps the agent location profile $\mathbf{x} = (x_1, \dots, x_n)$ to the location of a facility y . We define the cost incurred by agent i as its distance from the facility multiplied by the scaling factor: $c_i(q, y) := q(y)|y - x_i|$. Finally, we define the total cost of an instance as $TC(q, y, \mathbf{x}) := \sum_i c_i(q, y) = q(y) \sum_i |y - x_i|$, and the maximum cost of an instance as $MC(q, y, \mathbf{x}) := \max_i c_i(q, y) = q(y) \max_i |y - x_i|$. Respectively, we denote the optimal facility location which minimizes the total cost (resp. maximum cost) as y_{TC}^* (resp. y_{MC}^*).

It is typically ideal for the mechanism output to be independent of the agents' labelling, so we are concerned with mechanisms that satisfy *anonymity*, meaning that output does not change when the agents' labels are permuted.

As we assume agent locations are private information, there is a concern that agents may misreport their locations to unfairly attain a better facility placement. It is therefore ideal to implement a *strategyproof* mechanism, which does not incentivize agents to lie about their locations.

Definition 2.1 (Strategyproofness). *A mechanism $f(\cdot)$ is strategyproof if for every agent $i \in N$, we have, for every scaling function q and agent locations x'_i, \mathbf{x}_{-i} and x_i ,*

$$c_i(q, f(x_i, \mathbf{x}_{-i})) \leq c_i(q, f(x'_i, \mathbf{x}_{-i})).$$

3. COMPUTING THE OPTIMAL SOLUTION

We first investigate the properties of the optimal facility locations for total cost y_{TC}^* and maximum cost y_{MC}^* .

Theorem 3.1. *Let the scaling function q be a continuous, piecewise linear function. The facility location minimizing total cost y_{TC}^* is either on one of the agents' locations or on a local minimum of q .*

Proposition 3.2. *For continuous scaling functions, the optimal facility location for maximum cost satisfies*

$$y_{MC}^* \in \left\{ \arg \min_{y \in [0, \frac{x_1+x_n}{2}]} c_n(q, y), \arg \min_{y \in [\frac{x_1+x_n}{2}, 1]} c_1(q, y) \right\}.$$

Note that neither optimal solution is strategyproof.

4. ACHIEVING SINGLE-PEAKED PREFERENCES

When agents have single-peaked preferences along a compact domain, [3] characterized the set of anonymous and strategyproof mechanisms as *phantom mechanisms*, which place the facility at the median of the n agent locations and $n + 1$ constant 'phantom' points.

²If there is a point y where $q(y) = 0$, then it is trivial to place the facility at that point.

³For mechanisms, we define continuity with respect to the agents' locations. Formally, we say a mechanism is continuous if $\forall \mathbf{x} \in X^n : \forall \epsilon > 0 : \exists \delta > 0 : \forall \mathbf{x}' \in X^n, \forall q \in \mathcal{Q} : \|\mathbf{x} - \mathbf{x}'\|_1 < \delta \implies \|f(q, \mathbf{x}) - f(q, \mathbf{x}')\|_1 < \epsilon$.

Definition 4.1 (Phantom Mechanism). *Given \mathbf{x} and $n + 1$ constant values $0 \leq p_1 \leq p_2 \leq \dots \leq p_{n+1} \leq 1$, a phantom mechanism places the facility at $\text{med}\{x_1, \dots, x_n, p_1, \dots, p_{n+1}\}$.*

However, we find that in our setting, phantom mechanisms are not necessarily strategyproof.

Theorem 4.2. *Every phantom mechanism with $n + 1$ phantoms on at least two unique phantom locations is not strategyproof.*

This is because the scaling function may cause the agent's preference to no longer be single-peaked. However, we can characterize the conditions on the scaling functions which ensure that agents have single-peaked preferences. These conditions hold as long as q is continuous, even if it is not differentiable at a countable set of points D .

Theorem 4.3. *If q is continuous but not differentiable at a countable set of points $D := \{j \mid q'(j) \text{ does not exist}\}$, the agents' preferences are guaranteed to be single-peaked if and only if $q(y) - |q'(y)| \geq 0$ for all $y \in [0, 1] \setminus D$.*

5. CHARACTERIZATION OF STRATEGYPROOF AND ANONYMOUS MECHANISMS

As we have characterized the conditions for the scaling function to guarantee single-peaked agent preferences, it may seem immediate to apply the results by [3] to characterize strategyproof and anonymous mechanisms under these conditions as phantom mechanisms. However, recall that the mechanism takes the scaling function as an additional input, meaning that the constant value p_1, \dots, p_{n+1} may be dependent on the scaling function. Furthermore, the domain of single-peaked preferences induced by a scaling function meeting the key condition of Theorem 4.3 may be a strict subset of the domain of all arbitrary single-peaked preferences. As a result, the characterization by [3] does not immediately hold in our setting.

Nevertheless, we are still able to obtain a similar characterization of strategyproof and anonymous mechanisms in our setting when the scaling function guarantees single-peaked agent preferences, but we additionally require that the mechanism is continuous. We first define an adaptation of the phantom mechanism, in which each 'constant' value admits the scaling function as input.

Definition 5.1 (Phantom Mechanism with Scaling). *Given an agent location profile \mathbf{x} , a scaling function q , and $n + 1$ 'phantom values' $\{p_i(q)\}_{i \in [n+1]}$ defined by $p_i : \mathcal{Q} \rightarrow [0, 1]$, a phantom mechanism with scaling places the facility at*

$$\text{med}\{x_1, \dots, x_n, p_1(q), \dots, p_{n+1}(q)\}.$$

We now present our characterization, which leverages a key result by [4].

Theorem 5.2. *When agents are guaranteed to have single-peaked preferences, a continuous mechanism is strategyproof and anonymous if and only if it is a phantom mechanism with scaling.*

	General Pref.	Single-Peaked Pref.	
	Continuous and Piecewise Linear q	Continuous q	Piecewise Linear q
Total Cost	r_q	e	$(1 + \frac{1}{k})^k$
Max. Cost	$2r_q$	$2e$	$2(1 + \frac{1}{k})^k$

TABLE 1. Phantom mechanism approximation ratio results for total and maximum cost. Agents have single-peaked preferences when the scaling function q meets certain conditions, resulting in phantom mechanisms with scaling being strategyproof. The term r_q denotes the ratio between the max. and min. values of q , and k denotes the number of line segments in the piecewise linear q .

6. APPROXIMATION RATIO RESULTS

Finally, Table 1 gives results on the approximation ratios of phantom mechanisms with scaling. All results are tight in the sense that they are a lower bound for all phantom mechanisms, and that there exists a phantom mechanism with a matching approximation ratio.

REFERENCES

- [1] M. Ma, M. Xiao, T. Bai, and B. Khoussainov. Facility location games with entrance fees. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37, 5797–5804, 2023.
- [2] A. D. Procaccia and M. Tennenholtz. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation*, 1, 2013.
- [3] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 1980.
- [4] F. Brandt, M. Greger, E. Segal-Halevi, and W. Suksompong. Optimal Budget Aggregation with Star-Shaped Preferences. *CoRR*, abs/2402.15904, 2024.

MONG MAN WAI BUILDING, THE HONG KONG POLYTECHNIC UNIVERSITY, HUNG HOM, HONG KONG
Email address: alexander-a.lam@polyu.edu.hk

HOW FAIR CAN STRATEGYPROOF FAIR DIVISION BE?

JÉRÔME LANG

Classification AMS 2020: 68Q25, 91B14, 91B32

Keywords: computational social choice; fair division; strategyproofness; social welfare; fairness

This is a summary of a longer research article co-authored with Sylvain Bouveret, Hugo Gilbert, and Guillaume Méroué [1].

When allocating indivisible items to agents, it is known that the only strategyproof mechanisms that satisfy a set of rather mild conditions are *constrained serial dictatorships* (CSDs): given a fixed order over agents, at each step the designated agent chooses a given number of items (depending on her position in the sequence). Agents who come earlier in the sequence have a larger choice of items; however, this advantage can be compensated by a higher number of items received by those who come later. How to balance priority in the sequence and number of items received is a nontrivial question.

1. RELATED WORK

Various characterization theorems state that, under mild additional conditions, strategyproof allocation mechanisms all have a serial dictatorship flavour: with strict preferences over subsets, only serial dictatorships are strategyproof, neutral, and nonbossy [2], whereas only sequential dictatorships (a generalization of serial dictatorship where the identity of the agent picking in position k depends on the items assigned to the agents in positions 1 to $k - 1$) are strategyproof, Pareto-efficient, and nonbossy [4]. If preferences are quantity-monotonic (a bundle of larger cardinality is always preferred to one of lower cardinality) then a mechanism is strategyproof, nonbossy, Pareto-efficient and neutral if and only if it is a CSD [3]. Similar characterizations hold replacing quantity-monotonic by lexicographic preferences [12, 11]. Ignoring Pareto-efficiency or neutrality opens the door to more complex strategyproof mechanisms; a full characterization in the two-agent case is given in [13]. It is shown in [14] that the CSD where all agents except the last one pick only one item approximates the maxmin fair share criterion.

Sequential allocation of indivisible goods, also known as picking sequences, originates from [15, 10] and have been studied in a number of subsequent works. CSDs correspond to *non-interleaving* picking sequences, where agents pick all their items in a row.

A classical way of guaranteeing a level of fairness and/or efficiency consists in finding an allocation *maximizing social welfare* [16], under the assumption that the input contains, for each agent, her utility function over all bundles of goods (usually assumed additive). Egalitarian social welfare places fairness above all, utilitarian social welfare cares only about efficiency only, and Nash social welfare is considered as a sweet spot inbetween. Surveys of social welfare maximizing fair division can be found in [5, 6, 8]. These mechanisms are not strategyproof.

2. THE MODEL

Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a set of n agents with a_i the i^{th} agent to intervene in the allocation process and $\mathcal{G} = \{g_1, \dots, g_m\}$ a set of m goods. A preference profile $\mathbf{P} = (\succ_{a_1}, \dots, \succ_{a_n})$ describes the preferences of the agents: \succ_a is a ranking that specifies the preferences of agent a over the goods in \mathcal{G} . We denote by $\text{rk}_{\mathbf{P}}^a(g)$, the rank of item g in the ranking of a , given profile \mathbf{P} . *The preference profile is hidden, and therefore not part of the input:* we will assume that rankings are drawn independently according to some probabilistic model, that we denote by Ψ .

Two well-known probabilistic models are the *Mallows* (denoted by $\text{M11}_{\phi, \mu}$) and *Plackett-Luce* (denoted by PL_{ν}) models. These models generalize the two following sub-cases: *Impartial Culture*, denoted by IC, in which each preference ranking is drawn u.a.r. from the set of all possible rankings; The *Full Correlation* case, denoted by FC stipulates that all agents have exactly the same preference ranking.

The items are allocated to the different agents according to a CSD: given a vector $\mathbf{k} = (k_1, \dots, k_n)$ of n non-negative integers, agent a_1 will first pick k_1 goods, then a_2 will pick k_2 goods within the remaining ones, and so on until a_n picks k_n items. In most cases, we will consider complete CSDs, in the sense that $\sum_{i=1}^n k_i = m$. However, we may also consider incomplete CSDs such that $\sum_{i=1}^n k_i < m$. We assume that agents behave greedily by choosing their preferred goods within the remaining ones. This sequential process leads to an allocation that we denote by $\pi_{\mathbf{P}}^{\mathbf{k}}$. More formally, $\pi_{\mathbf{P}}^{\mathbf{k}}$ is a function such that $\pi_{\mathbf{P}}^{\mathbf{k}}(a)$ is the set of goods that agent a has obtained at the end of the sequential allocation process, given preference profile \mathbf{P} and vector \mathbf{k} .

Reusing a model from [17], the utility of an agent for obtaining an item i will be derived using a scoring vector $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{Q}^{+m}$ such that $s_i \geq s_{i+1}$ for all $i = 1, \dots, m-1$. The value received by an agent for obtaining her j^{th} preferred item is s_j . Different scoring vectors can be considered. An important example is the *Borda* scoring vector, where $s_i = m - i + 1$.

We denote by $U_{\mathbf{P}}^{\mathbf{k}}(a) = \sum_{g \in \pi_{\mathbf{P}}^{\mathbf{k}}(a)} s_{\text{rk}_{\mathbf{P}}^a(g)}$ the utility obtained by a when receiving $\pi_{\mathbf{P}}^{\mathbf{k}}(a)$ and by $EU_{\Psi}^{\mathbf{k}}(a) = \mathbb{E}_{\mathbf{P} \sim \Psi}[U_{\mathbf{P}}^{\mathbf{k}}(a)]$ her expected utility given model Ψ . This assumes that agents have *additive* preferences, which is very common in fair division. The utilitarian social welfare (USW) $W_{\Psi}^U(\mathbf{k})$, egalitarian social welfare (ESW) $SW_{\Psi}^E(\mathbf{k})$, and Nash social welfare (NSW) $SW_{\Psi}^N(\mathbf{k})$ are then defined by:

$$SW_{\Psi}^U(\mathbf{k}) = \sum_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a), \quad SW_{\Psi}^E(\mathbf{k}) = \min_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a),$$

$$SW_{\Psi}^N(\mathbf{k}) = \prod_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a).$$

Our objective is to study the following class of optimization problems: Given a number n of agents, a number m of goods, and a scoring vector \mathbf{s} , find a vector $\mathbf{k} = (k_1, \dots, k_n)$ of n non-negative integers with $\sum_{i=1}^n k_i = m$ maximizing $SW_{\Psi}^x(\mathbf{k})$.

3. RESULTS

Depending on the social welfare functional notion and the distribution over profiles, the optimal sequence can be:

- polynomial-time computable,
- efficiently approximated by sampling,
- or hard to approximate by sampling.

The following table summarizes the results obtained for the three social welfare functionals under different distributions; poly means “polynomial-time computable”, and approx means “efficiently approximable by sampling”.

Ψ	$EU_{\Psi}^k(a_i)$	Egal	Nash	Uti
FC	poly	poly	poly	poly
IC	poly	poly	poly	poly
PL_{ν}	approx	approx	approx	approx
$M11_{\phi,\mu}$	approx	approx	?	?

We also performed several experiments to explore further the properties of the CSDs obtained by maximizing USW, NSW or ESW.

REFERENCES

- [1] Sylvain Bouveret, Hugo Gilbert, Jérôme Lang and Guillaume Méroué Thou Shalt not Pick all Items if Thou are First: of Strategyproof and Fair Picking Sequences CoRR abs/2301.06086 (2023).
- [2] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. 16:557-567, 1999.
- [3] Szilvia Pápai. Strategyproof multiple assignment using quotas. Review of Economic Design, 5(1): 91-105, 2000.
- [4] Szilvia Pápai. Strategyproof and Nonbossy Multiple Assignments. Journal of Public Economic Theory, 3(3):257-271,
- [5] Sylvain Bouveret, Yann Chevaleyre and Nicolas Maudet. Fair Allocation of Indivisible Goods. Handbook of Computational Social Choice, 284–310, Cambridge University Press, 2016.
- [6] Jérôme Lang and Jörg Rothe, Fair Division of Indivisible Goods. Economics and Computation, Springer texts in business and economics, 493–550, Springer, 2024 (2nd edition).
- [7] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas and Alexandros A. Voudouris. Fair Division of Indivisible Goods: A Survey. IJCAI-2022, 5385–5393.
- [8] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. Artif. Intell. 322: 103965 (2023).
- [9] Thomas Kalinowski, Nina Narodytska and Toby Walsh. A Social Welfare Optimal Sequential Allocation Procedure, Proceedings IJCAI-2013, 227–233.
- [10] Steven J. Brams and Alan D. Taylor The Win-win Solution. Guaranteeing Fair Shares to Everybody, W. W. Norton & Company, 2000.
- [11] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Fair and efficient allocations under lexicographic preferences. AAAI 2021, 5472-5480.
- [12] Hadi Hosseini and Kate Larson. Multiple assignment problems under lexicographic preferences. AAMAS’19, 837-845.
- [13] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. EC’17, 545-562.
- [14] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On truthful mechanisms for maximin share allocations. IJCAI 2016, 31-37.
- [15] David A. Kohler and R. Chandrasekaran. A class of sequential games. Operations Research, 19(2):270-277, 1971.
- [16] Claude d’Aspremont and Louis Gevers. Social welfare functionals and interpersonal comparability. Chapter 10, volume 1 of Handbook of Social Choice and Welfare, pages 459-541. Elsevier, 2002.
- [17] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. IJCAI 2011, 73-78.

LAMSADE, CNRS AND UNIVERSITÉ PARIS-DAUPHINE, PARIS, FRANCE
E-mail address: lang@lamsade.dauphine.fr

IRRELEVANT ALTERNATIVES ARE RELEVANT.

JÉRÔME LANG

Classification AMS 2020: 68Q25, 91B14, 91B32

Keywords: computational social choice; fair division; strategyproofness; social welfare; fairness

This is a summary of a longer research article co-authored with Théo Delemazure and Grzegorz Pierczyński [1].

We give a quantitative analysis of the *independence of irrelevant alternatives* (IIA) axiom. IIA says that the society’s preference between x and y should depend only on individual preferences between x and y : we show that, in several contexts, if the individuals express their preferences about additional (or “irrelevant”) alternatives, this information helps to estimate better which of x and y has higher social welfare. Our contribution is threefold: (1) we provide a new tool to measure the impact of IIA on social welfare (*pairwise distortion*), based on the well-established notion of voting distortion, (2) we study the average impact of IIA in both general and metric settings, with experiments on synthetic data, and its impact with real datasets; and (3) we study the worst-case impact of IIA in the 1D-Euclidean metric space.

1. RELATED WORK

Distortion has been introduced by [4] as a means to evaluate whether it is reasonable to make a collective decision after eliciting only ordinal preferences. Assuming that cardinal preferences are represented by utilities, the social welfare of an alternative is the sum of the utilities it provides to the agents. The distortion of a voting rule f for a given profile is then defined as the ratio between the maximum social welfare of an alternative, and the social welfare of the alternative selected by f ; and the distortion of f is the maximum, over all profiles, of the distortion of f for that profile. *Metric distortion* [2] aims at minimizing social cost instead of maximizing social welfare: voters and alternatives belong to a metric space, and the cost of an alternative to a voter is the distance between them. See [3] for an extensive survey of the literature of voting distortion until 2021. Average-case analyses of distortion are far less common than worst-case analyses: [7, 8] for single-winner voting, [14] for multi-winner elections, [9] for district-based elections and [6] for social welfare functions.

The primary reason why Arrow imposed IIA was to prevent the implicit use of interpersonal comparisons [5]. However, it also prevents the use of information about *intensities of preferences between two alternatives* revealed by the positions of these alternatives with respect to *other* (“irrelevant”) alternatives. This has been previously discussed in many places, and examples such as the one presented in our introduction highlight its practical negative implications. Because Arrow’s theorem ruled out the existence of a social welfare function under “reasonable” conditions, it has had a negative impact on welfare economics [10, 11, 12]. As other properties stated in Arrow’s theorem can hardly be given up, IIA is the most debatable of the conditions of Arrow’s theorem, and is actually given up *de facto* when defining voting rules. Still, IIA is considered attractive for several reasons, such as avoiding vote splitting [13].

2. THE SETTING

2.1. Pairwise Voting Rules. Let V be a set of n voters and A a set of m alternatives. A ranking \succ of A is a linear order (irreflexive, antisymmetric, transitive and connected relation) of A . $\mathcal{L}(A)$ denotes the set of all rankings over A . A preference profile is a collection of rankings $P = (\succ_1, \dots, \succ_n)$. For a ranking \succ_i , we denote by σ_i the corresponding rank function: for each alternative $x \in A$, $\sigma_i(x) = |\{y \in A \mid y \succ_i x\}| + 1$ the rank of x in \succ_i .

A *pairwise (voting) rule* is a function f that, given a preference profile P over A and two alternatives $x, y \in A$, outputs $f(P \mid x, y) \in \{x, y\}$. Equivalently, f associates with every preference profile P a tournament (an irreflexive, antisymmetric and connected relation, but not necessarily transitive).

A pairwise rule f satisfies IIA if $f(P \mid x, y) = f(P' \mid x, y)$ for all $P = (\succ_1, \dots, \succ_n)$ and $P' = (\succ'_1, \dots, \succ'_n)$ such that for all voters $i \in V$, $x \succ_i y$ if and only if $x \succ'_i y$.

Among pairwise rules that satisfy IIA, the canonical one is the *pairwise majority rule*: $f_{maj}(P \mid x, y) = x$ (resp. y) if a majority of voters prefer x to y (resp. y to x). (Ties are broke by a priority relation over alternatives.)

Another prominent family of pairwise rules consists of those that output transitive tournaments, that is, if $f(P \mid x, y) = x$ and $f(P \mid y, z) = y$ then $f(P \mid x, z) = x$. In this case, f corresponds to a social welfare function g mapping every profile P to a ranking $g(P) \in \mathcal{L}(A)$ defined by $x \succ_{g(P)} y$ if and only if $f(P \mid x, y) = x$. Conversely, any social welfare function g induces a pairwise rule g_{PW} . Among pairwise rules of this class, are all those that are based on a score function Sc that maps every profile P and alternative x to a score $Sc(x, P)$. The pairwise rule f_{Sc} is then defined by $f_{Sc}(P \mid x, y) = \operatorname{argmax}(Sc(x, P), Sc(y, P))$. Examples of such pairwise rules are plurality, Borda, or Copeland.

2.2. Pairwise Distortion. In the unconstrained distortion setting, every voter $i \in V$ receives a utility $U_i(x) \in \mathbb{R}_{\geq 0}$ from alternative $x \in A$. A *utility profile* U is a collection $U = (U_i)_{i \in V}$. We say that a preference profile P and a utility profile U are consistent with each other if for all $x, y \in A$ and all voters i , if $U_i(x) > U_i(y)$, then $x \succ_i y$ in P and we denote it $P \approx U$. The *social welfare* of an alternative $x \in A$ is $SW(x) = \sum_{i \in V} U_i(x)$. The *pairwise distortion* of a pairwise rule f on a utility profile U for two alternatives $x, y \in A$ is the worst-case ratio over all $P \approx U$ between the social welfare of the optimal alternative and that of $f(P \mid x, y)$:

$$\operatorname{dist}(f, U \mid x, y) = \max_{P: P \approx U} \frac{\max(SW_U(x), SW_U(y))}{SW_U(f(P \mid x, y))}$$

Finally, we define the *average pairwise distortion* given a probability distribution \mathcal{D} over utility profiles U . When the distribution is sampled based on a real dataset, we refer to it as *empirical distortion*. Given a utility profile, we obtain a pairwise distortion for each pair of alternatives, which we have then to aggregate; for this we consider two possibilities: taking the *maximum* or the *average* over all pairs.

