Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)
Early Career Track

Algorithmic Fairness in Distribution of Resources and Tasks

Hadi Hosseini
Pennsylvania State University
hadi@psu.edu

Abstract

The widespread adoption of Artificial Intelligence (AI) systems has profoundly reshaped decision-making in social, political, and commercial contexts. This paper explores the critical issue of fairness in AI-driven decision-making, particularly in allocating resources and tasks. By examining recent advancements and key questions in computational social choice, I highlight challenges and prospects in designing fair systems in collective decision-making that are scalable, adaptable to intricate environments, and are aligned with complex and diverse human preferences.

1 Introduction

The rapid growth in the adoption of Artificial Intelligence (AI) systems has transformed the ways social, political, and commercial decisions are governed. These systems are primarily influenced by algorithmic processes, ranging from generative AI that relies on statistical patterns to symbolic and rule-based AI paradigms that require explicit representation of problems. With the rise of algorithmic methods, fairness has emerged as a pivotal concern in decision-making, particularly in application domains that involve the participation of multiple intelligent agents—autonomous agents, institutions, or humans—that often have conflicting preferences over the possible outcomes.

The field of computational social choice combines theoretical principles and axioms of fairness rooted in economics with the necessary algorithmic paradigms in computer science to offer the crucial building blocks of collective decision-making for a "*society of agents*". The recent efforts in this field have led to fascinating progress in fair preference aggregation methods to reach consensus or to rank outcomes [Brandt *et al.*, 2016], to develop democratic processes for deciding on public budget spending [Aziz and Shah, 2021], and to distribute divisible (e.g. computation) or indivisible (e.g. public housing) resources and tasks [Amanatidis *et al.*, 2023], to name a few. Fairness, in particular, plays an instrumental role in the allocation of public and private resources both in centralized and distributed settings, ranging from the distribution of scarce medical resources (e.g. vaccines) in federated healthcare and the distribution of tasks in digital gig economy (e.g. ridesharing platforms) to the collective development of foundation models in data-oriented systems.

These problems often have to deal with complex environments with incomplete or uncertain information: preference information may be uncertain or noisy, knowledge about the resources may be incomplete, unavailable, or simply expensive to elicit, participants' strategic behavior may critically influence the desirability of outcomes, and fairness metrics may not be aligned with social axioms. Moreover, complications arise when distributing tasks generating negative values for some agents (e.g. household chores or training a machine learning model), as opposed to positively-valued resources. Despite the recent advancements, achieving fair solutions in complex environments remains a challenging (and exciting) task that begs the following fundamental questions: (i) which fairness axioms are more suitable for different applications, and whether these axioms are compatible with other societal properties (e.g. social welfare)? (ii) how should we design scalable algorithms (or approximation algorithms) to compute such fair solutions? (iii) how do we align agents' incentives to behave truthfully? And which fairness axioms are aligned with human and societal values?

In this paper, I will demonstrate the nuances in investigating algorithmic and axiomatic boundaries of fairness in distributing resources and tasks, and discuss the impact of strategic behavior of agents on fairness of solutions. The goal is to provide an overview of recent results, primarily focused on my research, in addressing the above grand challenges in designing fair algorithms.

The paper is structured as follows: In Section 2, I will survey recent works in achieving fair allocations under a variety of constraints or models, In Section 3, I will discuss the strategic aspects of fair allocation and matching markets. In Section 4, I will argue that fairness axioms should be aligned with human value judgements. Finally, I will conclude, in Section 5, by laying out an overarching plan for research in fair collective decision-making.

2 Allocating Resources and Tasks

Achieving fairness when distributing *indivisible* resources or tasks is a notoriously difficult problem due to several axiomatic impossibility and computational intractability results. The primary fairness concepts can be seen as either *comparison-based* notions that rely on pairwise comparisons between the agents (e.g. envy-freeness [Foley, 1967]), or *threshold-based* criteria that deal with achieving a fair-share of the set of items (e.g. maximin share [Budish, 2011]). When dealing with indivisible items, none of the above fairness notions can be guaranteed, giving rise to a variety of deterministic or randomized relaxations [Caragiannis *et al.*, 2019; Lipton *et al.*, 2004; Bogomolnaia and Moulin, 2004]). These existing approaches to approximate fairness often impose a high cost on information transparency or social welfare, lead to strategic behavior by participating agents, or are not robust to uncertain or dynamic environments. In this section, I will survey some of the recent research in addressing these challenges.