We also define pairwise distortion (worst-case and average) in the *metric* distortion setting, replacing costs by distances and maximization by minimization.

3. RESULTS

3.1. Average (metric and nonmetric). We first investigate how the average pairwise distortion varies with the number of alternatives m . For all experiments, we use profiles of 30 voters and up to 15 alternatives. We compare distortion for pairwise majority (which satisfies IIA) and four

transitive pairwise rules. Our conclusion is that using information about additional alternatives can help a lot, provided that the way to use it is carefully chosen, and that Borda_{PW} and $\text{Copeland}_{\text{PW}}$ seem both to be good choices.

3.2. Worst-case, metric. In the metric setting, we consider worst-case pairwise distortion, assuming that voters are placed in the metric space so as to maximize the pairwise distortion of a specific pair of alternatives (x, y) , given the positions of all the alternatives. A key question is how to choose the positions of the other alternatives when determining the worst-case pairwise distortion of a pair (x, y) . This can be seen as a game: a first agent selects the positions of the alternatives, and a second agent responds in an adversarial manner by choosing the positions of the voters that maximize pairwise distortion. A *cooperative* (resp. *adversarial*) first agent that places the alternatives so as to minimize (resp. maximize) the worst-case pairwise distortion gives us a lower (resp. upper) bound, called inf-pairwise distortion (resp. sup-pairwise distortion).

The following table gives the worst-case pairwise distortion of various pairwise rules.

	IIA ($\text{majority}_{\text{PW}}$)	Borda_{PW}	k -Approval $_{\text{PW}}$	Plurality $_{\text{PW}}$
inf-pairwise distortion	3	$\frac{m+1}{m-1}$	2	3
sup-pairwise distortion	3	$2m - 1$	∞	∞

4. CONCLUSION

Our conclusions are mixed:

- using information about additional alternatives may help reducing average distortion, but it crucially depends on the choice of the pairwise voting rule used. We found out that — among the rules we studied — the Copeland and Borda pairwise rules are particularly good at decreasing average distortion, but the Plurality pairwise rule has the opposite effect and leads to a larger distortion than sticking to IIA and using pairwise majority.
- when it comes to worst-case distortion, a crucial parameter is the origin of additional alternatives. If they are chosen by the election designer, then the Borda pairwise rule is quite good, and one of its variant, called OddBorda (a rule that may be interesting on its own) is even better. However, if they are chosen adversarially, then better stick to IIA.

REFERENCES

- [1] T. Delemazure, J. Lang, and G. Pierczynski. Independence of Irrelevant Alternatives under the Lens of Pairwise Distortion. *AAAI 2024*: 9645-9652.
- [2] E. Anshelevich, O. Bhardwaj, E. Elkind, J. Postl, and P. Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264: 27-51.
- [3] E. Anshelevich, A. Filos-Ratsikas, N. Shah, and A. Voudouris. Distortion in Social Choice Problems: The First 15 Years and Beyond. *IJCAI-21*, 4294-4301.
- [4] A. D. Procaccia and J.S. Rosenschein. The distortion of cardinal preferences in voting. *International Workshop on Cooperative Information Agents*, 317-331, Springer, 2006.
- [5] K. J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4): 328-346.
- [6] G. Benadè, A. D. Procaccia, and M. Qiao. Low-Distortion Social Welfare Functions. *AAAI-2019*, 1788-1795.
- [7] C. Boutilier, I. Caragiannis, I. S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227: 190-213, 2015.
- [8] Y. A. Gonczarowski, G. Kehne, A. D. Procaccia, B. Schiffer, and S. Zhang. The Distortion of Binomial Voting Defies Expectation. *NeurIPS 2023*.

- [9] A. Filos-Ratsikas, E. Micha, and A. Voudouris. The distortion of distributed voting. *Artif. Intell.* 286: 103343 (2020).
- [10] M. Fleurbaey and P. Mongin. The news of the death of welfare economics is greatly exaggerated. *Soc. Choice Welf.* 25(2-3): 381-418, 2005.
- [11] D. Pearce. Individual and social welfare: a Bayesian perspective. In Frisch Lecture delivered to the World Congress of the Econometric Society, 2021.
- [12] H. Igersheim. The Death of Welfare Economics: History of a Controversy. *History of Political Economy*, 51(5): 827-865.
- [13] E. Maskin. A modified version of Arrow's IIA condition. *Soc. Choice Welf.* 54(2-3): 203-209, 2020.
- [14] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset Selection via Implicit Utilitarian Voting. *J. Artif. Int. Res.*, 58(1): 123-152, 2017.

LAMSADE, CNRS AND UNIVERSITÉ PARIS-DAUPHINE, PARIS, FRANCE

Email address: lang@lamsade.dauphine.fr

MMS ALLOCATION OF INDIVISIBLE CHORES WITH SUBADDITIVE VALUATIONS AND THE FAIR SURVEILLANCE ASSIGNMENT PROBLEM

BO LI

Classification AMS 2020: 68W25 Approximation algorithms

Keywords: Fair Division, Maximin Share, Subadditive, Vertex Cover.

We study the maximin share (MMS) fair allocation of m indivisible chores to n agents who have costs for completing the assigned items. It is known that exact MMS fairness cannot be guaranteed, and so far the best-known approximation for additive cost functions is $11/9$ by Huang and Segal-Halevi [1]; however, beyond additivity, very little is known. In our work [4], we first prove that no algorithm can ensure better than $\min\{n, \frac{\log m}{\log \log m}\}$ approximation if the cost functions are submodular. This result also shows a sharp contrast with the allocation of goods where constant approximations exist as shown by Barman and Krishnamurthy [2] and Ghodsi et al. [3]. We then prove that for subadditive costs, there always exists an allocation that is $\min\{n, \lceil \log m \rceil\}$ -approximation, and thus the approximation ratio is asymptotically tight. Due to these hardness results for the general subadditive setting, we study several specific problems, including job scheduling, bin-packing [4], and vertex cover [5]. For all problems, we show that constant approximate allocations exist.

In particular, the vertex cover cost function is motivated by a generalization of the surveillance problem, where the monitoring area, represented by a graph, is divided and assigned to a set of agents with personalized cost functions. In our work [5], each agent's patrolling cost towards receiving a subgraph is measured by the weight of the minimum vertex cover therein, and our objective is to design algorithms to compute fair assignments of the surveillance tasks. Our main result is an algorithm which ensures a 4.562-approximate MMS allocation for any number of agents with arbitrary vertex weights. We then prove that no algorithm can be better than 2-approximate MMS. For scenarios involving no more than four agents, we improve the approximation ratio to 2, which is thus the optimal achievable ratio.

REFERENCES

- [1] Huang Xin and Segal-Halevi Erel. A Reduction from Chores Allocation to Job Scheduling. Proceedings of the 24th ACM Conference on Economics and Computation. 2023.
- [2] Barman S, Krishnamurthy SK. Approximation algorithms for maximin fair division. ACM Transactions on Economics and Computation. 2020.
- [3] Ghodsi M, HajiAghayi M, Seddighin M, Seddighin S, Yami H. Fair allocation of indivisible goods: Beyond additive valuations. Artificial Intelligence. 2022.
- [4] Li B, Wang F, Zhou Y. Fair allocation of indivisible chores: Beyond additive costs. Advances in Neural Information Processing Systems. 2023.
- [5] Wang F, Li B. Fair Surveillance Assignment Problem. In Proceedings of the ACM Web Conference 2024. 2024.

DEPARTMENT OF COMPUTING, THE HONG KONG POLYTECHNIC UNIVERSITY
Email address: comp-bo.li@polyu.edu.hk

BEST-OF-BOTH-WORLDS FAIR ALLOCATION OF INDIVISIBLE AND MIXED GOODS

XINHANG LU

Classification AMS 2020: 91B32, 91B14

Keywords: Fair division; Indivisible goods; Mixed goods; Randomization and approximation.

This talk is based on a joint work with Xiaolin Bu, Zihao Li, Shengxin Liu, and Biaoshuai Tao [14]. We study the problem of fairly allocating either a set of *indivisible goods* or a set of mixed divisible and indivisible goods (i.e., *mixed goods*) to agents with additive utilities, taking the best-of-both-worlds perspective of guaranteeing fairness properties both ex ante and ex post.

1. INTRODUCTION

Two classic fairness notions in the literature are envy-freeness (EF) and proportionality. An allocation is said to be *envy-free* if each agent values her own bundle the most, and *proportional* if each of n agents gets a bundle of value at least $1/n$ times her value for the entire resources. Despite being desirable properties, neither can always be satisfied when (deterministically) allocating indivisible or mixed goods among the agents.

To circumvent the issue, relaxations of the notions have been proposed and studied. With indivisible goods, *envy-freeness up to any good (EFX)* (resp., *envy-freeness up to one good (EF1)*) requires that an agent's envy towards another agent should be eliminated after the hypothetical removal of any (resp., some) good from the latter agent's bundle [13, 15]. While the existence of EFX allocations is only known in special cases [4, Section 4], the weaker notion of EF1 can always be satisfied [23]. With mixed goods, Bei et al. [11] proposed *envy-freeness for mixed goods (EFM)*, which generalizes both envy-freeness and EF1 as follows: An agent is envy-free towards any agent whose bundle contains some divisible goods and EF1 towards the rest. An EFM allocation always exists for any number of agents with additive utilities.

An alternative and common method to achieve fairness is through randomization. Both envy-freeness and proportionality can be easily and trivially achieved by giving all goods to a single agent uniformly at random. The realized allocation, however, is patently unfair since all agents but one are left empty-handed.

Aziz et al. [5] timely introduced the *best-of-both-worlds* approach, which combines the two aforementioned methods with the goal of constructing a randomized allocation (i.e., a probabilistic distribution over deterministic allocations) that is exactly fair ex ante (before the randomness is realized) and approximately fair ex post (after the randomness is realized). They showed that ex-ante EF and ex-post EF1 can be simultaneously achieved when agents have additive utilities.

In our work, we aim to strengthen the ex-post fairness guarantee to EFX when allocating indivisible goods, and, for the first time, extend the study of best-of-both-worlds fairness to the mixed-goods setting. More specifically, we have the

following two main results. In both cases, we assume that agents have *bi-valued utilities*, i.e., each agent’s utility for each good belongs to one of two possible values.¹

Theorem 1.1. *In the mixed-goods setting where agents have bi-valued utilities, there exists an algorithm which can compute in polynomial time a deterministic allocation sampled from a randomized allocation that is ex-ante proportional and ex-post EFM.*

This result on the compatibility between EFM and ex-ante fairness notions adds to the growing literature revolved around EFM when allocating mixed resources [12, 20, 21, 22, 25]. Our next result concerns indivisible-goods allocation. In addition to fairness, it also takes economic efficiency into consideration. A deterministic allocation of indivisible goods is said to be *Pareto optimal* (resp., *fractionally Pareto optimal (fPO)*) if there is no other deterministic allocation (resp., deterministic or fractional allocation of the indivisible goods) that makes some agent better off without making another worse off.² Clearly, fPO implies PO.

Theorem 1.2. *In the indivisible-goods setting where agents have bi-valued utilities, there exists an algorithm which can compute in polynomial time a deterministic allocation sampled from a randomized allocation that is ex-ante EF, ex-post EFX and ex-post fPO.*

This result generalizes multiple results known in the literature. The compatibility between ex-ante EF, ex-post EFX, and ex-post PO was only known for binary utilities [9, 18].³ For bi-valued utilities, we only knew the compatibility of ex-post notions between EFX and PO [3] as well as between EFX and fPO [17].

2. FURTHER RELATED WORK ON BEST-OF-BOTH-WORLDS FAIRNESS

The fair allocation of indivisible goods has received extensive attention in the past decades [4, 26, 27]. Liu et al. [24] gave an overview of the recent developments of mixed-goods allocations. Below, we discuss more work on best-of-both-worlds fairness.

In addition to the compatibility between ex-ante EF and ex-post EF1, Aziz et al. [5] showed several impossibilities of achieving best-of-both-worlds fair and economically efficient allocations. For agents with unequal *entitlements*, ex-ante weighted EF (WEF) is compatible with ex-post weighted transfer EF up to one good, but not compatible with any stronger ex-post WEF relaxation [6, 19]. For agents with subadditive utilities, ex-ante $\frac{1}{2}$ -EF, ex-post $\frac{1}{2}$ -EFX and ex-post EF1 can be achieved simultaneously [16].

Best-of-both-worlds fairness has also been explored for fair-share-based notions like proportionality and the *maximin share (MMS)* guarantee for agents with additive [2, 10] or fractionally subadditive utilities [1]. Babaioff et al. [10] showed ex-ante proportionality and ex-post $\frac{1}{2}$ -MMS are compatible. The ex-post MMS approximation ratio was improved in [2], at the cost of weakening ex-ante fairness guarantees.

Slightly further afield, the best-of-both-worlds paradigm has also been applied to the contexts of collective choice, such as committee voting [7, 28] and participatory budgeting [8].

¹Binary utilities (the two possible values are 0 and 1) are special cases.

²In an fractional allocation, an indivisible good can be divided and allocated fractionally among agents.

³The result in [9] works for the more general *binary submodular* (also known as *matroid-rank*) utilities.

REFERENCES

- [1] Hannaneh Akrami, Kurt Mehlhorn, Masoud Seddighin, and Golnoosh Shahkarami. Randomized and Deterministic Maximin-share Approximations for Fractionally Subadditive Valuations. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS)*, pages 58821–58832, 2023.
- [2] Hannaneh Akrami, Jugal Garg, Eklavya Sharma, and Setareh Taki. Improving Approximation Guarantees for Maximin Share. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pages 198, 2024.
- [3] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris. Maximum Nash Welfare and Other Stories About EFX. *Theoretical Computer Science*, 863:69–85, 2021.
- [4] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair Division of Indivisible Goods: Recent Progress and Open Questions. *Artificial Intelligence*, 322:103965, 2023.
- [5] Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. *Operations Research*, 72(4):1674–1688, 2024.
- [6] Haris Aziz, Aditya Ganguly, and Evi Micha. Best of Both Worlds Fairness under Entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 941–948, 2023.
- [7] Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. Best-of-Both-Worlds Fairness in Committee Voting. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*, page 676, 2023.
- [8] Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. Fair Lotteries for Participatory Budgeting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, pages 9469–9476, 2024.
- [9] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and Truthful Mechanisms for Dichotomous Valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5119–5126, 2021.
- [10] Moshe Babaioff, Tomer Ezra, and Uriel Feige. On Best-of-Both-Worlds Fair-Share Allocations. In *Proceedings of the 18th Conference on Web and Internet Economics (WINE)*, pages 237–255, 2022.
- [11] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. Fair Division of Mixed Divisible and Indivisible Goods. *Artificial Intelligence*, 293:103436, 2021.
- [12] Umang Bhaskar, A. R. Sricharan, and Rohit Vaish. On Approximate Envy-Freeness for Indivisible Chores and Mixed Resources. In *Proceedings of the 24th International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 1:1–1:23, 2021.
- [13] Eric Budish. The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [14] Xiaolin Bu, Zihao Li, Shengxin Liu, Xinhang Lu, and Biaoshuai Tao. Best-of-both-worlds fair allocation of indivisible and mixed goods. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE)*, 2024. Forthcoming.
- [15] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3):12:1–12:32, 2019.
- [16] Michal Feldman, Simon Mauras, Vishnu V. Narayan, and Tomasz Ponitka. Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pages 1236–1266, 2024.
- [17] Jugal Garg and Aniket Murhekar. Computing Fair and Efficient Allocations with Few Utility Values. *Theoretical Computer Science*, 962:113932, 2023.
- [18] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 370–383, 2020.

- [19] Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. Best of Both Worlds: Agents with Entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 564–572, 2023.
- [20] Bo Li, Zihao Li, Shengxin Liu, and Zekai Wu. Allocating Mixed Goods with Customized Fairness and Indivisibility Ratio. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2868–2876, 2024.
- [21] Zihao Li, Shengxin Liu, Xinhang Lu, and Biaoshuai Tao. Truthful Fair Mechanisms for Allocating Mixed Divisible and Indivisible Goods. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2808–2816, 2023.
- [22] Zihao Li, Shengxin Liu, Xinhang Lu, Biaoshuai Tao, and Yichen Tao. A Complete Landscape for the Price of Envy-Freeness. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1183–1191, 2024.
- [23] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.
- [24] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. Mixed Fair Division: A Survey. *Journal of Artificial Intelligence Research*, 80:1373–1406, 2024.
- [25] Koichi Nishimura and Hanna Sumita. Envy-freeness and Maximum Nash Welfare for Mixed Divisible and Indivisible Goods. arXiv:2302.13342, 2024. URL <https://arxiv.org/abs/2302.13342>.
- [26] Warut Suksompong. Constraints in Fair Division. *ACM SIGecom Exchanges*, 19(2):46–61, 2021.
- [27] Warut Suksompong. Weighted Fair Division of Indivisible Items: A Review. *Information Processing Letters*, 187:106519, 2025.
- [28] Mashbat Suzuki and Jeremy Vollen. Maximum Flow is Fair: A Network Flow Approach to Committee Voting. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pages 964–983, 2024.

UNIVERSITY OF NEW SOUTH WALES, SYDNEY, AUSTRALIA
Email address: xinhang.lu@unsw.edu.au

PROPORTIONAL ALLOCATION OF INDIVISIBLE GOODS UP TO THE LEAST VALUED GOOD ON AVERAGE

RYOGA MAHARA

Classification AMS 2020: 91A68

Keywords: Discrete Fair Division, Indivisible Goods, Proportionality

We study the problem of fairly allocating a set of indivisible goods to multiple agents and focus on the proportionality, which is one of the classical fairness notions. Since proportional allocations do not always exist when goods are indivisible, approximate concepts of proportionality have been considered in the previous work. Among them, proportionality up to the maximin good (PROP_m) has been the best approximate notion of proportionality that can be achieved for all instances [1]. In this study, we introduce the notion of *proportionality up to the least valued good on average* (PROP_{avg}), which is a stronger notion than PROP_m, and show that a PROP_{avg} allocation always exists for all instances and can be computed in polynomial time. Our results establish PROP_{avg} as a notable non-trivial fairness notion that can be achieved for all instances.

Our algorithm can be seen as a generalization of *cut-and-choose* protocol, which is a well-known procedure to fairly allocate resources between two agents. In the cut-and-choose protocol, one agent partitions resources into two bundles for her valuation, and then the other agent chooses the best bundle of the two for her valuation. We generalize this protocol from two agents to n agents in the following way: some $n - 1$ agents partition the goods into n bundles, and then the remaining agent chooses the best bundle among them for her valuation. To apply this protocol, it suffices to show that there exists a partition of the goods into n bundles such that no matter which bundle the remaining agent chooses, the remaining $n - 1$ bundles can be allocated to the first $n - 1$ agents fairly. This technique is interesting by itself and seems to have a potential for further applications.

In our algorithm, we find such a partition by using an auxiliary graph called *PROP_{avg}-graph*. Let us emphasize that introducing the PROP_{avg}-graph is a key technical ingredient in this study. It is also worth noting that Hall's marriage theorem [2], a classical and famous result in discrete mathematics, plays an important role in our argument.

REFERENCES

- [1] Artem Baklanov, Pranav Garimidi, Vasilis Gkatzelis, and Daniel Schoepflin. PROP_m Allocations of Indivisible Goods to Multiple Agents. *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 24–30, 2021.
- [2] Philip Hall. On Representatives of Subsets *Journal of the London Mathematical Society*, 1, 26–30, 1935.

DEPARTMENT OF MATHEMATICAL INFORMATICS, UNIVERSITY OF TOKYO.
Email address: mahara@mist.i.u-tokyo.ac.jp

COUPLES CAN BE TRACTABLE: NEW ALGORITHMS AND HARDNESS RESULTS FOR THE HOSPITALS / RESIDENTS PROBLEM WITH COUPLES

DAVID MANLOVE

Classification AMS 2020: 05C70, 05C90, 68Q25, 68R10, 68W01, 68W25, 91B68

Keywords: junior doctor allocation, matching market, polynomial-time algorithm, NP-hardness, inapproximability

1. BACKGROUND

The STABLE MARRIAGE PROBLEM, and its many-to-one extension, the HOSPITALS / RESIDENTS PROBLEM (HR), are well-studied and central problems in Computational Social Choice, Computer Science, Game Theory and Economics, with numerous applications including in entry-level labour markets, school choice and higher education allocation [14]. In the medical sphere, centralised matching schemes that assign aspiring junior doctors to hospitals operate in many countries. One of the largest and best known examples is the National Resident Matching Program (NRMP) in the US [24]. There are analogous matching schemes for junior doctor allocation in other countries around the world.

The HR problem model represents a bipartite matching market with two-sided preferences, involving the preferences of doctors over acceptable hospitals, and those of hospitals over their applicants. Each hospital has a capacity, indicating the maximum number of doctors that it can admit. Roth [21] argued that a key property to be satisfied by a matching M in an instance I of HR is *stability*, which ensures that there is no *blocking pair*, comprising a doctor and a hospital, both of whom have an incentive to deviate from their assignments in M and become matched to one another, undermining the integrity of M . It is known that every instance of HR admits a stable matching, which may be found in time linear in the size of I [10].

In many of the above applications, there may be couples amongst the applying doctors, who wish to be allocated to hospitals that are geographically close to each other, for example. Indeed, the NRMP matching algorithm was redesigned in 1983 specifically to allow couples to provide preferences over pairs of hospitals, with each pair representing a simultaneous assignment that is suitable for both members of the couple. We thus obtain a generalisation of HR called the HOSPITALS / RESIDENTS PROBLEM WITH COUPLES (HRC). By modifying the definition of stability, taking into account how a couple could improve relative to a matching, Roth [21] showed that the addition of couples destroys the crucial property that a stable matching must always exist. In HRC, the problem therefore is to find a stable matching or report that none exists. Even when a stable matching does exist, such matchings may have different sizes [1]. Worse still, and again in contrast to the case for HR, Ronn [20] showed that HRC is NP-hard.