2.1 Allocating Goods and Resources

Epistemic Fairness. The most well-studied relaxation of envy-freeness such as *envy-freeness up to one item* (EF1) requires that any pairwise envy between the agents can be eliminated by the *counterfactual* removal of items (e.g. a single good from the envied agent) [Lipton *et al.*, 2004; Budish, 2011]. These 'up to one' relaxations have attracted significant attention from the AI/CS community in the past few years. These fairness concepts operate on two strong assumptions: (i) agents are fully aware of other agents' allocations, and (ii) the unfairness can be justified by reminding agents of "hypothetical" worse solutions. These assumptions—although theoretically appealing—do not reflect the inherent nature of fairness in practical domains.

To address these issues, we proposed a rather radical *epistemic* approach that utilizes *information asymmetry* which, in contrast to EF1, does not require counterfactual reasoning. This epistemic framework enables us to exploit information asymmetry when allocating indivisible goods (e.g. by carefully hiding minimum information) [Hosseini *et al.*, 2020]. We show that while minimizing this asymmetry to achieve envy-freeness remains intractable, an envy-free allocation can be computed in polynomial time by withholding information about only a small number of items (i.e. $n - 1$ items). Moreover, such allocations can be computed along with Pareto optimality (a welfare concept that guarantees no improving exchanges among the agents) in pseudo-polynomial time. In Section 4, I will discuss how epistemic fairness concepts are compared to their counterfactual counterparts based on perceived human values.

Guaranteeing a Fair Share. A compelling fairness axiom is maximin share (MMS), which requires that each agent receives a fair share that it is at least as good as if they were to create the best n-partition and receive the least preferred bundle [Budish, 2011]. Since such allocations do not always exist [Kurokawa *et al.*, 2018], several multiplicative approximations have been proposed that give each agent β fraction of their MMS value [Ghodsi *et al.*, 2018; Garg and Taki, 2020]. These approximation techniques are not *robust* to slight noise in agents' cardinal valuations, i.e., a small perturbation in the valuation of a rather *irrelevant* item induces a radical change in the final solution. In particular, multiplicative approximations of MMS are extremely sensitive to infinitesimal change in agents' valuations.

We proposed a technique for *ordinal approximation* of MMS that is robust to slight modifications or noise in valuations. Our work was the first approximation framework on two dimensions, allowing for α fraction of agents to receive β approximation of their fair share value [Hosseini and Searns, 2021; Searns and Hosseini, 2020]. In this vein, we show the existence of allocations that guarantee $2/3$ of agents receive their exact MMS value. In addition, we establish a connection between this approximation framework and ordinal approximation of MMS fair share that rely on adding (a small number of) dummy agents to the mix. In fact, we devise algorithms that ensure that each agent receive their MMS value if they were to distribute the goods among $\lfloor \frac{3n}{2} \rfloor$ agents [Hosseini *et al.*, 2022b].

Fairness in Dynamic Environments. Real-world applications often deal with complex environments that are inherently dynamic, that is, decisions must be made as new resources emerge *over time* and the information about future is not available apriori. Thus, it is often impossible to take history or prior decisions into account, and decisions about how (and what) to allocate need to be made when such information is unavailable (or uncertain). These problems are often modeled as online bipartite matching wherein upon arrival of items, preferences of agents are revealed as edges incident to the items, and the algorithm must irrevocably match the item to such agents. Here, agents may be partitioned into groups, and thus, a fair treatment demands that all groups receive a fair allocation.

In a recent work, we study deterministic and randomized algorithms for allocating indivisible items (leading to integral allocations) and for allocating divisible items (leading to fractional allocations) among such groups. In particular, we design online algorithms that achieve $1/2$ approximation of EF1 and MMS, and show that these guarantees are tight. For dealing with divisible items, we design an algorithm that achieves $(1 - \frac{1}{e})$-approximation of envy-freeness, and establish a $3/4$ upper bound [Hosseini *et al.*, 2023a]. Furthermore, we discuss *randomized algorithms* to improve the above approximation ratios, and introduce several intriguing open problems.

Of course, allocation problems may involve sequential allocations with dynamic preferences [Hosseini *et al.*, 2015] or dynamic population of agents [Kash *et al.*, 2014; Walsh, 2011]. These settings often require studying randomized algorithms with probabilistic guarantees on fairness (e.g. stochastic dominance envy-freeness) [Hosseini and Larson, 2019; Hosseini *et al.*, 2018].