2. RELATED WORK

In this work, we adopt the stability definition of McDermid and Manlove [17]. With respect to this definition, several algorithmic results for HRC hold. Firstly, Ronn’s NP-hardness result for HRC holds even in the case that each hospital has capacity 1 and there are no single doctors [20].

Ng and Hirschberg [18] also proved NP-hardness for a “dual market” restriction of HRC, which we refer to as HRC-DUAL MARKET, where the two sets H_1 and H_2 , comprising the hospitals appearing in the first (respectively second) positions of the acceptable pairs of the first (respectively second) members of each couple, are disjoint.

Further NP-hardness and polynomial-time solvability results for HRC when preferences lists are bounded in length were established by McDermid and Manlove [17], Biró et al. [4] and Manlove et al. [15]. A simple tractable case of HRC was given by Klaus and Klijn [12] (see also [13]). They required that each couple’s preference list must be *weakly responsive*, each couple must find acceptable all possible outcomes where at least one member is matched, and each hospital has capacity 1.

Given the prevalence of NP-hardness results for HRC, and the scarcity of polynomial-time algorithms, heuristics have been applied to the problem (see [14, Section 5.3,3] for a survey) as well as approaches based on parameterised complexity and local search [16, 3], integer programming [4, 8], constraint programming [15] and SAT solving [9].

Nguyen and Vohra [19] studied so-called *near-feasible* stable matchings in HRC. They showed that if one can modify the capacities of the hospitals by at most 2, then a stable matching with respect to the new capacities always exists (the new total hospital capacity is at least as large as before and increases by at most 4). They also provided an algorithm to find such a near-feasible solution, however their algorithm is not guaranteed to run in polynomial time, as in the first step it computes a stable fractional matching using Scarf’s algorithm [23], and the computation of stable fractional matchings is known to be PPAD-hard [5]. Biró et al. [2] also used Scarf’s algorithm to find stable matchings in HRC instances.

Another direction, to cope with the possible non-existence of a stable matching, is to consider matchings that are “as stable as possible”, i.e., admit the minimum number of blocking pairs. We refer to this problem as MIN BP HRC. Biró et al. [4] showed that this problem is NP-hard and not approximable within $n_D^{1-\varepsilon}$, for any $\varepsilon > 0$, unless $P=NP$, where n_D is the number of doctors in a given instance. Manlove et al. [15] presented integer and constraint programming formulations for MIN BP HRC.

3. OUR CONTRIBUTIONS

We present new polynomial-time algorithms for HRC in the case that the couples’ preference lists are *sub-complete* and *sub-responsive*. Informally, a couple’s preferences are sub-complete if there are underlying preferences for the couple members, such that if both members go to a hospital that is acceptable for them individually, then the pair of hospitals is also acceptable for the couple, and if one member is willing to be unmatched, then any assignment of the other member to an acceptable hospital is acceptable for the couple. The concept of sub-responsiveness is closely related to, but a lot less restrictive than weak responsiveness as described above, even together with sub-completeness, and has been extensively studied by economists [22].

Our main result is a novel and surprising polynomial-time algorithm to find a *near-feasible* stable matching in an HRC instance if the couples' preferences are sub-responsive and sub-complete. Our notion of *near-feasibility* is based on modifying the capacities of each hospital by at most 1. This strengthens Nguyen and Vohra's result [19], albeit for a special case of HRC, in two ways: (i) capacities are varied by at most 1 rather than 2, and (ii) our algorithm runs in polynomial time.

Next, we provide another polynomial-time algorithm for HRC in the presence of sub-responsive and sub-complete preferences that can find a stable matching if all couples are one of several possible types. One of these types corresponds to the very practical and natural restriction of HRC-DUAL MARKET in the case of sub-responsive and sub-complete preference lists, and gives a contrast to the NP-hardness result of Ng and Hirschberg for general HRC-DUAL MARKET instances as mentioned earlier. Using our approach, we argue that this algorithm can potentially be extended to other types of couples, depending on the specific application.

On the structural side, we prove that a version of the Rural Hospitals Theorem for HR remains true even in our HRC setting with sub-responsive and sub-complete preferences, and couples belonging to one of several possible types. These are the first non-trivial classes of HRC instances that we are aware of where these structural properties hold.

We complement our positive results with several hardness results. We show that HRC is NP-hard, even with unit hospital capacities, short preference lists and other strong (simultaneous) restrictions, including (i) sub-responsive and sub-complete couples, and (ii) dual markets and master preference lists [11]. Hence, even in these settings we may not hope to find an exact stable matching in polynomial time, so our algorithm to find a near-feasible stable solution becomes even more appealing.

Finally, we show that MIN BP HRC is not approximable within $m^{1-\varepsilon}$, for any $\varepsilon > 0$, where m is the total length of the hospitals' preference lists, unless $P=NP$, even if each couple applies to only one pair of hospitals.

The polynomial-time algorithms that we present substantially push forward our knowledge of tractable special cases of HRC. Moreover they give additional evidence as to why successful matching schemes such as the NRMP have continued to find stable matchings even in the presence of couples. Our algorithm for the case that the couples' preferences are sub-responsive and sub-complete, and the couples are one of several possible types, also helps to identify the frontier between polynomial-time solvable and NP-hard cases, as our hardness results show that if we have weaker restrictions on the couples' preferences, then it becomes NP-hard to find a stable matching.

This is joint work with Gergely Csáji, Iain McBride and James Trimble. A preliminary version of this paper appeared in the Proceedings of IJCAI 2024 [6]; see [7] for the full version. David Manlove gratefully acknowledges financial support from the Institute for Mathematical Sciences, National University of Singapore, from the School of Computing Science, University of Glasgow, and from the Engineering and Physical Sciences Research Council, grant number EP/P028306/1.

REFERENCES

- [1] B. Aldershof and O.M. Carducci. Stable matching with couples. *Discrete Applied Mathematics*, 68:203–207, 1996.

- [2] P. Biró, T. Fleiner, and R.W. Irving. Matching couples with Scarf’s algorithm. *Annals of Mathematics and Artificial Intelligence*, 77:303–316, 2016.
- [3] P. Biró, R.W. Irving, and I. Schlotter. Stable matching with couples: an empirical study. *ACM Journal of Experimental Algorithmics*, 16, 2011. Section 1, article 2, 27 pages.
- [4] P. Biró, D.F. Manlove, and I. McBride. The Hospitals/Residents problem with Couples: Complexity and integer programming models. In *Proceedings of SEA 2014: the 13th International Symposium on Experimental Algorithms*, volume 8504 of *Lecture Notes in Computer Science*, pages 10–21. Springer, 2014.
- [5] G. Csáji. On the complexity of stable hypergraph matching, stable multicommodity flow and related problems. *Theoretical Computer Science*, 931:1–16, 2022.
- [6] G. Csáji, D. Manlove, I. McBride, and J. Trimble. Couples can be tractable: New algorithms and hardness results for the Hospitals / Residents problem with Couples. In *Proceedings of IJCAI ’24: the 33rd International Joint Conference on Artificial Intelligence*, pages 2731–2739. ijcai.org, 2024.
- [7] Gergely Csáji, David Manlove, Iain McBride, and James Trimble. Couples can be tractable: New algorithms and hardness results for the Hospitals / Residents problem with Couples. arXiv eprint cs.DS 2311.00405. <https://arxiv.org/abs/2311.00405>, 2024.
- [8] M. Delorme, S. García, J. Gondzio, J. Kalcsics, D. Manlove, and W. Pettersson. Stability in the hospitals/residents problem with couples and ties: Mathematical models and computational studies. *Omega*, 103:102386, 2021.
- [9] J. Drummond, A. Perrault, and F. Bacchus. SAT is an effective and complete method for solving stable matching problems with couples. In *Proceedings of IJCAI ’15: the 24th International Joint Conference on Artificial Intelligence*, pages 518–525. AAAI Press, 2015.
- [10] D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [11] R.W. Irving, D.F. Manlove, and S. Scott. The stable marriage problem with master preference lists. *Discrete Applied Mathematics*, 156(15):2959–2977, 2008.
- [12] B. Klaus and F. Klijn. Stable matchings and preferences of couples. *Journal of Economic Theory*, 121(1):75–106, 2005.
- [13] B. Klaus, F. Klijn, and T. Nakamura. Corrigendum: Stable matchings and preferences of couples. *Journal of Economic Theory*, 144(5):2227–2233, 2009.
- [14] D.F. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific, 2013.
- [15] D.F. Manlove, I. McBride, and J. Trimble. “Almost-stable” matchings in the Hospitals / Residents problem with Couples. *Constraints*, 22:50–72, 2017.
- [16] D. Marx and I. Schlotter. Stable assignment with couples: parameterized complexity and local search. *Discrete Optimization*, 8:25–40, 2011.
- [17] E.J. McDermid and D.F. Manlove. Keeping partners together: Algorithmic results for the hospitals / residents problem with couples. *Journal of Combinatorial Optimization*, 19(3):279–303, 2010.
- [18] C. Ng and D.S. Hirschberg. Complexity of the stable marriage and stable roommate problems in three dimensions. Technical Report UCI-ICS 88-28, Department of Information and Computer Science, University of California, Irvine, 1988.
- [19] T. Nguyen and R. Vohra. Near-feasible stable matchings with couples. *American Economic Review*, 108(11):3154–3169, 2018.
- [20] E. Ronn. NP-complete stable matching problems. *Journal of Algorithms*, 11:285–304, 1990.
- [21] A.E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.
- [22] A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis*, volume 18 of *Econometric Society Monographs*. Cambridge University Press, 1990.
- [23] H.E. Scarf. The core of an N-person game. *Econometrica*, 35(1):50–69, 1967.
- [24] <https://www.nrmp.org> (National Resident Matching Program website).

SCHOOL OF COMPUTING SCIENCE, UNIVERSITY OF GLASGOW, GLASGOW G12 8QQ, UK
 Email address: david.manlove@glasgow.ac.uk

FAIR DIVISION FOR RANDOM UTILITIES

PASIN MANURANGSI

Classification AMS 2020: 05C70, 60C05, 68Q87, 68W40

Keywords: Fair Division; Random Model; Envy Freeness; Proportionality

In recent years, several works have studied fair division in a random model where the utilities are drawn from some distributions. A typical question is to determine when the allocation exists (with high probability); this question has been raised for many fairness notions and both in the individual setting—where each bundle is given to a single agent—and in the group setting—where each bundle is given to a group of agents. In this talk, we survey the results and techniques from this line of work.

1. PRELIMINARIES

We use the following notations, which are standard in fair division literature. Let N be the set of n agents and M the set of m goods. For each agent $i \in N$, let $u_i : \mathcal{P}(M) \rightarrow \mathbb{R}_{\geq 0}$ be their utility function where $\mathcal{P}(M)$ denote the power set of M . We assume throughout that the utility function is monotone: $u_i(S) \leq u_i(S')$ for all $S \subseteq S' \subseteq M$.

An allocation $\mathcal{A} = (A_i)_{i \in N}$ is a partition of goods, i.e., A_i 's are disjoint and $\bigcup_{i \in N} A_i = M$. There are many fairness criteria studied in literature; we will only focus on two well-studied notions: Envy-freeness and proportionality.

- An allocation $\mathcal{A} = (A_i)_{i \in N}$ is *envy-free (EF)* if $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$.
- An allocation $\mathcal{A} = (A_i)_{i \in N}$ is *proportional (PROP)* if $u_i(A_i) \geq \frac{1}{n}u_i(M)$ for all $i \in N$.

2. ADDITIVE UTILITIES

It is not hard to see that, in some “worst-case” instances, EF or PROP allocation may not exist. This motivates the study of “average-case” instances. The first—and simplest—such model is proposed by [5] where we assume that there is an underlying distribution \mathcal{D} such that the valuation $u_i(j)$ is drawn independently from \mathcal{D} for each $i \in N, j \in M$, and the utility is assumed to be additive (i.e. $u_i(S) = \sum_{j \in S} u_i(j)$). Much research has been done in this model [5, 7, 11, 8, 1, 9, 10]. We summarize the results for EF and PROP below; for convenience, we assume throughout that \mathcal{D} is the uniform distribution on $(0, 1)$. We say that an event occurs with high probability (w.h.p.) if the probability that it occurs approaches one as $n \rightarrow \infty$ (in this model).

2.1. Proportionality. It has been shown that a PROP allocation exists w.h.p. if $m \geq n$. Note that this is tight since no PROP allocation can exist if $m < n$.

Theorem 2.1 ([10]). *For any $m \geq n$, PROP allocation exists w.h.p.*

For $m = n$, the high-level idea of the proof is to select a certain threshold τ to be slightly above $1/2$, and construct a bipartite graph between the agents and the items such that agent i is connected to item j iff $u_i(j) \geq \tau$. A classic result in random graph

theory [6] states that a perfect matching exists in this “threshold graph” w.h.p. Such a perfect matching corresponds to an allocation where each agent receives a good with valuation at least τ . It can then be shown that such an allocation is PROP w.h.p.

It is not hard to extend the above algorithm to the case where m is divisible by n , e.g. by dividing the goods into m/n subsets of equal size n and using the above matching strategy on each subset. The case where m is not divisible by n is more challenging. Roughly speaking, it requires picking a larger threshold τ (very close to 1) and applies the matching strategy on the first $n \cdot \lfloor m/n \rfloor$ goods. The remaining $m \bmod n$ goods are then used to “fix” the agents whose proportional condition is not yet satisfied.

2.2. Envy-Freeness. For EF, the answer is more subtle and depends on whether m is divisible by n . We start with the divisible case, for which the high-probability existence is known for $m \geq 2n$ and non-existence is known for $m = n$.

Theorem 2.2 ([5, 9]). *For $m = n$, EF allocation does not exist w.h.p. For any $m \geq 2n$ such that m is divisible by n , EF allocation exists w.h.p.*

The non-existence follows by observing that w.h.p. some pair of agents will share the same most-preferred item. When $m = n$ (and all items have different valuations), this implies that EF allocation does not exist.

The existence follows a matching strategy similar to PROP. In fact, the exact same algorithm already suffices for $m \geq 3n$. The $m = 2n$ case requires more care. Specifically, with a constant probability, some pair of agents have the same top-two items, which can make the above matching algorithm fail. By modifying the graph “slightly”, we can avoid such bad cases while still maintaining the high-probability existence of perfect matchings.

Next, we discuss the indivisible case. In this case, it is known that an EF allocation exists as long as $m \geq \Theta\left(\frac{n \log n}{\log \log n}\right)$. Perhaps surprisingly, this is also known to be tight as long as m is not “nearly divisible” by n , as stated more formally below.

Theorem 2.3 ([10, 9]). *For any $m \geq \Theta\left(\frac{n \log n}{\log \log n}\right)$, EF allocation exists w.h.p. For any $m \leq \Theta\left(\frac{n \log n}{\log \log n}\right)$ such that the remainder of m when divided by n is between $[n^{0.1}, n - n^{0.1}]$, EF allocation does not exist w.h.p.*

The existence is shown via the widely-used round robin algorithm, while the non-existence is shown by carefully bounding the probability that each allocation is EF and taking the union bound over all the possible allocations.

3. OTHER MODELS

While the additive, independent and identical assumptions in the previous section allow for convenient analysis, these assumptions might be too strong and unrealistic. As such, several recent works have proposed different models with relaxed assumptions.

3.1. Non-identical Assumption. One such model, due to [3], relaxes the identical distribution assumption by allowing each agent i to have a different distribution \mathcal{D}_i from which $u_i(j)$ is drawn (independently). Among other results, they show that the round-robin algorithm still finds EF allocation w.h.p. for any $m \geq \Theta\left(\frac{n \log n}{\log \log n}\right)$.

3.2. Smoothed Analysis. Another approach taken is via *smoothed analysis*. In a smoothed model proposed by [2], we assume that there are base (worst-case) utility functions u_i . For each $i \in N, j \in M$, with some probability p , a noise $\zeta_{i,j}$ is drawn from some distribution and the final valuation is $u_i(j) = \underline{u}_i(j) + \zeta_{i,j}$. With the remaining probability $1 - p$, the valuation remains the same, i.e., $u_i(j) = \underline{u}_i(j)$. Bai et al. [2] show high-probability existence results for certain regimes of parameters p, m in this model.

3.3. Non-Additive Utilities. Finally, a random model for *non-additive* utilities has been proposed by Benade et al. [4]. Again, we assume that there are base utility functions u_i . Then, the items are randomly renamed for each agent. In particular, we randomly sample a permutation $\pi_i : M \rightarrow M$ independently for each agent i . The final utility function for agent i is $u_i(S) = \underline{u}_i(\pi_i(S))$. Without any additional assumption¹ on \underline{u}_i , for $n = 2$, Benade et al. show that EF allocation exists with probability approaching one if $m \rightarrow \infty$ is an even number (i.e. m is divisible by n). However, such an existence remains open for $n \geq 3$. In this case, the authors show a high-probability existence but only for a restricted class of “order-consistent” utility functions.

REFERENCES

- [1] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi. Approximation algorithms for computing maximin share allocations. *ACM Trans. Algorithms*, 13(4):52:1–52:28, 2017.
- [2] Yushi Bai, Uriel Feige, Paul Gözl, and Ariel D. Procaccia. Fair allocations for smoothed utilities. In *EC*, pages 436–465. ACM, 2022.
- [3] Yushi Bai and Paul Gözl. Envy-free and Pareto-optimal allocations for agents with asymmetric random valuations. In *IJCAI*, pages 53–59, 2022.
- [4] Gerdus Benade, Daniel Halpern, Alexandros Psomas, and Paritosh Verma. On the existence of envy-free allocations beyond additive valuations. In *EC*, page 1287. ACM, 2024.
- [5] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *AAAI*, pages 1405–1411, 2014.
- [6] Paul Erdős and Alfréd Rényi. On random matrices. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 8:455–461, 1964.
- [7] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. When can the maximin share guarantee be guaranteed? In *AAAI*, pages 523–529, 2016.
- [8] Pasin Manurangsi and Warut Suksompong. Asymptotic existence of fair divisions for groups. *Math. Soc. Sci.*, 89:100–108, 2017.
- [9] Pasin Manurangsi and Warut Suksompong. When do envy-free allocations exist? *SIAM J. Discret. Math.*, 34(3):1505–1521, 2020.
- [10] Pasin Manurangsi and Warut Suksompong. Closing gaps in asymptotic fair division. *SIAM J. Discret. Math.*, 35(2):668–706, 2021.
- [11] Warut Suksompong. Asymptotic existence of proportionally fair allocations. *Math. Soc. Sci.*, 81:62–65, 2016.

GOOGLE, PARK VENTURES ECOPLEX, 57 WITTHAYU RD, BANGKOK, THAILAND 10330
 Email address: pasin@google.com

¹In fact, this result holds even for non-monotonic utility functions.

TRUTHFUL BUDGET AGGREGATION: BEYOND MOVING-PHANTOM MECHANISMS

ULRIKE SCHMIDT-KRAEPELIN

Classification AMS 2020: 91B14

Keywords: Budget aggregation, Strategyproofness, Fairness

This is joint work with Mark de Berg, Rupert Freeman, and Markus Utke. For a full version of the paper, we refer to [de Berg et al., 2024].

1. EXTENDED ABSTRACT

Consider a scenario where a perfectly divisible resource, such as money or time, needs to be distributed among various alternatives while taking the preferences of a group of voters into account. This task, known as *portioning*, lies at the heart of participatory budgeting, a voting model gaining increasing attention thanks to its pivotal role in civic participation initiatives [Aziz and Shah, 2021, Cabannes, 2004]. We study a variant called *budget aggregation* [Lindner et al., 2008, Goel et al., 2019, Freeman et al., 2021], where voters express their favorite allocation over alternatives (also called *projects*), and their dissatisfaction with an outcome is measured by its ℓ_1 -distance from their ideal distribution. Unlike many traditional voting scenarios, budget aggregation with ℓ_1 -utilities opens the door to mechanisms that incentivize truthfulness among voters. In fact, Freeman et al. [2021] present a whole class of *truthful* mechanisms called *moving-phantom mechanisms*. Moving-phantom mechanisms are an extension of (the neutral subclass of) *generalized median rules* [Moulin, 1980], characterized as the only truthful mechanisms in the two-alternative setting meeting the criteria of anonymity and continuity [Moulin, 1980, Massó and De Barreda, 2011], to elections with more than two alternatives. However, the question of whether moving-phantom mechanisms are the *only* truthful, continuous, anonymous, and neutral¹ mechanisms in the general case has remained open and was repeatedly mentioned in recent literature [Freeman et al., 2021, Caragiannis et al., 2022, Freeman and Schmidt-Kraepelin, 2024, Brandt et al., 2024]. We resolve this question.

Theorem 1.1 (Informal). *There exists a budget-aggregation mechanism that is truthful, anonymous, neutral, and continuous but not a moving-phantom mechanism.*

To prove Theorem 1.1, we define the class of *cutoff-phantom* mechanisms by combining moving-phantom mechanisms with a cutoff function that redistributes

¹Informally, a mechanism is *truthful* if a voter can never decrease the ℓ_1 -distance between its favorite allocation and the mechanism's outcome by reporting an allocation that is not its favorite. A mechanism is *anonymous* (*neutral*, respectively) if the outcome does not depend on the identity of the voters (alternatives, respectively), and it is *continuous* if it is continuous according to the standard definition, when interpreted as a function.

budget away from any project that the moving-phantom mechanism assigns more than a certain threshold share of the budget. Cutoff-phantoms are well defined for any moving-phantom mechanism and any threshold at least $1/2$, but in general both components need to be chosen carefully to preserve truthfulness. We identify one novel moving-phantom mechanism, GREEDYMAX, for which all of the corresponding cutoff-phantoms (one per choice of threshold) are truthful.

While cutoff-phantoms significantly expand the class of known (anonymous, neutral, continuous) truthful mechanisms, they fail to satisfy unanimity, which prescribes that, whenever the voters all agree on their most preferred distribution, the mechanism should output that distribution. What happens when we add unanimity to our list of properties? We provide a partial answer to this question, by giving a mechanism that is not a moving-phantom mechanism but is unanimous and truthful (as well as anonymous, neutral, and continuous) for instances with two voters and three alternatives. While primarily a proof of concept, this result suggests that the class of truthful budget-aggregation mechanisms satisfying other desirable properties might not allow for a concise description.

One motivation to search for alternative truthful mechanisms stems from the desire for mechanisms that are not only truthful but also fair. To this end, one might consider the MEAN mechanism—the mechanism that averages the voters’ reports for every alternative—as a benchmark. This mechanism appears to be particularly appealing since it is equivalent to assigning each of the n voters their equal *fair share* of the budget and letting them allocate this budget according to their ideal distribution. However, doing so might not be in the voter’s best interest: it is well-known and intuitive that mean aggregation incentivizes voters to extremize their reported preference in order to bring the mean closer to their true preference. In other words, the mean mechanism violates truthfulness. A natural question, then, is how much fairness needs to be compromised in order to restore truthfulness.

To quantify (violations of) fairness, Caragiannis et al. [2022] proposed measuring the worst-case deviation, in terms of ℓ_1 -distance, of a mechanism’s outcome from the mean. They introduced the PIECEWISEUNIFORM mechanism, ensuring an ℓ_1 -distance of $2/3 + \varepsilon$ for some constant $\varepsilon < 10^{-5}$, for the case of three alternatives. Freeman and Schmidt-Kraepelin [2024] presented the LADDER mechanism, establishing an upper bound of $2/3$ for three alternatives and non-trivial bounds for up to six alternatives. They demonstrate that the ladder mechanism results in a worst-case ℓ_∞ -distance from the mean of $\frac{m-1}{2m}$, where m is the number of alternatives. While a lower bound provided by Caragiannis et al. [2022] indicates the tightness of these results within the class of moving-phantom mechanisms, the best lower bounds for ℓ_1 -approximation (and ℓ_∞ -approximation, respectively) for three alternatives within the class of all truthful mechanisms stand at $1/2$ (and $1/4$, respectively).

We show lower bounds matching the best known lower bounds for the class of moving-phantom mechanisms. These bounds are known to be tight for ℓ_∞ -approximation and for ℓ_1 -approximation with $m = 3$, but a gap remains for ℓ_1 with $m > 3$.

Theorem 1.2 (informal). *There exists a budget-aggregation instance for which every truthful, anonymous, neutral, and continuous mechanism returns an outcome with ℓ_∞ -distance of $\frac{m-1}{2m}$ and ℓ_1 -distance of $\frac{m-1}{m}$ from the mean.*

Related Work. Our work contributes to a growing literature on budget aggregation. Lindner et al. [2008] and Goel et al. [2019] study the rule that maximizes utilitarian welfare, the neutral version of which turns out to be the unique Pareto-efficient moving-phantom mechanism [Freeman et al., 2021]. Caragiannis et al. [2022] introduce the paradigm of mean approximation for the budget aggregation problem, which is built upon by Freeman and Schmidt-Kraepelin [2024]. Brandt et al. [2024] show that no mechanism can be truthful, Pareto-efficient, and proportional,² generalizing a result of Freeman et al. [2021] that held only for moving-phantom mechanisms. Elkind et al. [2023] axiomatically study several budget-aggregation mechanisms, and find that the mean performs well relative to the other rules they consider. Goyal et al. [2023] work in a similar setting to ours (except that every alternative’s funding is capped by its predefined cost) and study mechanisms with low sample complexity in terms of their distortion.

For the special case of two alternatives, it is known that truthful and anonymous budget-aggregation mechanisms are characterized by generalized median rules [Moulin, 1980, Massó and De Barreda, 2011]. Generalized median rules are parameterized by $n + 1$ “phantom” votes, with the output being the median of these phantom votes and the n submitted votes. Several papers have used generalized median mechanisms to truthfully approximate the mean in the two-alternative setting [Renault and Trannoy, 2005, 2011, Caragiannis et al., 2016, Jennings et al., 2023], with the optimal approximation stated explicitly by Caragiannis et al. [2022]. However, for higher numbers of alternatives, Caragiannis et al. [2022] obtained a lower bound for general truthful mechanisms that diverged from their lower bound for moving-phantom mechanisms. Our Theorem 1.2 closes this gap.

Beyond budget aggregation, portioning has been studied with other input models including ordinal preferences [Airiau et al., 2023], dichotomous preferences [Bogomolnaia et al., 2005, Brandl et al., 2021, Michorzewski et al., 2020], or more general cardinal utility functions over alternatives [Fain et al., 2016, Wagner and Meir, 2023]. We refer the reader to the survey of Aziz and Shah [2021] for additional discussion of the participatory budgeting literature.

Finally, we remark that a weaker version of our Theorem 1.1 was recently independently obtained by Brandt et al. [2024, Appendix D]. Specifically, the authors show that for the case of a single voter, there exists a truthful, anonymous, neutral, and continuous mechanism that is not a moving-phantom mechanism. Such a mechanism can then be extended to a mechanism satisfying the four properties by concatenating it with any moving-phantom mechanism for the cases of more voters. However, this seems unsatisfactory, since such a mechanism is only not a moving-phantom mechanism for the case of one voter. Our result is significantly stronger since we design mechanisms that are not moving-phantom mechanisms for any number of voters and alternatives.

²Proportionality [Freeman et al., 2021] says that, on instances where every voter prefers to spend the entire budget on a single alternative, the mechanism should output the mean of the votes.

REFERENCES

- S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, J. Lang, and D. Peters. Portioning using ordinal preferences: Fairness and efficiency. *Artificial Intelligence*, 314:103809, 2023.
- H. Aziz and N. Shah. Participatory budgeting: Models and approaches. *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236, 2021.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- F. Brandl, F. Brandt, D. Peters, and C. Stricker. Distribution rules under dichotomous preferences: Two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 158–179, 2021.
- F. Brandt, M. Greger, E. Segal-Halevi, and W. Suksompong. Optimal budget aggregation with single-peaked preferences. *arXiv:2402.15904*, 2024.
- Y. Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and urbanization*, 16(1):27–46, 2004.
- I. Caragiannis, A. Procaccia, and N. Shah. Truthful univariate estimators. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 127–135, 2016.
- I. Caragiannis, G. Christodoulou, and N. Protopapas. Truthful aggregation of budget proposals with proportionality guarantees. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pages 4917–4924, 2022.
- M. de Berg, R. Freeman, U. Schmidt-Kraepelin, and M. Utke. Truthful Budget Aggregation: Beyond Moving Phantom Mechanisms. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE)*, 2024. Extended Abstract. Full version at <https://arxiv.org/abs/2405.20303>.
- E. Elkind, W. Suksompong, and N. Teh. Settling the score: Portioning with cardinal preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence*, pages 621–628, 2023.
- B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th International Conference on Web and Internet Economics*, pages 384–399, 2016.
- R. Freeman and U. Schmidt-Kraepelin. Project-fair and truthful mechanisms for budget aggregation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pages 9704–9712, 2024.
- R. Freeman, D. M. Pennock, D. Peters, and J. W. Vaughan. Truthful aggregation of budget proposals. *Journal of Economic Theory*, 193(3):105234, 2021.
- A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation*, 7(2):1–27, 2019.
- M. Goyal, S. Sakshuwong, S. Sarmasarkar, and A. Goel. Low Sample Complexity Participatory Budgeting. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming*, pages 70:1–70:20, 2023.
- A. B. Jennings, R. Laraki, C. Puppe, and E. M. Varloot. New characterizations of strategy-proofness under single-peakedness. *Mathematical Programming*, pages 1–32, 2023.
- T. Lindner, K. Nehring, and C. Puppe. Allocating public goods via the midpoint rule. In *Proceedings of the 9th International Meeting of the Society for Social Choice and Welfare*,

- 2008.
- J. Massó and I. M. De Barreda. On strategy-proofness and symmetric single-peakedness. *Games and Economic Behavior*, 72(2):467–484, 2011.
- M. Michorzewski, D. Peters, and P. Skowron. Price of fairness in budget division and probabilistic social choice. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*, pages 2184–2191, 2020.
- H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, pages 437–455, 1980.
- R. Renault and A. Trannoy. Protecting minorities through the average voting rule. *Journal of Public Economic Theory*, 7(2):169–199, 2005.
- R. Renault and A. Trannoy. Assessing the extent of strategic manipulation: the average vote example. *SERIEs*, 2:497–513, 2011.
- J. Wagner and R. Meir. Strategy-proof budgeting via a vcg-like mechanism. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory*, pages 401–418, 2023.

TU EINDHOVEN, DE GROENE LOPER 5, 5612 AZ EINDHOVEN
Email address: u.schmidt.kraepelin@tue.nl

FAIRNESS IN REAL ESTATE DIVISION: REBUILD AND DIVIDE

NOGA KLEIN-ELMALEM, RICA GONEN, AND EREL SEGAL-HALEVI

Classification AMS 2024: 91B32

Keywords: Real Estate, Envy-Freeness, Entitlements

At time of this writing, the full paper has not been uploaded to arXiv yet. Please contact us for the uptodate full version.

Urban renewal processes, such as *Rebuild and Divide*, are becoming pivotal in addressing housing shortages and improving infrastructure in densely populated city centers. These processes aim to transform aging residential infrastructure into modern, safe, and more spacious housing while simultaneously increasing urban density.

The process involves demolishing old buildings and constructing new ones, with original homeowners receiving upgraded apartments as compensation. The primary goals include enhancing urban housing availability and improving disaster resilience.

However, disagreements over the assignment of new units often hinder progress, as disputes regarding fairness and equity arise. These disagreements commonly stem from perceptions of unfairness, as homeowners compare the value of their newly assigned units to others. For example, differences in size, floor levels, or unique features of apartments can lead to envy among stakeholders. Courts and legal systems frequently address such disputes, but their resolutions are often unsatisfactory, leading to delays or even project cancellations. For illustration, in a recent court discussion, one homeowner claimed that their new apartment is only 19 square meters larger than their old one, while other homeowners with larger original apartments received upgrades of 23 square meters. The court responded by stating that homeowners with larger original apartments belong to a different “class” and are therefore incomparable to others. This example illustrates the subjective nature of fairness and the challenges in resolving disputes in a way that satisfies all stakeholders.

Current regulations require at least two-thirds of homeowners in a building to agree to participate in a renewal project. This highlights the urgent need for systematic, equitable solutions to apartment allocation.

Monetary compensation in indivisible resource allocation has been widely studied. Demange et al. [2] introduced an ascending auction for envy-free allocation with payments. Halpern and Shah [10] studied the amount of subsidy required for envy-freeness under additive valuations. Brustle et al. [11] and Kawase et al. [17] explored subsidies for envy-freeness under monotone valuations.

Caragiannis and Ioannidis [13] approximated minimum subsidies, though exact computation is hard for many agents. Aziz [12] proposed conditions for envy-freeness and equitability with monetary transfers. Barman et al. [15] studied envy-freeness under dichotomous valuations. Goko et al. [16] developed a truthful subsidy-based allocation for agents with submodular binary valuations.

REFERENCES

The house allocation problem involves assigning m houses to n agents based on preferences, ensuring each agent gets one house [1, 3]. Pareto optimality is a key efficiency criterion [7, 8].

Early work focused on strategy-proofness and stability [6, 4], while recent studies emphasize fairness. Gan et al. [9] provided a polynomial-time algorithm for envy-free assignments, and Kamiyama et al. [14] proved fairness-related computational hardness results.

Rental Harmony defined by Su in [5] is a formal approach to resource allocation that seeks to assign n rooms to n tenants with a fixed total rent R , in such a way that the total prices equals R and no tenant envies the room and price of another tenant. Inspired by this concept, we adapt and extend these principles to the context of *Rebuild and Divide*. In a Rebuild and Redivide problem we have n old units, n new units and n agents, each with subjective valuations. Each agent owns exactly one old unit, and the goal is to assign one new unit to each agent, charging a positive or negative price to each agent such that the total sum of all prices is zero (balanced budget). Additionally, the assignment should ensure that no agent envies another agent's allocation.

In addition to assessing the value of apartments based solely on unit values, we focus on the values of their characteristics, following the approach commonly used in real estate appraisal — consider a collection of apartment characteristics, such as floor level, parking availability, airflow direction, and natural light. Each agent assigns a score to each characteristic. We assume that the value of an apartment is the sum of the values of its characteristics. Under this model, agents may have reduced incentives to manipulate their valuations, as truthful reporting aligns better with their preferences.

In our work, we introduce three distinct models:

- (1) *The Difference Model* examines how each agent compares their own improvement to that of other agents. We provide a necessary and sufficient condition for an allocation to be EF-able. Additionally, we show that in certain scenarios, no EF-able solution exists. In such cases, we aim to identify an allocation and a payment vector that minimize positive envy.
- (2) *The Entitlement Model* evaluates the perceived entitlements of agents, determined by the value of their original apartments. We establish a necessary and sufficient condition for an allocation to be EF-able. However, we also demonstrate that there are cases where no EF-able allocation is feasible.
- (3) *The Envy Sum Model* permits agents to tolerate some envy toward others, as long as the total envy they experience does not exceed the level present in the initial allocation. We establish a necessary and sufficient condition for an allocation to be EF-able. We also prove that it is possible to determine in polynomial time whether an EF-able allocation and payment vector exist and, if so, to compute them.

REFERENCES

- [1] Aanund Hylland and Richard Zeckhauser. “The efficient allocation of individuals to positions”. In: *Journal of Political economy* 87.2 (1979), pp. 293–314.
- [2] Gabrielle Demange, David Gale, and Marilda Sotomayor. “Multi-item auctions”. In: *Journal of political economy* 94.4 (1986), pp. 863–872.

REFERENCES

- [3] Lin Zhou. “On a conjecture by Gale about one-sided matching problems”. In: *Journal of Economic Theory* 52.1 (1990), pp. 123–135.
- [4] Atila Abdulkadiroğlu and Tayfun Sönmez. “House allocation with existing tenants”. In: *Journal of Economic Theory* 88.2 (1999), pp. 233–260.
- [5] Francis Edward Su. “Rental harmony: Sperner’s lemma in fair division”. In: *The American mathematical monthly* 106.10 (1999), pp. 930–942.
- [6] Lars-Gunnar Svensson. “Strategy-proof allocation of indivisible goods”. In: *Social Choice and Welfare* 16.4 (1999), pp. 557–567.
- [7] David J Abraham, Katarína Ceclárová, David F Manlove, and Kurt Mehlhorn. “Pareto optimality in house allocation problems”. In: *International Symposium on Algorithms and Computation*. Springer. 2004, pp. 3–15.
- [8] David Manlove. *Algorithmics of matching under preferences*. Vol. 2. World Scientific, 2013.
- [9] Jiarui Gan, Warut Suksompong, and Alexandros A Voudouris. “Envy-freeness in house allocation problems”. In: *Mathematical Social Sciences* 101 (2019), pp. 104–106.
- [10] Daniel Halpern and Nisarg Shah. “Fair division with subsidy”. In: *Algorithmic Game Theory: 12th International Symposium, SAGT 2019, Athens, Greece, September 30–October 3, 2019, Proceedings 12*. Springer. 2019, pp. 374–389.
- [11] Johannes Brustle, Jack Dippel, Vishnu V Narayan, Mashbat Suzuki, and Adrian Vetta. “One dollar each eliminates envy”. In: *Proceedings of the 21st ACM Conference on Economics and Computation*. 2020, pp. 23–39.
- [12] Haris Aziz. “Achieving envy-freeness and equitability with monetary transfers”. In: 35.6 (2021), pp. 5102–5109.
- [13] Ioannis Caragiannis and Stavros D Ioannidis. “Computing envy-freeable allocations with limited subsidies”. In: *International Conference on Web and Internet Economics*. Springer. 2021, pp. 522–539.
- [14] Naoyuki Kamiyama, Pasin Manurangsi, and Warut Suksompong. “On the complexity of fair house allocation”. In: *Operations Research Letters* 49.4 (2021), pp. 572–577.
- [15] Siddharth Barman, Anand Krishna, Yadati Narahari, and Soumyarup Sadhukhan. “Achieving envy-freeness with limited subsidies under dichotomous valuations”. In: *arXiv preprint arXiv:2201.07419* (2022).
- [16] Hiromichi Goko, Ayumi Igarashi, Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, Yu Yokoi, and Makoto Yokoo. “A fair and truthful mechanism with limited subsidy”. In: *Games and Economic Behavior* 144 (2024), pp. 49–70.
- [17] Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, and Makoto Yokoo. “Towards optimal subsidy bounds for envy-freeable allocations”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 9. 2024, pp. 9824–9831.

MY ADDRESS

Email address: erelsgl@gmail.com

REDUCING LEXIMIN FAIRNESS TO UTILITARIAN OPTIMIZATION

EDEN HARTMAN, YONATAN AUMANN, AVINATAN HASSIDIM, AND EREL SEGAL-HALEVI*

Classification AMS 2024: 91B32

Keywords: Egalitarianism, Utilitarianism, Leximin

Full Version: <https://arxiv.org/abs/2409.10395>

1. MOTIVATION

In social choice, the goal is to find the best choice for society, but 'best' can be defined in many ways. Two frequent, and often contrasting definitions are the *utilitarian best*, which focuses on maximizing the total welfare (i.e., the sum of utilities); and the *egalitarian best*, which focuses on maximizing the least utility. The *leximin best* generalizes the egalitarian one. It first aims to maximize the least utility; then, among all options that maximize the least utility, it chooses the one that maximizes the second-smallest utility, among these — the third-smallest utility, and so forth. Leximin is often the solution of choice in social choice applications, and frequently used (e.g., [7, 4, 5, 6]).

Calculating a choice that maximizes utilitarian welfare is often easier than finding one that maximizes egalitarian welfare, while finding one that is leximin optimal is typically even more complex. For example, when allocating indivisible goods among agents with additive utilities, finding a choice (in this case, an allocation) that maximizes the utilitarian welfare can be done by greedily assigning each item to the agent who values it most. Finding an allocation that maximizes the egalitarian welfare, however, is NP-hard [3], even in this relatively simple case.

2. THE REDUCTION IDEA

In the paper, we present a general reduction from leximin to utilitarian. Specifically, for any social choice problem with non-negative utilities, we prove that given a black-box that returns a state (deterministic outcome) that maximizes the utilitarian welfare, one can obtain a polynomial-time algorithm that returns a lottery over states that is leximin-optimal with respect to the agents' expected utilities.

Our reduction *extends to approximations*, meaning that given an α -approximate solver for utilitarian welfare for some $\alpha \in (0, 1)$, the output is an α -leximin-approximation, preserving the same approximation factor. The reduction also *extends to randomized solvers*, which means that given a solver that returns a utilitarian-optimal state only with some probability $p \in (0, 1)$, then the output is leximin-optimal with the same probability p . Furthermore, the reduction applies even when the utilitarian solver is both approximate and randomized simultaneously.

In all, with our reduction at hand, optimizing leximin in expectation is no more difficult than optimizing the utilitarian welfare.

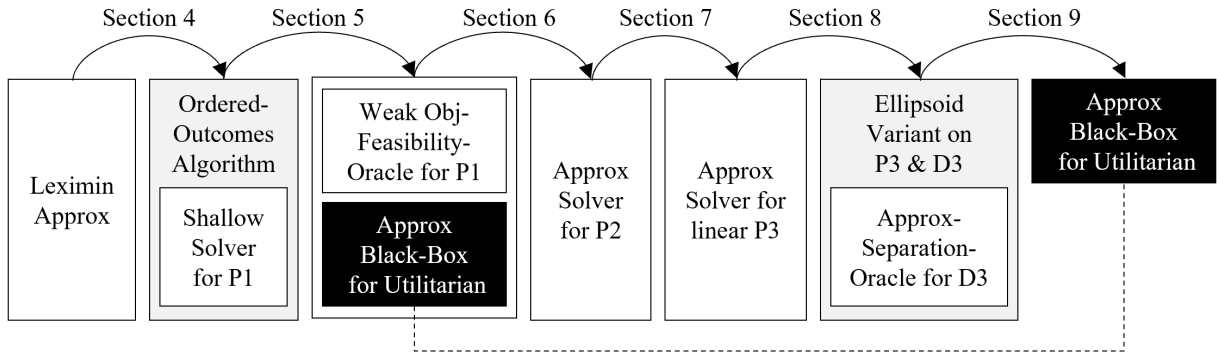


FIGURE 1. High level description of the reduction algorithm. An arrow from element A to B denotes that the corresponding section reduces problem A to B. White components are implemented in the paper; gray components represent existing algorithms; the black component is the black-box for the utilitarian welfare.

3. HIGH-LEVEL DESCRIPTION OF THE REDUCTION

The reduction is done step by step. At each step, we simplify the problem further until we rely only on a solver for utilitarian welfare.

We start with leximin-optimization and consider an iterative algorithm introduced by [9], called the *Ordered Outcomes Algorithm*. This algorithm iteratively solves an evolving mathematical program, P_1 , which, in our context, is non-trivial to solve or approximate. We define a new type of solver for P_1 , which we call *shallow*, and prove that when using a shallow solver for P_1 , the Ordered Outcomes algorithm returns a leximin-approximation. Next, we show that such a solver can be designed by perform a binary search over the potential objective values – where we use the black-box for the utilitarian welfare to get an upper bound for the search; and operate a procedure we call *Weak Feasibility Oracle for P_1* on each value. We then prove that this oracle can be obtained by approximating a different mathematical program, P_2 , and subsequently show that this can be achieved by approximating an equivalent linear program, P_3 . Finally, we prove that P_3 can be approximated using a variant of the ellipsoid method applied to this program and its dual, D_3 ; this method requires an approximate separation oracle for the dual, which can be implemented using the given black-box.

A schematic description of the reduction structure is provided in Figure 1.

4. APPLICATIONS

We demonstrate the significance of this reduction by applying it to three social choice problems as follows.

First, we consider the classic problem of *allocations of indivisible goods* [8], where one seeks to fairly distribute a a set of indivisible goods among a set of heterogeneous agents. Maximizing the utilitarian welfare in this case is well-studied. Using our reduction, the previously mentioned greedy algorithm for agents with additive utilities, allows us to achieve a leximin optimal lottery over the allocations in polynomial time. For submodular utilities, approximating leximin to a factor better than $(1 - \frac{1}{e})$ is NP-hard. However, by

applying our reduction, existing approximation algorithms for utilitarian welfare can be leveraged to prove that a 0.5-approximation can be obtained deterministically, while the best-possible $(1 - \frac{1}{e})$ -approximation can be obtained with high probability.