Degree of Fairness. In some constrained allocation problems, agents are limited to receive only a single resource (e.g. assigning public housing to families). These problems are often known as "Shapley-Scarf Housing Markets" where the goal is often to achieve economic efficiency by resolving mutually-beneficial exchanges among agents with initial endowments [Shapley and Scarf, 1974]. Focusing on fairness, the 'one house per agent' constraint immediately renders previous approximate fairness axioms meaningless: any 'up to one' style axiom or fair-share variants are satisfied vacuously. This motivates the study of fairness metrics that measure the degree of envy among a society of agents (e.g. the aggregate envy of all agents or the number of envious agents). Our recent work has focused on studying the degree of fairness under various measures of economic efficiency (e.g. utilitarian or egalitarian welfare) [Hosseini *et al.*, 2024a]. When the set of agents are placed along the vertices of a social network, we showed theoretical upper bounds on minimizing aggregate envy, and further exploit the graph structures to illuminate the axiomatic and computational boundaries of fairness [Hosseini *et al.*, 2023e; Hosseini *et al.*, 2024c].

Preference Elicitation. It is often challenging, if not impossible, to elicit complete preference information due to uncertainty in the environment, high elicitation cost, or simply cognitive burden on agents to assign *exact* rankings to each and every one of the items. Instead of asking agents to report their preferences over item (or bundles) at once, it is possible to use a *preference elicitation* mechanism, in which agents repeatedly respond to queries about their preferences. When agents are constrained to receive only one item, preferences can be elicited through pairwise comparisons or through reporting the k most-preferred alternatives [Aziz *et al.*, 2015; Borodin *et al.*, 2022]. In this vein, our work investigated the *query complexity* of elicitation mechanisms, and designed algorithms that can achieve *competitive* social welfare (measured by Pareto optimality or rank maximality). In particular, our online algorithm requires $O(\sqrt{n})$ queries, with a matching lower bound, which implies that our algorithm is asymptotically optimal [Hosseini *et al.*, 2021a].

2.2 Distributing Tasks (and Mixtures)

Unlike allocating desirable resources where the more is often preferred, receiving more tasks is never preferred. These problems range from allocating household chores to allocating computational or sampling tasks to servers to train AI models. Interestingly, societal problems that involve collective ownership of responsibility (e.g. limiting gas emissions or managing nuclear waste) at their core involve distribution of negatively valued tasks.

Fair Allocation of Tasks. Distributing tasks is crucially different from allocating goods or resources. While goods are freely disposable, tasks must be completed fully. These fundamental differences along with technical challenges motivates the study of fairness in this domain. For instance, while for positively valued items (goods) an EF1 and PO allocation always exist [Caragiannis *et al.*, 2019] and can be computed in pseudo-polynomial time [Barman *et al.*, 2019], for tasks even the existence of EF1 allocations together with PO remains an open problem.

Our work aims at providing a more fine-grained notion for distributing chores. Under the aforementioned epistemic framework, we introduce 'dubious' tasks to achieve envy-freeness through minimal information asymmetry. In this setting, we prove the existence of envy-free and PO allocations with only $2n-2$ dubious tasks and strengthen it to $n-1$ in some special classes of valuations.

Similar to the allocation of resources, one may also take a threshold-based approach such as MMS to distribute tasks. The non-existence and computational intractability once again strike [Aziz *et al.*, 2017], motivating the investigation of approximate fair share. In particular, an agent's ordinal fair share, aka 1-out-of-k MMS, is determined by partitioning the indivisible tasks into k bundles in a counterfactual world where there are fewer agents ($k < n$) available. We prove the existence of allocations with $k \leq \lfloor \frac{3n}{4} \rfloor$, and develop efficient algorithms to compute such allocations for $k \leq \lfloor \frac{2n}{3} \rfloor$ [Hosseini *et al.*, 2022a].

Goods, Chores, and Mixtures. In many domains, the preferences over items (or outcomes) may be *subjective*, i.e., some agents may consider a resource as a *good* (with non-negative utility) while others may see the same as a *task* (with negative utility). For instance, in peer reviewing, reviewers may consider a paper to be a chore if it is outside of their immediate expertise while another subset of reviewers consider it as a good due its proximity to their own field. A gold-standard fairness relaxation is *envy-freeness up to any item* (EFX) [Caragiannis *et al.*, 2019]. Yet, its existence and computation has remained a notable open problem even when all items are considered to be goods. Focusing on mixtures of goods and chores, we showed, for the first time, that an allocation satisfying EFX could fail to exist [Hosseini *et al.*, 2023f] and its corresponding decision problem is computationally hard [Hosseini *et al.*, 2023c] even under a restricted domain of lexicographic preferences—a subdomain of additive preferences. These results stand in contrast to goods-only [Hosseini *et al.*, 2021b; Hosseini and Larson, 2019] and chores-only [Hosseini *et al.*, 2023f] settings where efficient algorithms exist that can compute fair (EFX) and efficient (Pareto optimality) allocations under the lexicographic domain of preferences.