Second, we consider the problem of *giveaway lotteries* [1], where there is an event with limited capacity and groups wish to attend, but only-if they can all be admitted together. Maximizing the utilitarian welfare in this setting can be seen as a knapsack problem, for which there is a well-known FPTAS (fully polynomial-time approximation scheme). Using our reduction, we obtain an FPTAS for leximin as well.

Lastly, we consider the problem of *fair lotteries for participatory budgeting* [2], where one seeks to find a fair lottery over the possible budget allocations. When agents have additive utilities, maximizing the utilitarian welfare can also be interpreted as a knapsack problem (albeit in a slightly different way), which allows us to provide an FPTAS for leximin.

REFERENCES

- [1] Tal Arbib and Yonatan Aumann. Fair and Truthful Giveaway Lotteries. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4785–4792, June 2022.
- [2] Haris Aziz, Xinhang Lu, Mashbat Suzuki, Jeremy Vollen, and Toby Walsh. Fair lotteries for participatory budgeting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9469–9476, 2024.
- [3] Nikhil Bansal and Maxim Sviridenko. The santa claus problem. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, page 31–40, New York, NY, USA, 2006. Association for Computing Machinery.
- [4] Xiaohui Bei, Xinhang Lu, and Warut Suksompong. Truthful Cake Sharing. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):4809–4817, June 2022.
- [5] Hao Cheng, Shufeng Kong, Yanchen Deng, Caihua Liu, Xiaohu Wu, Bo An, and Chongjun Wang. Exploring leximin principle for fair core-selecting combinatorial auctions: Payment rule design and implementation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, pages 2581–2588. ijcai.org, 2023.
- [6] Bailey Flanigan, Jennifer Liang, Ariel D. Procaccia, and Sven Wang. Manipulation-robust selection of citizens’ assemblies. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan, editors, *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9696–9703. AAAI Press, 2024.
- [7] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Equitable allocations of indivisible goods. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-2019*, pages 280–286. ijcai.org, 2019.
- [8] Yasushi Kawase and Hanna Sumita. On the Max-Min Fair Stochastic Allocation of Indivisible Goods. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(02):2070–2078, April 2020.
- [9] Włodzimierz Ogryczak and Tomasz Śliwiński. On direct methods for lexicographic min-max optimization. In Marina Gavrilova, Osvaldo Gervasi, Vipin Kumar, C. J. Kenneth Tan, David Taniar, Antonio Laganá, Youngsong Mun, and Hyunseung Choo, editors, *Computational Science and Its Applications - ICCSA 2006*, pages 802–811, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

ARIEL UNIVERSITY, ISRAEL

Email address: erelsgl@gmail.com

FAIR ALLOCATION WITH BINARY VALUATIONS FOR MIXED DIVISIBLE AND INDIVISIBLE GOODS VIA HYBRID CONVEX OPTIMIZATION

HANNA SUMITA

Classification AMS 2020: 91B14, 91B32

Keywords: Fair allocation, Mixed goods, Maximum Nash welfare

In this talk, we study fair allocation problem of mixed goods. We are given a set $N = \{1, 2, \dots, n\}$ of n agents and a set $E = C \cup M$, where C and M are the sets of divisible and indivisible goods, respectively. We are also given a *symmetric strictly convex* function $\Phi: \mathbb{R}^N \rightarrow \mathbb{R}$. Each agent i has a binary valuation $v_{ie} \in \{0, 1\}$ for each good e , meaning that for each good, each agent either desires it or not. An allocation is a matrix $\pi \in [0, 1]^{N \times E}$ such that $\pi_{ie} \in \{0, 1\}$ for all agent $i \in N$ and indivisible good $e \in M$. The entry π_{ie} means the allocated amount of good e to agent i . Throughout this paper, we only consider utilitarian optimum allocations, that is, $\pi_{ie} > 0$ only if $v_{ie} = 1$. Agents have additive utility, and the utility of agent i in allocation π is $\pi_i(E) = \sum_{e \in E} v_{ie} \pi_{ie}$. For an allocation π , a vector $z = (\pi_1(E), \dots, \pi_n(E))$ is called a utility vector of π . We say that an allocation π is Φ -fair if its utility vector z minimizes $\Phi(z)$ among allocations. The goal of our problem is to find a Φ -fair allocation.

By appropriately setting a function Φ , this problem is equivalent to finding an allocation achieving *maximum Nash welfare* (MNW). There also exists a function Φ such that a Φ -fair allocation is *leximin*, i.e., the smallest utility is maximized and subject to that, the second smallest utility is maximized, and so on.

There is a vast body of literature on the allocation of goods in cases where only divisible or indivisible goods are present. In such cases, it suffices to find a feasible minimizer (utility vector) of Φ because we can find an allocation achieving the utilities by solving the maximum flow problem. In the continuous case, where there are only divisible goods, the set of possible utility vectors forms an *integral base-polyhedron*. This is a polyhedron represented as $\bar{\mathbf{B}} = \{x \in \mathbb{R}^N \mid x(N) = f(N), x(X) \leq f(X) (\forall X \subseteq N)\}$, where $f: 2^N \rightarrow \mathbb{Z}$ is an integer-valued submodular function with $f(\emptyset) = 0$. It is known that an integral base-polyhedron has a common unique minimizer independent of Φ , and the minimizer can be characterized by a structure called the *principal partition* [5, 7]. The minimizer can also be found in polynomial time [5, 6]. In the discrete case, where there are only indivisible goods, the set of possible utility vectors forms an *M-convex set*, which is the set $\bar{\mathbf{B}}$ of integral vectors in an integral base-polyhedron. It is known that a minimizer of Φ on an M-convex set can be characterized by the *canonical partition* [2], which is an aggregation of the principal partition. Additionally, the set of minimizers of a symmetric strictly convex function does not depend on the function [2] and a minimizer (utility vector) can be found in polynomial time [3]. Furthermore, a *proximity theorem* has been established [4]. This theorem states that a minimizer of Φ in an M-convex set

lies within a unit hypercube that contains the minimizer in the corresponding integral base-polyhedron.

Our problem is regarded as the hybrid of continuous and discrete optimization problems: finding an allocation whose utility vector z minimizes $\Phi(z)$ under the constraint that $z \in \mathbf{B}_E := \bar{\mathbf{B}}_C + \ddot{\mathbf{B}}_M$, where \mathbf{B}_E is the Minkowski sum of a given integral base-polyhedron $\bar{\mathbf{B}}_C$ and a given M-convex set $\ddot{\mathbf{B}}_M$. Note that the convex hull of \mathbf{B}_E , denoted by $\bar{\mathbf{B}}_E$, is an integral base-polyhedron, and the set of integral points in \mathbf{B}_E is an M-convex set.

In this talk, we show the following results.

Structure of Φ -fair allocations. First, we investigate the structure of the fair allocation of mixed goods. Unfortunately, the set of minimizers $\arg \min_{z \in \mathbf{B}_E} \Phi(z)$ depends on Φ in general. Nevertheless, we show that the hybrid problem still retains a structure of the canonical partition. By using this fact, we prove a proximity theorem. We remark that this is not specialized to the fair allocation setting, and thus it extends the existing result for the discrete case [4].

Theorem 1. *Let Φ be a symmetric strictly convex function. For any $z^* \in \arg \min_{z \in \mathbf{B}_E} \Phi(z)$ and $\bar{z} \in \arg \min_{z \in \bar{\mathbf{B}}_E} \Phi(z)$, we have $\lfloor \bar{z}_i \rfloor \leq z_i^* \leq \lceil \bar{z}_i \rceil$ for all $i \in N$.*

By using the canonical partition, we also show that there exists partitions E_1, \dots, E_q of goods E and N_1, \dots, N_q of agents N such that goods in E_j are allocated to agents in N_j for each $j = 1, 2, \dots, q$ in any Φ -fair allocation.

NP-hardness. Next, by utilizing the proximity theorem, we show that our problem is NP-hard even when indivisible goods are *identical*, i.e., for each agent i , either $v_{ie} = 1$ ($\forall e \in M$) or $v_{ie} = 0$ ($\forall e \in M$).

Theorem 2. *For any fixed symmetric strictly convex function Φ , finding a Φ -fair allocation is NP-hard even when indivisible goods are identical.*

We also prove that computing an MNW allocation and an leximin allocation are both NP-hard. These results highlight the difficulty of the mixed goods case.

Polynomial-time solvable case. In contrast, we show the following tractability when divisible goods are identical.

Theorem 3. *Let Φ be a symmetric strictly convex function. There exists a polynomial-time algorithm that finds a Φ -fair allocation if all the divisible goods are identical.*

A key tool to construct the algorithm is the structure of Φ -fair allocation. We can find the above-mentioned partitions E_1, \dots, E_q of E and N_1, \dots, N_q of N in polynomial time by using the canonical partition. Because of the structure, we can find a minimizer of Φ by independently solving the subproblems of assigning goods in E_j to agents in N_j for $j = 1, 2, \dots, q$. When divisible goods are identical, there exists j^* such that E_{j^*} has all the divisible goods, and the utilities of the agents receiving a piece of divisible goods are the same. By using this property, we can enumerate possible utility vectors of agents and check the existence of the corresponding allocation by submodular flow [6].

Connection to relaxed envy-freeness. As another consequence of our structural result, we also show the connection between MNW allocations and relaxed envy-freeness. For binary valuations, we prove that MNW allocations satisfy *envy-freeness up to any good for mixed goods* (EFXM): if agent i envies agent j , then agent j has no piece of divisible goods and the envy can be eliminated by removing any indivisible good in agent j 's bundle. The notion of EFXM coincides with envy-freeness (EF) when only divisible goods exists, and envy-freeness up to any good (EFX) when only indivisible goods exists. It is known that any MNW allocation for divisible goods is EF [8] and the one for indivisible goods is EFX [1]. Thus, our result is a generalization of these results. We also mention results on other classes of valuations.

This talk is based on the joint work with Yasushi Kawase and Koichi Nishimura.

REFERENCES

- [1] G. Amanatidis, G. Birmpas, A. Filos-Ratsikas, A. Hollender, and A. A. Voudouris. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863:69–85, 2021.
- [2] A. Frank and K. Murota. Decreasing minimization on M-convex sets: background and structures. *Mathematical Programming*, 195(1–2):977–1025, 2022.
- [3] A. Frank and K. Murota. Decreasing minimization on M-convex sets: algorithms and applications. *Mathematical Programming*, 195(1–2):1027–1068, 2022.
- [4] A. Frank and K. Murota. Decreasing minimization on base-polyhedra: Relation between discrete and continuous cases. *Japan Journal of Industrial and Applied Mathematics*, 40(1):183–221, 2023.
- [5] S. Fujishige. Lexicographically optimal base of a polymatroid with respect to a weight vector. *Mathematics of Operations Research*, 5(2):186–196, 1980.
- [6] S. Fujishige. *Submodular Functions and Optimization*. Elsevier, 2nd edition, 2005.
- [7] F. Maruyama. A unified study on problems in information theory via polymatroids (in Japanese), 1978. Graduation Thesis, University of Tokyo.
- [8] E. Segal-Halevi and B. R. Sziklai. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory*, 68(2):363–401, 2019.

INSTITUTE OF SCIENCE TOKYO, TOKYO, JAPAN
Email address: sumita@comp.isct.ac.jp

STABLE MATCHING IN PRACTICE: DAYCARE MATCHING MARKETS IN JAPAN

ZHAOHONG SUN

Classification AMS 2020:

Keywords: Matching, Stability, Fairness, Siblings, Constraint Programming

With the prevalence of dual-career households in recent years, the demand for daycare facilities in Japan, especially in metropolitan areas, has soared. Unfortunately, scarce space and insufficient teachers lead to a long waiting list each year, leaving numerous children unable to enroll in daycare centers. The waiting child problem becomes one of the critical social challenges nowadays. The allocation of children to daycare centers in Japan is not done on a first-come, first-served basis. Instead, the allocation process starts with families submitting applications to their local government office, which contains basic information such as children's age, guardians' health conditions and work schedule, as well as preferences over acceptable daycare centers. Each municipality adopts a unique scoring system that strictly prioritizes children. The scoring system is designed in a way that children who may have greater needs or face additional challenges have more chances of utilizing daycare services. Typically, children from low-income or single-parent households, and those whose guardians are suffering from diseases or disabilities take precedence over others. Some local governments formulate lottery rules for tie-breaking when children have identical priority scores. The allocation is then computed by a centralized matching algorithm that allows for both families' preferences over daycares and daycares' priorities over children.

The daycare matching process has faced criticism due to long waiting lists and the Japanese government has made considerable efforts to address this issue, including improving working conditions and increasing salaries for childcare workers to engage more people in early childhood education, and providing financial assistance for families in need to enlarge their options of affordable daycare centers. Although the number of children on the waiting list significantly decreased recently, the shortage of daycare facilities continues. This is because not all unmatched children are counted in the waiting list, such as those who live near daycare centers with vacant slots but are only willing to attend certain oversubscribed daycare centers and those whose parents have to suspend their jobs or extend their childcare leave. Thus, despite these measures, the waiting child problem remains a major social challenge and long waiting lists have a profound impact on young couples' careers and lives.

The daycare matching problem bears similarities with conventional two-sided matching problems such as college admissions and job hunting, where one side of the market consists of families who submit applications on behalf of their children and the other side consists of daycare centers with limited resources (e.g., room space, teachers). However,

	initial enrollments	# siblings			any joint preferences	transferable quotas
		2	3	any		
(McDermid and Manlove, 2010)	×	✓			×	×
(Ashlagi, Braverman, and Hassidim, 2014)	×	✓			✓	×
(Manlove, McBride, and Trimble, 2017)	×	✓			✓	×
(Suzuki et al., 2023)	✓	×			×	×
(Okumura, 2019)	×	×			×	✓
(Kamada and Kojima, 2023)	×	×			×	✓
(Dur, Morrill, and Phan, 2022)	×	✓			×	×
(Correa et al., 2022)	✓			✓	×	×
(Sun et al., 2023)	✓		✓		✓	×
This Work	✓			✓	✓	✓

TABLE 1. Comparison with some recent papers. ✓ indicates that the proposed approach in that paper is applicable to a particular feature.

the daycare matching market possesses three features that set it apart from classical matching models. These features include i) *transfers* (i.e., some children who are already enrolled prefer to be transferred to other daycare centers), ii) *siblings* (i.e., several children from the same family report joint preferences and only consent to an assignment if all of them are matched), and iii) *transferable quotas* (i.e., a daycare center may partition grades into grade groups and available spots can be used by any child within the same grade group). It is well-known that when couples exist, there may not exist any matching satisfying stability (Roth, 1984) (i.e., one of the most important solution concepts in matching theory) and determining whether there exists a stable matching is NP-complete (Ronn, 1990). The presence of these complexities poses more significant challenges.

The objective of this research is to develop a trustworthy algorithm to help municipalities tackle the waiting children problem. The key research question is *how to design and implement practical matching algorithms that minimize the number of children on the waiting list in a transparent, stable and computationally efficient manner*.

Our contributions are summarized as follows: Firstly, we formalize three particular features of the market and develop a comprehensive model that encompasses other important matching markets.

Secondly, we present some new computational complexity results, including that it is NP-complete to check the existence of feasible and individual rational matching that differs from the initial enrollments. Thirdly, we propose a practical algorithm based on constraint programming (CP) which is a powerful technique for coping with NP-hard problems. Fourthly, we evaluate the effectiveness of our algorithm by conducting experiments on real-world data sets and summarize interesting findings on the factors that could increase the number of matched children. Lastly, we release our implementation that could benefit future work in related domains.

RELATED WORK

There is a large body of literature on matching problems under preferences and we give a more detailed review of related work in Appendix. We next compare our work with some recent papers that also consider some of the features mentioned above. The daycare matching problem can be seen as a generalization of hospital-doctor matching

with couples, where pairs of doctors participate in the market and aim to secure a pair of positions (Kojima, Pathak, and Roth, 2013; Biró, Manlove, and McBride, 2014; Nguyen and Vohra, 2018). While some papers focus on school choice with initial enrollment, they do not take siblings into consideration (Hamada et al., 2017; Suzuki, Tamura, and Yokoo, 2018; Suzuki et al., 2023). On the other hand, two recent papers delve into school choice with siblings, but assuming restrictive preferences of families (Dur, Morrill, and Phan, 2022; Correa et al., 2022). Regarding the Japanese daycare matching problem, two prior works investigate transferable quotas but do not consider initial enrollments or siblings in their models (Okumura, 2019; Kamada and Kojima, 2023). Another study also examines the Japanese daycare matching market; however, they do not allow for transferable quotas and propose an algorithm suitable for a specific scenario with a maximum of three children per family (Sun et al., 2023). Several papers tackle the practical daycare matching problem in European countries, but it is important to note that the settings and objectives of these studies differ significantly from ours (Veski et al., 2017; Geitle et al., 2021; Reischmann, Klein, and Giegerich, 2021).

REFERENCES

- A. E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89(4): 748–780, 1999.
- A. E. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, Cambridge, UK, 1990.
- Z. Sun, Y. Takenami, D. Moriwaki, Y. Tomita, and M. Yokoo. Daycare matching in Japan: Transfers and siblings. In *Proceedings of the 37th AAI Conference on Artificial Intelligence, AAI 2023*, pages 14487–14495, 2023.
- Z. Sun, N. Yamada, Y. Takenami, D. Moriwaki, and M. Yokoo. Stable matchings in practice: A constraint programming approach. In *Proceedings of the 38th AAI Conference on Artificial Intelligence, AAI 2024*, pages 22377–22384, 2024.
- T. Suzuki, A. Tamura, K. Yahiro, M. Yokoo, and Y. Zhang. Strategyproof allocation mechanisms with endowments and m-convex distributional constraints. *Artificial Intelligence*, 315:103825, 2023.
- Abdulkadiroğlu, A.; and Sönmez, T. 2003. School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3): 729–747.
- Ashlagi, I.; Braverman, M.; and Hassidim, A. 2014. Stability in Large Matching Markets with Complementarities. *Operations Research*, 62(4): 713–732.
- Biró, P.; Irving, R. W.; and Schlotter, I. 2011. Stable matching with couples: an empirical study. *Journal of Experimental Algorithmics (JEA)*, 16: 1–2.
- Biró, P.; Manlove, D. F.; and McBride, I. 2014. The hospitals/residents problem with couples: Complexity and integer programming models. In *International Symposium on Experimental Algorithms*, 10–21. Springer.
- Correa, J.; Epstein, N.; Epstein, R.; Escobar, J.; Rios, I.; Aramayo, N.; Bahamondes, B.; Bonet, C.; Castillo, M.; Cristi, A.; Epstein, B.; and Subiabre, F. 2022. School choice in Chile. *Operations Research*, 70(2): 1066–1087.
- Dur, U.; Morrill, T.; and Phan, W. 2022. Family Ties: School Assignment with Siblings. *Theoretical Economics*, 17(1): 89–120.

- Gale, D.; and Shapley, L. S. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1): 9–15.
- Garey, M. R.; and Johnson, D. S. 1979. Computers and Intractability: A Guide to the Theory of NP-Completeness.
- Geitle, A. H.; Johnsen, Ø. K.; Ruud, H. F. E.; Fagerholt, K.; and Julsvoll, C. A. 2021. Kindergarten allocation in Norway: An integer programming approach. *Journal of the Operational Research Society*, 72(7): 1664–1673.
- Hamada, N.; Hsu, C.; Kurata, R.; Suzuki, T.; Ueda, S.; and Yokoo, M. 2017. Strategy-proof school choice mechanisms with minimum quotas and initial endowments. *Artificial Intelligence*, 249: 47–71.
- Hatfield, J. W.; and Milgrom, P. R. 2005. Matching with contracts. *American Economic Review*, 95(4): 913–935.
- Kamada, Y.; and Kojima, F. 2023. Fair Matching under Constraints: Theory and Applications. *The Review of Economic Studies*.
- Kennes, J.; Monte, D.; and Tumennasan, N. 2014. The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics*, 6(4): 362–406.
- Kojima, F.; Pathak, P. A.; and Roth, A. E. 2013. Matching with couples: Stability and incentives in large markets. *The Quarterly Journal of Economics*, 128(4): 1585–1632.
- Manlove, D. F.; McBride, I.; and Trimble, J. 2017. “Almost-stable” matchings in the Hospitals/Residents problem with Couples. *Constraints*, 22(1): 50–72.
- Marx, D.; and Schlotter, I. 2011. Stable assignment with couples: Parameterized complexity and local search. *Discret. Optim.*, 8(1): 25–40.
- McDermid, E. J.; and Manlove, D. F. 2010. Keeping partners together: algorithmic results for the hospitals/residents problem with couples. *Journal of Combinatorial Optimization*, 19(3): 279–303.
- Nguyen, T.; and Vohra, R. 2018. Near-feasible stable matchings with couples. *American Economic Review*, 108(11): 3154–69.
- Okumura, Y. 2019. School choice with general constraints: a market design approach for the nursery school waiting list problem in Japan. *The Japanese Economic Review*, 70(4): 497–516.
- Reischmann, T.; Klein, T.; and Giegerich, S. 2021. An iterative deferred acceptance mechanism for decentralized, fast and fair childcare assignment. *ZEW-Centre for European Economic Research Discussion Paper*, (21-095).
- Ronn, E. 1990. NP-complete stable matching problems. *Journal of Algorithms*, 11(2): 285–304.
- Roth, A. E. 1984. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 92(6): 991–1016.
- Veski, A.; Biró, P.; Pöder, K.; and Lauri, T. 2017. Efficiency and fair access in kindergarten allocation policy design. *The Journal of Mechanism and Institution Design*, 2(1): 57–104.
- Suzuki, T.; Tamura, A.; and Yokoo, M. 2018. Efficient allocation mechanism with endowments and distributional constraints. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 50–58.

KYUSHU UNIVERSITY, JAPAN

Email address: sunzhaohong1991@gmail.com

TRUTHFUL AND ALMOST ENVY-FREE MECHANISM OF ALLOCATING INDIVISIBLE GOODS: THE POWER OF RANDOMNESS

BIAOSHUAI TAO

Classification AMS 2020: 91A05, 91A06, 91A80

Keywords: Fair Division, Mechanism Design, Envy-Freeness

We study the problem of *fairly* and *truthfully* allocating m indivisible items to n agents with additive preferences. Specifically, we consider truthful mechanisms outputting allocations that satisfy EF_{-v}^{+u} , where, in an EF_{-v}^{+u} allocation, for any pair of agents i and j , agent i will not envy agent j if u items were added to i 's bundle and v items were removed from j 's bundle. Previous work easily indicates that, when restricted to deterministic mechanisms, truthfulness will lead to a poor guarantee of fairness: even with two agents, for any u and v , EF_{-v}^{+u} cannot be guaranteed by truthful mechanisms when the number of items is large enough. In this work, we focus on randomized mechanisms, where we consider *ex-ante* truthfulness and *ex-post* fairness. For two agents, we present a truthful mechanism that achieves EF_{-1}^{+0} (i.e., the well-studied fairness notion EF1). For three agents, we present a truthful mechanism that achieves EF_{-1}^{+1} (i.e., envy-freeness up to one good more-and-less proposed by Barman and Krishnamurthy (2019)). For n agents in general, we show that there exist truthful mechanisms that achieve EF_{-v}^{+u} for some u and v that depend only on n (not m).

We further consider fair and truthful mechanisms that also satisfy the standard efficiency guarantee: Pareto-optimality. We provide a mechanism that simultaneously achieves truthfulness, EF1, and Pareto-optimality for bi-valued utilities (where agents' valuation on each item is either p or q for some $p > q \geq 0$). For tri-valued utilities (where agents' valuations on each item belong to $\{p, q, r\}$ for some $p > q > r \geq 0$) and any u, v , we show that truthfulness is incompatible with EF_{-v}^{+u} and Pareto-optimality even for two agents.

This is a joint work with Xiaolin Bu.

1. INTRODUCTION

Fair division studies how to allocate a set of resources to a set of agents with heterogeneous preferences. In this paper, we study the fair division problem when resources are *indivisible* items. Specifically, we aim to fairly allocate m items to n agents, where each agent has her own valuations on those m items.

Among various fairness criteria, *envy-freeness* is arguably the most studied notion, which says that, for any pair of agents i and j , agent i should value her own allocated share weakly more than agent j 's, i.e., agent i does not envy agent j . However, when indivisible items are concerned, envy-free allocation may not exist (e.g., all the agents value the items equally, but m is not a multiple of n). It is then natural to define relaxations of envy-freeness that are tractable. The most popular line of research considers envy-freeness up to the addition or/and removal of a small number of items.

In particular, an allocation is “almost envy-free” if, for each pair of agents i and j , agent i will no longer envy agent j if a small number of items is (hypothetically) added to agent i ’s allocated bundle and a small number of items is (hypothetically) removed from agent j ’s allocated bundle. Among this type of relaxation, *envy-freeness up to one item* (EF1) receives the most significant attention. It is well-known that an EF1 allocation always exists, and it can be computed efficiently Budish (2011); Lipton et al. (2004).

When deploying a fair division algorithm in practice, agents may not honestly report their valuation preferences to the algorithm if they can benefit from strategic behaviors. This motivates the study of the fair division problem from the mechanism design point of view. Other than guaranteeing fairness, we would also like an algorithm, or a mechanism, to be *truthful*, where truth-telling is each agent’s dominant strategy. Unfortunately, it is known that truthfulness is incompatible with most of the meaningful fairness notions for deterministic mechanisms Lipton et al. (2004); Amanatidis et al. (2017, 2016); Caragiannis et al. (2009); Dobzinski et al. (2023), including those above-mentioned variants of envy-freeness Amanatidis et al. (2017). In particular, Amanatidis et al. (2017) give a characterization of truthful mechanisms with two agents. Their observation implies that no truthful mechanism can achieve envy-freeness even up to adding/removing an arbitrary number of items. Truthfulness and (almost) envy-freeness are compatible only for very restrictive valuation functions Halpern et al. (2020); Babaioff et al. (2021); Barman and Verma (2022); Christodoulou and Christoforidis (2024). Further, it is shown that under mild additional assumptions, the only deterministic truthful mechanism is serial/sequential-quota dictatorship Pápai (2000, 2001); Bouveret et al. (2023); Babaioff and Morag (2024), where each agent is asked to take a predefined number of items one by one¹. Such mechanisms clearly lack fairness guarantees.

In this paper, we seek to resolve the incompatibility of truthfulness and fairness by applying *randomness* in mechanisms. We aim to design *randomized mechanisms* that is *truthful in expectation*—truth-telling maximizes each agent’s *expected utility*, and meanwhile guaranteeing that every allocation possibly output by the mechanism is almost envy-free. Although the use of randomness to achieve truthfulness has been proven successful in other problems Mossel and Tamuz (2010); Dobzinski and Dughmi (2013), our understanding of the power of randomness for fair division of indivisible items is still limited, especially for envy-based fairness notions.

2. OUR RESULTS

In this paper, we mainly focus on envy-based fairness notions, and we show that randomized mechanisms provide significantly better fairness guarantees than their deterministic counterpart.

For $n = 2$ agents, we provide a simple truthful randomized mechanism based on the equal division rule that outputs EF1 allocations. We show that the equal division rule fails to guarantee the EF1 fairness property for $n = 3$. For $n = 3$ agents, we provide a truthful randomized mechanism that outputs EF_{-1}^{+1} allocations (envy-freeness up to adding and removing one item, also called envy-freeness up to one good more-and-less

¹In serial dictatorship, the order of the agents is predefined and independent of the agents’ preferences. In sequential dictatorship, the order of the agents depends on the agents’ preferences.

proposed by Barman and Krishnamurthy (2019)). This is achieved by some carefully designed fractional allocation rule, and the decomposition to EF_{-1}^{+1} allocations applies a series of techniques including proper coloring of regular bipartite graphs and rounding of vertex solutions of linear programs. For general numbers of agents, we design two mechanisms based on the equal division rule: a truthful randomized mechanism that outputs $EF_{-(n-1)}^{+(n-1)^2}$ allocations and a truthful randomized mechanism that outputs allocations simultaneously satisfying two share-based fairness notions—PROP1 and $\frac{1}{n}$ -MMS.

Finally, we study efficient randomized truthful mechanisms that satisfy Pareto-optimality in addition to fairness. We show that the truthful EF1 Pareto-optimal mechanism for binary valuations Halpern et al. (2020); Babaioff et al. (2021); Barman and Verma (2022) generalizes to the bi-valued valuations (where an agent’s value to an item can only take two values p or q) if randomization is allowed. This is complemented by the impossibility result that, for any u and v , EF_{-v}^{+u} is incompatible with Pareto-optimality for randomized truthful mechanisms for two agents with tri-valued valuations.

REFERENCES

- Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On truthful mechanisms for maximin share allocations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 31–37, 2016.
- Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 545–562, 2017.
- Moshe Babaioff and Noam Manaker Morag. On truthful mechanisms without pareto-efficiency: Characterizations and fairness. *arXiv preprint arXiv:2411.11131*, 2024.
- Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5119–5126, 2021.
- Siddharth Barman and Sanath Kumar Krishnamurthy. On the proximity of markets with integral equilibria. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1748–1755, 2019.
- Siddharth Barman and Paritosh Verma. Truthful and fair mechanisms for matroid-rank valuations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4801–4808, 2022.
- Sylvain Bouveret, Hugo Gilbert, Jérôme Lang, and Guillaume Méroué. Thou shalt not pick all items if thou are first: of strategyproof and fair picking sequences. *arXiv preprint arXiv:2301.06086*, 2023.
- Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On low-envy truthful allocations. In *Algorithmic Decision Theory: First*

- International Conference, ADT 2009, Venice, Italy, October 20-23, 2009. Proceedings 1*, pages 111–119. Springer, 2009.
- George Christodoulou and Vasilis Christoforidis. Fair and truthful allocations under leveled valuations, 2024. URL <https://arxiv.org/abs/2407.05891>.
- Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. *SIAM Journal on Computing*, 42(6):2287–2304, 2013.
- Shahar Dobzinski, Sigal Oren, and Jan Vondrak. Fairness and incentive compatibility via percentage fees. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 517–535, 2023.
- Daniel Halpern, Ariel D Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Web and Internet Economics: 16th International Conference, WINE 2020, Beijing, China, December 7–11, 2020, Proceedings 16*, pages 370–383. Springer, 2020.
- Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 125–131, 2004.
- Elchanan Mossel and Omer Tamuz. Truthful fair division. In *International Symposium on Algorithmic Game Theory*, pages 288–299. Springer, 2010.
- Szilvia Pápai. Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5:91–105, 2000.
- Szilvia Pápai. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory*, 3(3):257–271, 2001.

NO. 800 DONG CHUAN ROAD, MINHANG DISTRICT, SHANGHAI 200240 P. R. CHINA
Email address: bstao@sjtu.edu.cn

TEMPORAL FAIR DIVISION OF INDIVISIBLE GOODS AND CHORES

NICHOLAS TEH

Classification AMS 2020: 91B14, 91B32

Keywords: Fair Division, Envy-Freeness, Temporal Fairness, Indivisible Items

This talk is based on joint work with Edith Elkind, Alexander Lam, Mohamad Latifian, and Tzeh Yuan Neoh [14].

1. PRELIMINARIES

For each positive integer k , let $[k] := \{1, \dots, k\}$. We consider the problem of fairly allocating indivisible items to agents over multiple rounds. Let an instance of the problem be denoted by $\mathcal{I} = \langle N, T, \{O_t\}_{t \in [T]}, \mathbf{v} = (v_1, \dots, v_n) \rangle$, in which we have a set of agents $N = [n]$ and a set O of m items which arrive over T rounds, and are to be allocated to the agents. For each $t \in [T]$, we denote the set of items that arrive at round t by O_t , and define the cumulative set of items that arrived in rounds $1, \dots, t$ by $O^t := \bigcup_{\ell \in [t]} O_\ell$. Note that $O = O^T$.

We assume that each agent $i \in N$ has an additive *valuation function* $v_i : 2^O \rightarrow \mathbb{R}$ over the items, i.e., for $S \subseteq O$, $v_i(S) = \sum_{o \in S} v_i(\{o\})$. For notational convenience, we write $v_i(o)$ instead of $v_i(\{o\})$ for a single item $o \in O$, and v instead of v_i when valuation functions are identical. Denote $\mathbf{v} = (v_1, \dots, v_n)$ as the *valuation profile*. In this work, we consider two cases: *goods allocation* where for each $i \in N$ and $o \in O$, $v_i(o) \geq 0$, and *chores allocation* where for each $i \in N$ and $o \in O$, $v_i(o) \leq 0$. For clarity, we use g instead of o and refer to items as goods when explicitly referring to the goods setting, and c instead of o and refer to the items as chores when considering the chores setting.

An allocation $\mathcal{A} = (A_1, \dots, A_n)$ of items in O to the agents is an ordered partition of O , i.e. for $i, j \in N$, $A_i \cap A_j = \emptyset$ and $\bigcup_{i \in N} A_i = O$. In addition, for $t \in [T]$ we denote the allocation after round t by $\mathcal{A}^t = (A_1^t, \dots, A_n^t)$ where $A_i^t = A_i \cap O^t$. For $t < T$, we sometimes refer to \mathcal{A}^t as a *partial allocation*. Note that $\mathcal{A} = \mathcal{A}^T$. Our goal is to find an allocation that is *fair* after each round.

Definition 1.1. *In a goods (resp. chores) allocation instance, an allocation $\mathcal{A} = (A_1, \dots, A_n)$ is said to be EF1 if for all pairs of agents $i, j \in N$, there exists a good $g \in A_j$ (resp. chore $c \in A_i$) such that $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ (resp. $v_i(A_i \setminus \{c\}) \geq v_i(A_j)$).*

In this work we target fairness in a *cumulative* sense, introducing the notion of *temporal envy-freeness up to one item (TEF1)* which requires that at every round prefix, the cumulative allocation of items that have arrived so far satisfies EF1.

Definition 1.2 (Temporal EF1). *For any $t \in [T]$, an allocation $\mathcal{A}^t = (A_1^t, \dots, A_n^t)$ is said to be temporal envy-free up to one item (TEF1) if for all $t' \leq t$, the allocation $\mathcal{A}^{t'}$ is EF1.*

However, the possible non-existence of TEF1 allocations has been shown in the general goods allocation setting, as He et al. (2019) illustrated using a counterexample with 3 agents and 23 items, which can be generalized to $n > 3$ agents.

2. EXISTENCE OF TEF1 ALLOCATIONS IN SPECIAL CASES

In this section we identify several restricted classes of instances under which a TEF1 allocation is guaranteed to exist.

He et al. (2019) showed the existence of a TEF1 allocation for goods when $n = 2$, by presenting a polynomial-time algorithm that returns such an allocation. We extend this result to the case of chores, and combine the results into a single theorem, as follows.

Theorem 2.1. *For $n = 2$, a TEF1 allocation for goods or chores exists and can be computed in polynomial time.*

The next setting we consider is one where items can be divided into two *types*, and each agent values all items of a particular type equally. Formally, let $S_1, S_2 \subseteq O$ be a partition of the set of items, so that $S_1 \cap S_2 = \emptyset$, and $S_1 \cup S_2 = O$. Then, for any $r \in \{1, 2\}$, two items $o, o' \in S_r$, and agent $i \in N$, we have that $v_i(o) = v_i(o')$.

Theorem 2.2. *When there are two types of items, a TEF1 allocation for goods or chores exists and can be computed in polynomial time.*

Another setting we consider is one where agents have *generalized binary valuations* (also known as *restricted additive valuations* [1, 8]). Formally, we say that agents have *generalized binary valuations* if for every agent $i \in N$ and item $o_j \in O$, $v_i(o_j) \in \{0, p_j\}$, where $p_j \in \mathbb{R} \setminus \{0\}$.

Theorem 2.3. *When agents have generalized binary valuations, a TEF1 allocation for goods or chores exists and can be computed in polynomial time.*

3. HARDNESS RESULTS FOR TEF1 ALLOCATIONS

The non-existence of TEF1 goods allocations for $n \geq 3$ prompts us to explore whether we can determine if a given instance admits a TEF1 allocation for goods. Unfortunately, we show that this problem is NP-hard, with the following result.

Theorem 3.1. *Given an instance with goods and $n \geq 3$, determining whether there exists a TEF1 allocation is NP-hard.*

However, we note that the approach used in proving the above result cannot be extended to show hardness for the setting with chores. Nevertheless, we are able to show a similar, though weaker, intractability result for the case of chores in general.

Theorem 3.2. *For every $t \in [T]$, given any partial TEF1 allocation \mathcal{A}^t for chores, deciding if there exists an allocation \mathcal{A} that is TEF1 is NP-hard.*

4. COMPATIBILITY BETWEEN TEF1 AND EFFICIENCY

In traditional fair division, many papers have focused on the existence and computation of fair and efficient allocations for goods or chores, with a particular emphasis on simultaneously achieving EF1 and *Pareto-optimality (PO)*. In this section, we explore the compatibility between TEF1 and PO. We begin by defining PO as follows.

Definition 4.1 (Pareto-optimality). *We say that an allocation \mathcal{A} is Pareto-optimal (PO) if there does not exist another allocation \mathcal{A}' such that for all $i \in N$, $v_i(A'_i) \geq v_i(A_i)$, and for some $j \in N$, $v_j(A'_j) > v_j(A_j)$. If such an allocation \mathcal{A}' exists, we say that \mathcal{A}' Pareto-dominates \mathcal{A} .*

Observe that for any \mathcal{A} that is PO, any partial allocation \mathcal{A}^t for $t \leq [T]$ is necessarily PO as well. We demonstrate that PO is incompatible with TEF1 in this setting, even under very strong assumptions (of two agents and two types of items). Despite this non-existence result, one may still wish to obtain a TEF1 and PO outcome when the instance admits one. However, we show that this is not computationally tractable.

Theorem 4.2. *Determining whether there exists a TEF1 allocation that is PO for goods or chores is NP-hard, even when $n = 2$.*

5. CONCLUSION

Numerous potential directions remain for future work, including revisiting variants of the standard fair division model. Examples include studying the existence (and polynomial-time computability) of allocations satisfying a temporal variant of the weaker *proportionality up to one item* property, which would be implied by EF1; studying group fairness [2, 4, 7, 11, 19, 23]; considering the more general class of *submodular* valuations [18, 22, 24, 25, 26]; examining the *house allocation* model where each agent gets a single item [10, 17], which was partially explored in [21], or even looking at more general settings with additional size constraints [6, 5, 12]. It would also be interesting to extend our results, which hold for the cases of goods and chores separately, to the more general case of mixed manna, in which items can be simultaneously viewed as goods by some agents and as chores by others (see, e.g., [3]). The concept of achieving cumulative fairness while knowing all future information could also be applied to other temporal models in the social choice literature [9, 13, 16, 15, 20, 27].

REFERENCES

- [1] Hannaneh Akrami, Rojin Rezvan, and Masoud Seddighin. An EF2X allocation protocol for restricted additive valuations. In *IJCAI*, pages 17–23, 2022.
- [2] Martin Aleksandrov and Toby Walsh. Group envy freeness and group pareto efficiency in fair division with indivisible items. In *KI*, pages 57–72, 2018.
- [3] H. Aziz, I. Caragiannis, A. Igarashi, and T. Walsh. Fair allocation of indivisible goods and chores. In *IJCAI*, pages 53–59, 2019.
- [4] Haris Aziz and Simon Rey. Almost group envy-free allocation of indivisible goods and chores. In *IJCAI*, pages 39–45, 2020.
- [5] S. Barman, A. Khan, S. Shyam, and K. V. N. Sreenivas. Guaranteeing envy-freeness under generalized assignment constraints. In *EC*, pages 242–269, 2023.
- [6] Siddharth Barman, Arindam Khan, Sudarshan Shyam, and K.V.N. Sreenivas. Finding fair allocations under budget constraints. In *AAAI*, pages 5481–5489, 2023.
- [7] N. Benabbou, M. Chakraborty, E. Elkind, and Y. Zick. Fairness towards groups of agents in the allocation of indivisible items. In *IJCAI*, pages 95–101, 2019.

- [8] Franklin Camacho, Rigoberto Fonseca-Delgado, Ramón Pino Pérez, and Guido Tapia. Generalized binary utility functions and fair allocations. *Mathematical Social Sciences*, 121:50–60, 2023.
- [9] Nikhil Chandak, Shashwat Goel, and Dominik Peters. Proportional aggregation of preferences for sequential decision making. In *38th*, pages 9573–9581, 2024.
- [10] D. Choo, Y. H. Ling, W. Suksompong, N. Teh, and J. Zhang. Envy-free house allocation with minimum subsidy. *Operations Research Letters*, 54:107103, 2024.
- [11] Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Vaughan. Group fairness for the allocation of indivisible goods. In *AAAI*, pages 1853–1860, 2019.
- [12] E. Elkind, A. Igarashi, and N. Teh. Fair division of chores with budget constraints. In *SAGT*, pages 55–71, 2024.
- [13] Edith Elkind, Sonja Kraiczy, and Nicholas Teh. Fairness in temporal slot assignment. In *SAGT*, pages 490–507, 2022.
- [14] Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. Temporal fair division of indivisible items. In *AAMAS*, 2025.
- [15] Edith Elkind, Tzeh Yuan Neoh, and Nicholas Teh. Temporal elections: Welfare, strategyproofness, and proportionality. In *ECAI*, 2024.
- [16] Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. Temporal fairness in multiwinner voting. In *38th*, pages 22633–22640, 2024.
- [17] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. Envy-freeness in house allocation problems. *Mathematical Social Sciences*, 101:104–106, 2019.
- [18] M. Ghodsi, M. HajiAghayi, M. Seddighin, S. Seddighin, and H. Yami. Fair allocation of indivisible goods: Beyond additive valuations. *Art. Int.*, 303:103633, 2022.
- [19] Maria Kyropoulou, Warut Suksompong, and Alexandros Voudouris. Almost envy-freeness in group resource allocation. In *IJCAI*, pages 400–406, 2019.
- [20] Andrew Mackenzie and Vilmos Komornik. Fairly taking turns. *Games and Economic Behavior*, 142:743–764, 2023.
- [21] K. J. Micheel and A. Wilczynski. Fairness in repeated house allocation. In *AAMAS*, pages 2390–2392, 2024.
- [22] L. Montanari, U. Schmidt-Kraepelin, W. Suksompong, and N. Teh. Weighted envy-freeness for submodular valuations. In *AAAI*, pages 9865–9873, 2024.
- [23] J. Scarlett, N. Teh, and Y. Zick. For one and all: Individual and group fairness in the allocation of indivisible goods. In *AAMAS*, pages 2466–2468, 2023.
- [24] W. Suksompong and N. Teh. On maximum weighted nash welfare for binary valuations. *Math Soc. Sci.*, 117:101–108, 2022.
- [25] W. Suksompong and N. Teh. Weighted fair division with matroid-rank valuations: Monotonicity and strategyproofness. *Math Soc. Sci.*, 126:48–59, 2023.
- [26] Gilad Ben Uziah and Uriel Feige. On fair allocation of indivisible goods to submodular agents. *arXiv preprint arXiv:2303.12444*, 2023.
- [27] Valentin Zech, Niclas Boehmer, Edith Elkind, and Nicholas Teh. Multiwinner temporal voting with aversion to change. In *ECAI*, 2024.

UNIVERSITY OF OXFORD, UK

Email address: nicholas.teh@cs.ox.ac.uk

FAIR INTERVAL SCHEDULING OF INDIVISIBLE CHORES

ROHIT VAISH

Classification AMS 2020:

Keywords: Fair Division, Indivisible Chores, Scheduling, Conflict Constraints, Envy-Freeness

This talk considers a problem in the emerging area of “fair division under constraints” [1].

As a motivating example, consider a scheduling problem where each task, or *chore*, has a designated start and finish time, and each agent can perform at most one chore at any time (in other words, agents cannot multitask). The agents have possibly differing values for subsets of chores, and should be allocated disjoint subsets of chores so that the final allocation is fair and economically efficient. In our model, the chores are represented as vertices of an interval graph whose edges capture temporal overlaps between the chores. A feasible allocation corresponds to a subpartition of the vertices into independent subsets. Note that due to the conflict constraints, some chores may have to be left unassigned.

An allocation is deemed “fair” if it is *envy-free up to one chore* (EF1), which means that any pairwise envy can be eliminated by the hypothetical removal of some chore from the envious agent’s bundle. We also study a strengthening of EF1 called *envy-freeness up to any chore* (EFX), wherein pairwise envy can be eliminated by dropping *any* individual chore from the envious agent’s bundle. On the economic efficiency front, we consider two criteria: *maximality* and *Pareto optimality*. Maximality requires that any unallocated chore generates a conflict when assigned to any agent. Pareto optimality, on the other hand, stipulates that an allocation, in addition to being maximal, should not be Pareto dominated by another maximal allocation.

We note that EFX is not compatible with the weaker efficiency notion of maximality, and EF1 is incompatible with Pareto optimality. This leaves us with only the combination of EF1 and maximality as the potential candidates for universal existence. Notably, several techniques from unconstrained fair division that achieve EF1, such as round-robin and top-trading envy-cycle algorithms, fail to maintain EF1 in the constrained setting.

Towards showing the existence of EF1 and maximal allocations, we contribute two results:

- We show that for two agents, an EF1 and maximal schedule always exists and can be efficiently computed when the conflicts are represented by an interval graph. This result uses an alternate coloring technique and holds even under monotone valuations.
- We show that for an arbitrary number of agents with identical additive valuations, an EF1 and maximal allocation always exists when the conflict graph is a path graph. This result is largely non-algorithmic and uses an application of

the “cycle-plus-triangles” theorem (originally conjectured by Erdős) for achieving approximate envyfreeness.

The talk also highlights several open problems and directions for future work.

Based on joint work with Sarfaraz Equbal, Rohit Gurjar, Yatharth Kumar, Swaprava Nath, and Raghuvansh Saxena. An early version of the paper is available as [2].

REFERENCES

- [1] Warut Suksompong. Constraints in Fair Division. *ACM SIGecom Exchanges*, 19(2), pages 46-61, 2024. <https://dl.acm.org/doi/abs/10.1145/3505156.3505162>
- [2] Sarfaraz Equbal, Rohit Gurjar, Yatharth Kumar, Swaprava Nath, and Rohit Vaish. Fair Interval Scheduling of Indivisible Chores. *arXiv*, 2024. <https://arxiv.org/abs/2402.04353>

INDIAN INSTITUTE OF TECHNOLOGY DELHI

Email address: rvaish@iitd.ac.in

SIX CANDIDATES SUFFICE TO WIN A VOTER MAJORITY

ADRIAN VETTA

(This talk concerned joint work with Moses Charikar, Alexandra Lassota, Prasanna Ramakrishnan, and Kangning Wang)

Keywords: Social Choice, Voting, Condorcet Winner, Condorcet Winning Sets, Committee Selection.

Abstract: A cornerstone of social choice theory is Condorcet’s paradox which says that in an election where n voters rank m candidates it is possible that, no matter which candidate is declared the winner, a majority of voters would have preferred an alternative candidate. Instead, can we always choose a small *committee* of winning candidates that is preferred to any alternative candidate by a majority of voters?

Elkind, Lang, and Saffidine raised this question and called such a committee a *Condorcet winning set*. They showed that winning sets of size 2 may not exist, but sets of size logarithmic in the number of candidates always do. In this work, we show that Condorcet winning sets of size 6 always exist, regardless of the number of candidates or the number of voters. More generally, we show that if $\frac{\alpha}{1-\ln \alpha} \geq \frac{2}{k+1}$, then there always exists a committee of size k such that less than an α fraction of the voters prefer an alternate candidate. These are the first nontrivial positive results that apply for all $k \geq 2$.

Our proof uses the probabilistic method and the minimax theorem, inspired by recent work on *approximately stable committee selection*. We construct a distribution over committees that performs sufficiently well (when compared against any candidate on any small subset of the voters) so that this distribution must contain a committee with the desired property in its support.

1. INTRODUCTION

Voting is a versatile model for the aggregation of individual preferences to reach a collective decision. Disparate situations, such as constituents choosing representatives, organizations hiring employees, judges choosing prize winners, and even friends choosing games to play, can all be understood as a group of voters choosing from a pool of candidates. Voting theory seeks to understand how winning candidates can be selected in a fair and representative manner.

One of the longest known challenges with voting is *Condorcet’s paradox*, discovered by Nicolas de Condorcet around the French Revolution [2].¹ The paradox is that in an election where voters have ranked preferences over candidates, the preferences of the “majority” can be contradictory — no matter which candidate is declared the winner, a majority of the voters would have preferred another candidate. In fact, the contradiction

¹It is plausible that in early academic explorations of voting, 13th-century philosopher Ramon Llull had already discovered the possibility of this paradoxical situation [8, 6].

can be even more dramatic, with “majority” replaced by a fraction arbitrarily close to 1. An illustrative example is when the voters have cyclic preferences as, for example, displayed in 1.

v_1	v_2	v_3	v_4	v_5	v_6
1	2	3	4	5	6
2	3	4	5	6	1
3	4	5	6	1	2
4	5	6	1	2	3
5	6	1	2	3	4
6	1	2	3	4	5

TABLE 1. An election where voters have cyclic preferences. The column headed with v_i represents the i th voter’s ranking of the candidates (labeled $1, 2, \dots, 6$ from top to bottom). For each candidate, another candidate is preferred by every voter except one.

Though it is impossible to always find a single candidate that is always preferred over the others by a majority (called a *Condorcet winner*), one hope is that relaxations of this condition are still possible to achieve. A natural relaxation arises in the setting of *committee selection*, where rather than choosing a single winner, the goal is to choose a *committee* of k winners. For example, a political system may have districts with multiple representatives, organizations may make many hires at once, and friends might play more than one game in an evening. Another view is that committee selection can be used as an filtering step in a process with more than one round, like primaries or runoffs, choosing interviewees for a position, or nominations for a prize.

In this context, Elkind, Lang, and Saffidine [3] asked: is it always possible to find a small committee of candidates such that no other candidate is preferred by a majority of voters over each member of the committee? They called this committee-analogue of a Condorcet winner a *Condorcet winning set*, and defined the *Condorcet dimension* of an election as the size of its smallest Condorcet winning set. For example, the election depicted in 1 has Condorcet dimension 2, since any pair of diametrically opposite candidates such as $\{3, 6\}$ would be a Condorcet winning set. More generally, [3] raised the following question for an arbitrary threshold of α in place of $\frac{1}{2}$, and a target committee size k .

Question 1.1 ([3]). *A committee S is α -undominated if for all candidates $a \notin S$, less than an α fraction of voters prefer a over each member of S . For what values of $k \in \mathbb{Z}^+$ and $\alpha \in (0, 1]$ does every election have an α -undominated committee of size k ?*

In particular, we would like to know, for each k , what is the smallest α for which α -undominated committees of size k always exist (and, equivalently, for each α , the smallest k such that these committees always exist).

Condorcet’s paradox (or rather, its aforementioned generalization) shows that for $k = 1$ and any α bounded away from 1, there are elections with no α -undominated singleton candidates. For the threshold of $\alpha = \frac{1}{2}$, [3] constructed instances with Condorcet dimension 3 by taking the Kronecker product of two elections with cyclic

preferences. This construction can be easily extended to give a lower bound of $\frac{2}{k+1}$ on the smallest α such that there always exists an α -undominated committee of size k . They also showed that an election with m candidates has Condorcet dimension at most $\lceil \log_2 m \rceil$; to see this, note that some candidate beats a majority of the other candidates, so we can iteratively add such a candidate to our committee and remove all the candidates that it beats.

1.1. Our Contributions. We prove that every election has Condorcet dimension at most 6. This result is a corollary of our main theorem, which gives a nontrivial existence result for α -undominated committees of size $k \geq 2$. We note that the final result we prove is stronger, but the approximation below is easier to get a handle on.

Theorem 1.2. *If $\frac{\alpha}{1-\ln \alpha} \geq \frac{2}{k+1}$, then in any election, there exists an α -undominated committee of size k .*

For the specific threshold of $\alpha = \frac{1}{2}$, 1.2 applies as long as $k \geq 3 + 4 \ln 2 \approx 5.77$, and so any election has Condorcet dimension at most 6 (which is not far from the lower bound of 3). Taking $k = 2$, 1.2 implies that there always exists a pair of candidates such that no third candidate is preferred by more than roughly 80% of the voters. Even replacing 80% with 99%, this was previously unknown.

These results show that just by having a few winners instead of one, the most dramatic failures of Condorcet's paradox are avoidable. We emphasize that these results hold for *any election*, regardless of the number of voters, the number of candidates, or the preferences that the voters have over candidates.

Our starting point for proving Theorem 1.2 is the observation that Question 1.1 is closely linked to the problem of *approximate stability* in committee selection [7]. The principle behind stability is that a subset of voters should have control over a subset of the committee of proportional size. That is, a committee of size k is *stable* (also referred to as *in the core* [9, 4, 5]) if the fraction of voters that prefers any committee of size k' is less than $\frac{k'}{k}$. We note that in this setting, voters have preferences over *committees* rather than candidates. This more expressive space of preferences gives it the power to model a wide variety of preference structures, such as approval voting and participatory budgeting.

Unfortunately, in many settings, stable committees do not always exist. To remedy this, [7] put forth the following approximate notion of stability, and showed the surprising result that for any monotone preference structure and any k , a 32-stable committee of size k exists.

Definition 1.3 (Approximately stable committees [7]). *A committee S of k candidates is c -stable if for any committee S' of size k' , the fraction of voters that prefers S' over S is less than $c \cdot \frac{k'}{k}$.*

Consider the natural preference order over committees induced by rankings over candidates, where v prefers S' over S if and only if she prefers her favorite candidate in S' over her favorite in S . A simple observation shows that a committee of size k is c -stable if and only if it is $\frac{c}{k}$ -undominated. For this ranked preference structure, the constant of 32 in the result of [7] can be improved to 16 using the existence of stable lotteries for these preferences [1]. Then, as a black box, [7] implies that $\frac{16}{k}$ -undominated committees of size k always exist, which in turn implies that we can

always find Condorcet winning sets of size at most 32. Since this conclusion follows easily from [7], we attribute the first constant upper bound on the size of Condorcet winning sets to their work.

One can interpret the approximately stable committee problem as a version of Question 1.1 focused on the asymptotics of α as the committee size k grows large. For this purpose, [7] implies a result that is optimal up to a constant factor, but it says nothing nontrivial for committees of size at most 16. In contrast, Theorem 1.2 gives results even for $k = 2$, and outperforms the bound implied by [7] for $k \leq 1.75 \times 10^4$, despite only implying the existence of $O(\log k)$ -stable committees.

Nonetheless, we show that our techniques can be applied to the asymptotic setting as well, giving an improvement over [7].

Theorem 1.4. *In any election, there exists a $\frac{9.8217}{k}$ -undominated committee of size k .*

As a corollary, Theorem 1.4 implies the existence of 9.8217-stable committees for preferences induced by rankings over candidates. We note that Theorem 1.4 improves Theorem 1.2 for $k \geq 496$.

REFERENCES

- [1] Y. Cheng, Z. Jiang, K. Munagala and K. Wang, Group Fairness in Committee Selection *ACM Transactions on Economics and Computation*, **8(4)**, 23:1–18, 2020.
- [2] N. de Condorcet, Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. *Imprimerie royale*, 1785.
- [3] E. Elkind, J. Lang and A. Saffidine, Condorcet winning sets. *Social Choice and Welfare*, **44(3)**, pp493–517, 2015.
- [4] D. Foley, Lindahl's Solution and the Core of an Economy with Public Goods *Econometrica*, pp66–72, 1970.
- [5] B. Fain, K. Munagala and N. Shah, Fair Allocation of Indivisible Public Goods *Proceedings of ACM Conference on Economics and Computation (EC)*, pp575–592, 2018.
- [6] G. Hägele and F. Pukelsheim, Lulls Schriften zu Wahlverfahren. *Augsburg Universitätsbibliothek und Institut fr Mathematik*, 2000.
- [7] Z. Jiang, K. Munagala and K. Wang, Approximately stable committee selection *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp463–472, 2020.
- [8] R. Llull. Artificio electionis personarum. Manuscript, 1274 –1283
- [9] H. Scarf. The core of an N person game *Econometrica*, pp50–69, 1967.

MCGILL UNIVERSITY, MONTRÉAL, CANADA
E-mail address: adrian.vetta@mcgill.ca

PRESENTATION REPORT: FAIR ALLOCATION OF CHORES WITH SUBSIDY

XIAOWEI WU

ABSTRACT. The fair allocation problem has gained significant attention recently in the fields of theoretical computer science, artificial intelligence, and economics. In this presentation, I will discuss our latest research on ensuring fairness for the allocation of chores using subsidies. We consider the allocation of m indivisible chores among n agents with subsidies. Specifically, we focus on scenarios where agents have additive cost functions and assume that the maximum cost of an item to an agent can be offset by one dollar, we show that a total subsidy of $n/4$ dollars is sufficient to achieve a proportional allocation. Furthermore, we prove that $n/4$ is the minimum necessary subsidy, as there exists an instance with n agents where any proportional allocation requires at least $n/4$ dollars in subsidies. Additionally, we explore the weighted case and show that a total subsidy of $n/3$ dollars is sufficient to ensure weighted proportionality.

Classification AMS 2020:

- 91-08 Computational methods for problems pertaining to game theory, economics, and finance
- 91B14 Social choice
- 91B32 Resource and cost allocation (including fair division, apportionment, etc.)

Keywords: Fair Allocation of Chores, Proportionality, Allocation with Subsidy

1. INTRODUCTION

We study the fair allocation problem of allocating a set of m indivisible items M to a group of n heterogeneous agents N . When all items M give positive values to all agents, the problem refers to the fair allocation of *goods*, which represents the situation of distributing resources or public goods. Analogous but contrary, when all items M give negative values (positive costs) to all agents, the problem refers to the fair allocation of *chores*, which covers the task allocation among employees. In this paper, we mainly focus on the allocation of chores, where each agent i has a cost function $c_i : 2^M \rightarrow \mathbb{R}^+ \cup \{0\}$, but our main result also applies to the allocation of goods. We consider the general weighted setting in which each agent i has a $w_i > 0$ that represents her obligation to undertake the chores. We normalize the weights of agents such that $\sum_{i \in N} w_i = 1$. Traditionally, the fair allocation problem considers how to fairly allocate the items into n bundles (X_1, \dots, X_n) , while each agent receives exactly one bundle that guarantees some fairness criteria for her. We say that agent i has cost $c_i(S)$ for bundle $S \subseteq M$. One natural and well-studied fairness notion is *proportionality* (PROP) [33]: an allocation is called weighted proportional (WPROP) [31] if for all agent i , $c_i(X_i) \leq w_i \cdot c_i(M)$, where we refer to $w_i \cdot c_i(M)$ as the proportional share of agent i . Unfortunately, when items are indivisible, PROP allocations are not guaranteed to exist, e.g., considering allocating a single item to two agents. One possible way to circumvent this non-existence result is by introducing money to eliminate the inevitable unfairness. Maskin [30] first proposed the setting

called fair allocation with money that allows each agent to receive a subsidy $s_i \geq 0$ to eliminate unfairness, under which the objective is to minimize the total subsidized money. Halpern and Shah [24] considered the fair allocation of goods with money, for another well-known fairness notion, called *envy-freeness* [22]. Assuming that each item has a value of at most 1 to each agent, they conjecture that a total subsidy of $n-1$ suffices to guarantee envy-freeness for the allocation of goods. The conjecture was later verified by Brustle et al. [13]. Very recently, Aziz et al. [4] considered the weighted setting and characterized the property of guaranteeing weighted envy-free allocation with subsidy.

1.1. Our Results for the Unweighted Case. We aim to compute an allocation (X_1, \dots, X_n) and subsidies (s_1, \dots, s_n) such that $c_i(X_i) - s_i \leq w_i \cdot c_i(M)$ for all agent $i \in N$, with a small amount of total subsidy $\|s\|_1 = \sum_{i \in N} s_i$. When all agents have equal weights ($w_i = 1/n$ for all $i \in N$), we propose an algorithm that computes a proportional allocation with a total subsidy of at most $n/4$ [35]. Our algorithm is based on rounding a fractional allocation returned by the Moving Knife Algorithm, where we use two rounding schemes: Up Rounding and Threshold Rounding (Greedy).

1.2. Our Results for the Weighted Case. For the weighted case, we revisit the fractional bid-and-take algorithm proposed by our recent work [26] and devise a new rounding scheme. We characterize the structure of the fractional allocations returned by FBAT, by introducing the *item-sharing* graph where agents are nodes and fractional items are edges. We show that the item-sharing graph of the allocation returned by FBAT is a directed tree¹. We introduce a general rounding framework based on tree splitting, which allows us to split the directed tree graph into canonical components and round each component independently. We show that there exists a rounding scheme that guarantees weighted proportionality with a total subsidy of at most $n/3 - 1/6$.

1.3. Other Related Works. Other than proportionality, another well-studied fairness notion is *envy-freeness* (EF) [22], that is, no agent wants to exchange her bundle of items with another agent to improve her utility. Since both EF and PROP are not guaranteed to exist when items are indivisible, a line of literature focused on some relaxed fairness notions. *Envy-freeness up to one item* (EF1) [28] and *envy-freeness up to any item* (EFX) [15] are two widely studied relaxations of EF, which require that the envy between any two agents can be eliminated by removing some item; any item respectively. Similarly, we have *proportionality up to one item* (PROP1) [18, 9] and *proportionality up to any item* (PROPX) [6] to be two relaxations of PROP. In addition, *maximin share* (MMS) [14] is another popular relaxation of PROP.

Weighted Setting. Motivated by real-world applications where agents are usually not equally obliged, Chakraborty et al. [16] proposed the *weighted* setting. They introduced the *weighted envy-freeness up to one item* (WEF1) for the allocation of goods and show that WEF1 allocations always exist. Lately, WPROP1 allocations have been proved to exist for chores [12], and the mixture of goods and chores [6]. Li et al. [26] showed the existence and computation of WPROPX allocations of chores. Wu et al. [36] and Springer et al. [32] proposed algorithms for the computation of WEF1 allocations for chores. We refer to the recent survey [34] for a review of weighted fair allocation.

¹As a comparison, the item-sharing graph defined by the allocation returned by the Moving Knife Algorithm is a path; while that for the Eating Algorithm can be any complex dense multi-graph.

Fair Allocation with Money. Beyond additive valuation functions, Brustle et al. [13] showed that a subsidy of $2(n - 1)$ dollars per agent suffices to guarantee envy-freeness for monotonic functions. The result was recently improved by Kawase et al. [25] to $n - 1$ per agent. Barman et al. [8] considered the dichotomous valuations and showed that envy-freeness can be guaranteed with a per-agent subsidy of at most 1. Regarding truthfulness, Goko et al. [23] showed that for submodular functions, there exists a truthful mechanism that guarantees envy-freeness with a subsidy of at most one dollar per agent. The subsidy setting was recently extended to the weighted setting [17, 19]. A similar setting is called fair allocation with monetary transfers, introduced by Aziz [2], that allows agents to transfer money to each other. Instead of minimizing the total subsidy, their result focused on the characterization of allocations that are equitable and envy-free with monetary transfers. When considering money as a divisible good, the setting of fair allocation with money is similar to the fair allocation of mixed divisible and indivisible items, which also receives much attention [10, 11, 27]. For a more detailed review of the existing works on mixed fair allocation, please refer to the recent survey [29].

REFERENCES

- [1] Amanatidis, G., Aziz, H., Birmpas, G., Filos-Ratsikas, A., Li, B., Moulin, H., Voudouris, A.A., Wu, X.: Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.* **322**, 103965 (2023)
- [2] Aziz, H.: Achieving envy-freeness and equitability with monetary transfers. In: *AAAI*. pp. 5102–5109. AAAI Press (2021)
- [3] Aziz, H., Chan, H., Li, B.: Weighted maxmin fair share allocation of indivisible chores. In: *IJCAI*. pp. 46–52. *ijcai.org* (2019)
- [4] Aziz, H., Huang, X., Kimura, K., Saha, I., Sun, Z., Suzuki, M., Yokoo, M.: Weighted envy-free allocation with subsidy. *CoRR* **abs/2408.08711** (2024)
- [5] Aziz, H., Li, B., Moulin, H., Wu, X.: Algorithmic fair allocation of indivisible items: a survey and new questions. *SIGecom Exch.* **20**(1), 24–40 (2022)
- [6] Aziz, H., Moulin, H., Sandomirskiy, F.: A polynomial-time algorithm for computing a pareto optimal and almost proportional allocation. *Oper. Res. Lett.* **48**(5), 573–578 (2020)
- [7] Babaioff, M., Ezra, T., Feige, U.: Fair-share allocations for agents with arbitrary entitlements. In: *EC*. p. 127. *ACM* (2021)
- [8] Barman, S., Krishna, A., Narahari, Y., Sadhukhan, S.: Achieving envy-freeness with limited subsidies under dichotomous valuations. In: *IJCAI*. pp. 60–66. *ijcai.org* (2022)
- [9] Barman, S., Krishnamurthy, S.K.: On the proximity of markets with integral equilibria. In: *AAAI*. pp. 1748–1755 (2019)
- [10] Bei, X., Li, Z., Liu, J., Liu, S., Lu, X.: Fair division of mixed divisible and indivisible goods. *Artif. Intell.* **293**, 103436 (2021)
- [11] Bhaskar, U., Sricharan, A.R., Vaish, R.: On approximate envy-freeness for indivisible chores and mixed resources. In: *APPROX-RANDOM. LIPIcs*, vol. 207, pp. 1:1–1:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021)
- [12] Brânzei, S., Sandomirskiy, F.: Algorithms for competitive division of chores. *Mathematics of Operations Research* (2023)
- [13] Brustle, J., Dippel, J., Narayan, V.V., Suzuki, M., Vetta, A.: One dollar each eliminates envy. In: *EC*. pp. 23–39. *ACM* (2020)
- [14] Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In: *BQGT*. p. 74:1. *ACM* (2010)
- [15] Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.* **7**(3), 12:1–12:32 (2019)

- [16] Chakraborty, M., Igarashi, A., Suksompong, W., Zick, Y.: Weighted envy-freeness in indivisible item allocation. *ACM Trans. Economics and Comput.* **9**(3), 18:1–18:39 (2021)
- [17] Choo, D., Ling, Y.H., Suksompong, W., Teh, N., Zhang, J.: Envy-free house allocation with minimum subsidy. *Oper. Res. Lett.* **54**, 107103 (2024)
- [18] Conitzer, V., Freeman, R., Shah, N.: Fair public decision making. In: *EC* (2017)
- [19] Dai, S., Chen, Y., Wu, X., Xu, Y., Zhang, Y.: Weighted envy-freeness in house allocation. *CoRR abs/2408.12523* (2024)
- [20] Farhadi, A., Ghodsi, M., Hajiaghayi, M.T., Lahaie, S., Pennock, D.M., Seddighin, M., Seddighin, S., Yami, H.: Fair allocation of indivisible goods to asymmetric agents. *J. Artif. Intell. Res.* **64**, 1–20 (2019)
- [21] Feige, U., Huang, X.: On picking sequences for chores. In: *EC*. pp. 626–655. ACM (2023)
- [22] Foley, D.: Resource allocation and the public sector. *Yale Economic Essays* pp. 45–98 (1967)
- [23] Goko, H., Igarashi, A., Kawase, Y., Makino, K., Sumita, H., Tamura, A., Yokoi, Y., Yokoo, M.: A fair and truthful mechanism with limited subsidy. *Games and Economic Behavior* **144**, 49–70 (2024)
- [24] Halpern, D., Shah, N.: Fair division with subsidy. In: *SAGT. Lecture Notes in Computer Science*, vol. 11801, pp. 374–389. Springer (2019)
- [25] Kawase, Y., Makino, K., Sumita, H., Tamura, A., Yokoo, M.: Towards optimal subsidy bounds for envy-freeable allocations. In: *AAAI*. pp. 9824–9831. AAAI Press (2024)
- [26] Li, B., Li, Y., Wu, X.: Almost (weighted) proportional allocations for indivisible chores. In: *WWW*. pp. 122–131. ACM (2022)
- [27] Li, Z., Liu, S., Lu, X., Tao, B.: Truthful fair mechanisms for allocating mixed divisible and indivisible goods. In: *IJCAI*. pp. 2808–2816. ijcai.org (2023)
- [28] Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: *EC*. pp. 125–131. ACM (2004)
- [29] Liu, S., Lu, X., Suzuki, M., Walsh, T.: Mixed fair division: A survey. In: *AAAI*. pp. 22641–22649. AAAI Press (2024)
- [30] Maskin, E.S.: On the fair allocation of indivisible goods. In: *Arrow and the Foundations of the Theory of Economic Policy*, pp. 341–349. Springer (1987)
- [31] Robertson, J.M., Webb, W.A.: *Cake-cutting algorithms - be fair if you can*. A K Peters (1998)
- [32] Springer, M., Hajiaghayi, M., Yami, H.: Almost envy-free allocations of indivisible goods or chores with entitlements. In: *AAAI*. pp. 9901–9908. AAAI Press (2024)
- [33] Steihaus, H.: The problem of fair division. *Econometrica* **16**, 101–104 (1948)
- [34] Suksompong, W.: Weighted fair division of indivisible items: A review. *Information Processing Letters* p. 106519 (2024)
- [35] Wu, X., Zhang, C., Zhou, S.: One quarter each (on average) ensures proportionality. In: *International Conference on Web and Internet Economics*. pp. 582–599. Springer (2023)
- [36] Wu, X., Zhang, C., Zhou, S.: Weighted EF1 allocations for indivisible chores. In: *EC*. p. 1155. ACM (2023)

STATE KEY LAB OF IOTSC, UNIVERSITY OF MACAU, MACAU, CHINA

Email address: xiaoweiwu@um.edu.mo

POPULAR MATCHING UNDER MATROID AND OPTIMALITY CONSTRAINTS

YU YOKOI

Classification AMS 2020: 91B68, 90C27.

Keywords: matching under preferences, popularity, matroids

We present recent algorithmic results on popular matching problems with constraints, such as matroid, size, and weight constraints. These results have appeared in [3, 9].

In general, popular matchings are defined in bipartite graphs and have two models. One is the *one-sided preferences model*, where only vertices on one side have preferences, and the other is the *two-sided preferences model*, where all vertices have preferences. In the former model, a popular matching may not exist, but it is tractable to test its existence. In the latter model, a popular matching always exists, as any stable matching is popular. For each of these models, a matroid generalization has been proposed [4, 5].

As a variant of a popular matching, the notion of a *popular maximum matching* has been proposed. This is a popular solution among the set of all maximum matchings, and many algorithmic results have been developed [6, 8]. A more general *popular maximum-weight matching* has also been introduced and shown to be computable efficiently [7]. Our results are common generalizations of those in the above settings.

We first describe our result on the one-sided preferences model. For a positive integer k , we write $[k] = \{1, 2, \dots, k\}$. Let $\{S_1, S_2, \dots, S_n\}$ be a partition of a finite set S , and let M_1 be a 1-partition matroid defined by this partition. That is, $M_1 = (S, \mathcal{I}_1)$ is a matroid with ground set S and independent set family $\mathcal{I}_1 = \{I \subseteq S : |I \cap S_i| \leq 1 \ (i \in [n])\}$. Each index $i \in [n]$ represents an agent and has a partial order \succ_i on $S_i \cup \{\emptyset\}$ satisfying $u \succ_i \emptyset$ for each element $u \in S_i$. Additionally, we have another matroid $M_2 = (S, \mathcal{I}_2)$, which can be arbitrary and has no associated orders. For a common independent set $I \in \mathcal{I}_1 \cap \mathcal{I}_2$ and an agent $i \in [n]$, let $I(i)$ denote the unique element in $I \cap S_i$ if it exists, and \emptyset otherwise. Given common independent sets $I, J \in \mathcal{I}_1 \cap \mathcal{I}_2$, define $\Delta(I, J) \in \mathbb{Z}$ as

$$\Delta(I, J) = |\{i \in [n] : I(i) \succ_i J(i)\}| - |\{i \in [n] : J(i) \succ_i I(i)\}|.$$

We are also given a weight function $w: S \rightarrow \mathbb{R}$. For a common independent set $I \subseteq S$, its weight $w(I)$ is defined as $w(I) = \sum_{u \in I} w(u)$. Let $\text{opt}(w)$ denote the maximum weight of a common independent set, i.e., $\text{opt}(w) = \max\{w(I) : I \in \mathcal{I}_1 \cap \mathcal{I}_2\}$.

Definition 0.1. A common independent set $I \in \mathcal{I}_1 \cap \mathcal{I}_2$ is called a popular maximum-weight common independent set if $w(I) = \text{opt}(w)$ and $\Delta(I, J) \geq 0$ holds for every $J \in \mathcal{I}_1 \cap \mathcal{I}_2$ with $w(J) = \text{opt}(w)$.

We provide a polynomial-time algorithm to determine the existence of such a solution.

Theorem 0.2 (Tractability in the one-sided preferences model). *Given a 1-partition matroid $M_1 = (S, \mathcal{I}_1)$ associated with partial orders $\{\succ_i\}_{i \in [n]}$, any matroid $M_2 = (S, \mathcal{I}_2)$, and a weight function $w: S \rightarrow \mathbb{R}$, one can determine the existence of a popular maximum-weight common independent set and find one, if it exists, in polynomial time.*

This result is obtained by a reduction to the *popular common base* problem, for which a polynomial-time algorithm was recently provided in [9]. The popular common base problem includes the *popular arborescence* problem as a notable special case.

For the two-sided preferences model, we address the problem of finding a popular maximum-weight matching in a many-to-many setting with two-sided preferences and matroid constraints. In this model, two matroids are given on the same ground set S , both as direct sums: $M_1 = M_1^1 \oplus M_2^1 \oplus \cdots \oplus M_{k_1}^1$, $M_2 = M_1^2 \oplus M_2^2 \oplus \cdots \oplus M_{k_2}^2$. Each summand $M_j^i = (S_j^i, \mathcal{I}_j^i)$ corresponds to an agent, so there are $k_1 + k_2$ agents. A set $I \subseteq S$ is feasible if $I \cap S_j^i \in \mathcal{I}_j^i$ for each $i \in \{1, 2\}$ and $j \in [k_i]$. The simple bipartite matching model is a special case where each M_j^i is a uniform matroid of rank 1. Since the definition of popularity in general matroid intersection is nontrivial, we omit it and refer the reader to [3]. We provide the following tractability result in this setting.

Theorem 0.3 (Tractability in the two-sided preferences model). *In the two-sided preferences model, if preferences are total orders, then a popular maximum-weight common independent set always exists and can be found in polynomial time.*

This theorem assumes that preferences are total orders. Note that if ties are allowed, finding a popular matching is NP-hard, even in the simple bipartite matching model [2].

To contrast these tractability results, we also provide some hardness results on *popular near-maximum-weight matching*, which is a matching that is popular among all matchings whose weights are at least a given threshold. See [3] for details.

REFERENCES

- [1] David J Abraham, Robert W Irving, Telikepalli Kavitha, and Kurt Mehlhorn. Popular matchings. *SIAM Journal on Computing*, 37(4):1030–1045, 2007.
- [2] Péter Biró, Robert W Irving, and David F Manlove. Popular matchings in the marriage and roommates problems. In *Proc. of the Seventh International Conference on Algorithms and Complexity (CIAC)*, pages 97–108,, 2010.
- [3] Gergely Csáji, Tamás Király, Kenjiro Takazawa and Yu Yokoi. Popular maximum-utility matchings with matroid constraints. *arXiv preprint arXiv:2407.09798*, 2024.
- [4] Naoyuki Kamiyama. Popular matchings with ties and matroid constraints. *SIAM Journal on Discrete Mathematics*, 31(3):1801–1819, 2017.
- [5] Naoyuki Kamiyama. Popular matchings with two-sided preference lists and matroid constraints. *Theoretical Computer Science*, 809:265–276, 2020.
- [6] Telikepalli Kavitha. Maximum matchings and popularity. *SIAM Journal on Discrete Mathematics*, 38(2):1202–1221, 2024.
- [7] Telikepalli Kavitha. Popular solutions for optimal matchings. In *Proc. of 50th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 297–311, 2024.
- [8] Telikepalli Kavitha, Tamás Király, Jannik Matuschke, Ildikó Schlotter, and Ulrike Schmidt-Kraepelin. The popular assignment problem: when cardinality is more important than popularity. In *Proc. of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 103–123. SIAM, 2022.
- [9] Telikepalli Kavitha, Kazuhisa Makino, Ildikó Schlotter, and Yu Yokoi. Arborescences, colorful forests, and popularity. In *Proc. of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3724–3746. SIAM, 2024.

INSTITUTE OF SCIENCE TOKYO, TOKYO 152-8552, JAPAN
 Email address: yokoi@comp.isct.ac.jp

ATTAINING EF1 ALLOCATIONS BY EXCHANGING GOODS

SHEUNG MAN YUEN

Classification AMS 2020: 91B32

Keywords: fair division, EF1, exchanging goods, reconfiguration, optimal control

1. INTRODUCTION

Fair division refers to the study of how to allocate resources fairly among competing agents. The fair division literature typically assumes that there is a set of unallocated goods and the objective is to allocate them fairly. We take a different perspective by assuming that an initial allocation is already given. In the first problem of this report, we initiate the study of reachability in fair division: given two fair allocations—an initial allocation and a target allocation—we are interested in whether the target allocation can be reached from the initial allocation via a sequence of operations such that every intermediate allocation is also fair. In the second problem of this report, we are instead given an initial unfair allocation, and we are interested in whether a fair allocation can be reached from the initial allocation via a sequence of operations. In both problems, the fairness benchmark used is *envy-freeness up to one good (EF1)*, and we allow any two agents to *exchange* a pair of goods in each operation.

This is a joint work with Ayumi Igarashi, Naoyuki Kamiyama, and Warut Suksompong [1, 2].

1.1. Related Work. Closest to our work is perhaps a line of work initiated by Gourvès *et al.* [3]. These authors considered the “housing market” setting, where the number of agents is the same as the number of goods and each agent receives exactly one good. In their model, a pair of agents is allowed to exchange goods if the two agents are neighbors in a given social network and the exchange benefits both agents. Their paper, along with a series of follow-up papers [4, 5, 6, 7], explored the complexity of determining whether an allocation can be reached from another allocation in this model and its variants. More broadly, reachability problems are also known as *reconfiguration* problems [8]; examples of such problems that have been studied include minimum spanning tree [9], graph coloring [10], and perfect matching [11].

Boehmer *et al.* [12] studied the problem of *discarding* goods from an initial allocation in order to reach an envy-free or EF1 allocation. In a similar vein, Dorn *et al.* [13] investigated deleting goods to attain another fairness notion called *proportionality*. Aziz *et al.* [14] focused on reallocating goods to make agents better off, but did not delve into the aspect of fairness. Chandramouleeswaran *et al.* [15] examined *transferring* goods starting from a “near-EF1” allocation with the goal of reaching an EF1 allocation. Segal-Halevi [16] considered the reallocation of a *divisible* good and explored the trade-off between guaranteeing a minimum utility for every agent and ensuring each agent a

Date: 4 February 2025.

certain fraction of her original utility. Chevaleyre *et al.* [17] also strived to reach fair allocations but via exchanges with money.

1.2. Preliminaries. Let N be a set of $n \geq 2$ agents, and M be a set of $m \geq 1$ goods. A *bundle* is a subset of goods. An *allocation* $\mathcal{A} = (A_1, \dots, A_n)$ is an ordered partition of M into n bundles such that bundle A_i is allocated to agent $i \in N$. An *(allocation) size vector* $\vec{s} = (s_1, \dots, s_n)$ is a vector of non-negative integers such that $\sum_{i \in N} s_i = m$. We say that an allocation \mathcal{A} has size vector \vec{s} if $|A_i| = s_i$ for all $i \in N$. A size vector \vec{s} is *balanced* if $|s_i - s_j| \leq 1$ for all $i, j \in N$, and an allocation is *balanced* if it has a balanced size vector. Each agent $i \in N$ has an additive *utility function* u_i that maps bundles to non-negative real numbers. The utility functions are *identical* if $u_i = u_j$ for all $i, j \in N$, and the utility functions are *binary* if $u_i(\{g\}) \in \{0, 1\}$ for all $i \in N$ and $g \in M$. An allocation \mathcal{A} is *EF1* if for all $i, j \in N$ and $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that $u_i(A_i) \geq u_i(A_j \setminus \{g\})$. A *(fair division) instance* \mathcal{I} consists of a set of agents N , a set of goods M , the agents' utility functions $(u_i)_{i \in N}$, and a size vector \vec{s} .

Given an instance, define the *exchange graph* G as a simple undirected graph with the following properties: the set of vertices consists of all allocations \mathcal{A} with size vector \vec{s} , and the set of edges consists of all pairs $\{\mathcal{A}, \mathcal{B}\}$ such that $\mathcal{B} = (B_1, \dots, B_n)$ can be obtained from $\mathcal{A} = (A_1, \dots, A_n)$ by having two agents exchange one pair of goods with each other—that is, there exist distinct agents $i, j \in N$ and goods $g \in A_i$ and $g' \in A_j$ such that $B_i = (A_i \cup \{g'\}) \setminus \{g\}$, $B_j = (A_j \cup \{g\}) \setminus \{g'\}$, and $B_k = A_k$ for all $k \in N \setminus \{i, j\}$. Note that the exchange graph is a non-empty connected graph. A path from one allocation to another on the graph is called an *exchange path*. The *distance* (or the *optimal number of exchanges*) between two allocations is the length of a shortest exchange path between them. Define the *EF1 exchange graph* H as the subgraph of the exchange graph G induced by all EF1 allocations. An exchange path using only the edges in H is called an *EF1 exchange path*. An EF1 exchange path is *optimal* if its length is equal to the distance between the two corresponding allocations (in G).

2. MAIN RESULTS

We first state the results related to the first problem of our work. The first five results are on the connectivity of the EF1 exchange graph.

Theorem 2.1. *There exists an instance with two agents such that the EF1 exchange graph is not connected.*

Theorem 2.2. *For two agents with identical (resp. binary) utilities, the EF1 exchange graph is always connected. Moreover, there always exists an optimal EF1 exchange path between any two EF1 allocations, and this path can be found in polynomial time.*

Theorem 2.3. *For each $n \geq 3$, there exists an instance with n agents with identical (resp. binary) utilities such that the EF1 exchange graph is not connected.*

Theorem 2.4. *For three or more agents with identical binary utilities, the EF1 exchange graph is always connected. Moreover, an EF1 exchange path between any two EF1 allocations can be found in polynomial time.*

Theorem 2.5. *Deciding the existence of an EF1 exchange path between two EF1 allocations is PSPACE-complete.*

Without EF1 considerations, finding the distance between two allocations is NP-hard.

Theorem 2.6. *Finding the distance between two allocations is NP-hard.*

The next three results concern the existence of an optimal EF1 exchange path. Note that Theorem 2.2 already gives the result for two agents with identical (resp. binary) utilities.

Theorem 2.7. *There exists an instance with two agents satisfying the following properties: the EF1 exchange graph is connected, but for some pair of EF1 allocations, no optimal EF1 exchange path exists between them.*

Theorem 2.8. *For each $n \geq 3$, there exists an instance with n agents with identical binary utilities satisfying the following properties: the EF1 exchange graph is connected, but for some pair of EF1 allocations, no optimal EF1 exchange path exists between them.*

Theorem 2.9. *Deciding the existence of an optimal EF1 exchange path between two EF1 allocations is NP-hard, even for four agents with identical utilities.*

We now state the results related to the second problem of our work. It can be shown that two allocations can be reached from each other via sequential exchanges if and only if they have the same size vector. Deciding whether there exists an EF1 allocation that can be reached from a given initial allocation is equivalent to deciding whether an EF1 allocation with a given size vector exists in a given instance—we shall refer to this latter problem as REFORMABILITY.

Theorem 2.10. REFORMABILITY is in P for (i) two agents with identical utilities, (ii) a constant number of agents with binary utilities, and (iii) identical binary utilities.

Theorem 2.11. REFORMABILITY is weakly NP-complete for (i) n agents where $n \geq 2$ is a constant, and (ii) n agents with identical utilities where $n \geq 3$ is a constant.

Theorem 2.12. REFORMABILITY is strongly NP-complete for (i) identical utilities, and (ii) binary utilities.

We refer to as OPTIMAL EXCHANGES the problem of computing the optimal number of exchanges to reach an EF1 allocation from a given initial allocation.

Theorem 2.13. OPTIMAL EXCHANGES is in P for (i) two agents with identical utilities, (ii) a constant number of agents with binary utilities, and (iii) identical binary utilities.

Theorem 2.14. OPTIMAL EXCHANGES is NP-hard for (i) two agents, (ii) three or more agents with identical utilities, and (iii) binary utilities, even when the given initial allocation is balanced.

Given n and s , let $f(n, s)$ be the smallest integer such that for every instance with n agents, ns goods, and size vector $\vec{s} = (s, \dots, s)$, and for every allocation \mathcal{A} with size vector \vec{s} in the instance, there exists an EF1 allocation with distance at most $f(n, s)$ from \mathcal{A} .

Theorem 2.15. $f(n, s) \approx s(n - 1)/2$.

REFERENCES

- [1] Ayumi Igarashi, Naoyuki Kamiyama, Warut Suksompong, and Sheung Man Yuen. Reachability of fair allocations via sequential exchanges. *Algorithmica*, 86(12):3653–3683, 2024.
- [2] Sheung Man Yuen, Ayumi Igarashi, Naoyuki Kamiyama, and Warut Suksompong. Reforming an unfair allocation by exchanging goods. *Computing Research Repository (CoRR)*, abs/2412.19264, 2024.
- [3] Laurent Gourvès, Julien Lesca, and Anaëlle Wilczynski. Object allocation via swaps along a social network. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 213–219, 2017.
- [4] Sen Huang and Mingyu Xiao. Object reachability via swaps under strict and weak preferences. *Autonomous Agents and Multi-Agent Systems*, 34(2):51:1–51:33, 2020.
- [5] Fu Li, C. Gregory Plaxton, and Vaibhav B. Sinha. Object allocation over a network of objects: Mobile agents with strict preferences. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1578–1580, 2021.
- [6] Luis Müller and Matthias Bentert. On reachable assignments in cycles. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, pages 273–288, 2021.
- [7] Takehiro Ito, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, Yuta Nozaki, Yoshio Okamoto, and Kenta Ozeki. On reachable assignments under dichotomous preferences. *Theoretical Computer Science*, 979:114196, 2023.
- [8] Naomi Nishimura. Introduction to reconfiguration. *Algorithms*, 1(4):52:1–52:25, 2018.
- [9] Takehiro Ito, Erik D. Demaine, Nicholas J. A. Harvey, Christos H. Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12–14):1054–1065, 2011.
- [10] Matthew Johnson, Dieter Kratsch, Stefan Kratsch, Viresh Patel, and Daniël Paulusma Finding shortest paths between graph colourings. *Algorithmica*, 75(2):295–321, 2016.
- [11] Marthe Bonamy, Nicolas Bousquet, Marc Heinrich, Takehiro Ito, Yusuke Kobayashi, Arnaud Mary, Moritz Mühlenthaler, and Kunihiko Wasa. The perfect matching reconfiguration problem. In *Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 80:1–80:14, 2019.
- [12] Niclas Boehmer, Robert Brederbeck, Klaus Heeger, Dušan Knop, and Junjie Luo. Multivariate algorithmics for eliminating envy by donating goods. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 127–135, 2022.
- [13] Britta Dorn, Ronald de Haan, and Ildikó Schlotter. Obtaining a proportional allocation by deleting items. *Algorithmica*, 83(5):1559–1603, 2021.
- [14] Haris Aziz, Péter Biró, Jérôme Lang, Julien Lesca, and Jérôme Monnot. Efficient reallocation under additive and responsive preferences. *Theoretical Computer Science*, 790:1–15, 2019.
- [15] Harish Chandramouleeswaran, Prajakta Nimbhorkar, and Nidhi Rathi. Fair division in a variable setting. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Forthcoming.
- [16] Erel Segal-Halevi. Redividing the cake. *Autonomous Agents and Multi-Agent Systems*, 36(1):14:1–14:36, 2022.
- [17] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Reaching envy-free states in distributed negotiation settings. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1239–1244, 2007.

SCHOOL OF COMPUTING, NATIONAL UNIVERSITY OF SINGAPORE, COM1, 13 COMPUTING DR, SINGAPORE 117417

Email address: ysm@u.nus.edu