Despite the challenges posed in settings with mixture of resources and tasks, we demonstrate that it is possible to identify natural domains where fairness notions (e.g. EFX, MMS, EF1) in conjunction with optimality always exist and can be computed efficiently [Hosseini *et al.*, 2023c]. Interestingly, the results on restricted domains can often, but not always, be lifted up to include more complex preferences. For instance, in a recent work we develop algorithms to compute EFX and MMS allocation of goods under weakly lexicographic preferences; and demonstrate that the problem becomes far more challenging when it comes to distributing tasks when preferences contain ties [Hosseini *et al.*, 2024b].

Distributing Tasks on Graphs. When agents, resources, or tasks are located on vertices of a connected graph, often the graph structure plays an instrumental role in the design of fair algorithmic solutions. Fair distribution of tasks or resources among agents that are either connected by a network or need to cover a path on a network is particularly critical with the rise of digital marketplaces (e.g. package delivery tasks). We investigated balancing fairness (e.g. EF1 and MMS) and efficiency (measured by delivery time or distance travelled) and study the *price of fairness* [Bertsimas *et al.*, 2011; Dickerson *et al.*, 2014] when distributing delivery orders among agents on a network with possibly submodular costs [Hosseini *et al.*, 2023d].

3 Incentives and Strategic Behavior

In many novel application domains (e.g. federated healthcare or data/compute sharing frameworks), decisions heavily rely on eliciting data (e.g. preference information) from the involving parties. Yet, participants may act strategically to exploit the system or influence the outcome to their own benefit, which could drastically hinder the quality of decisions (measured by stability, fairness, and efficiency). In fact, incentivizing truthful behavior (e.g. strategyproofness) stands in contrast with other socially desirable axioms such as fairness and stability in both one-sided and two-sided matching markets [Roth, 1982; Svensson, 1999].

3.1 One-Sided Allocation Mechanisms

The interaction between economic efficiency and fairness with truthful incentives has been a subject of extensive studies within economics and computer science. Unfortunately, with a few exceptions, incentivizing truthful reporting stands in contrast with fairness (and efficiency) [Zhou, 1990; Bogomolnaia and Moulin, 2004]. When agents are restricted to receive a single item, randomized mechanisms can sometimes strike a balance between truthfulness (aka strategy-proofness) and weaker axioms of ex ante fairness [Abdulkadiroğlu and Sönmez, 1998; Hosseini *et al.*, 2018].

It turns out that when allocating multiple resources or tasks, no truthful mechanism can guarantee any 'up to one' style fairness axiom ex post even for two agents [Amanatidis *et al.*, 2017]. These impossibility results raise the question of whether truthfulness and fairness remain incompatible under some natural restrictions. For binary valuations, truthfulness is compatible with fairness and economic efficiency [Halpern *et al.*, 2020], which as we show extends to lexicographic preferences. In fact, we characterize such mechanisms and develop algorithms that satisfy the stronger fairness axiom of envy-freeness up to any item (EFX) [Hosseini *et al.*, 2021b].

3.2 Two-Sided Allocation Mechanisms

One of the most prominent application domains involves allocation markets that are inherently two-sided, i.e., two disjoint sets have preferences over the alternatives on the opposite side with some capacity constraints: drivers are matched with one or few riders according to both preferences; freelancers may prefer certain tasks while job requesters prefer higher-rated freelancers. Here, the stability of an allocation relies on an individual fairness notion known as *justified envy-freeness*, i.e., any envy towards another agent should be justified by the preferences of the matched partner.

In these settings—similar to their one-sided counterparts—no mechanism can simultaneously satisfy individual fairness while incentivizing truthful behavior of agents [Roth, 1982] (see our recent survey on strategic manipulation in these markets [Hosseini and Pathak, 2024]). Thus, the strategic behavior (by single or coalition of agents) plays a key role on the fairness of the resulting outcome. In a series of work, we proposed a novel framework for "two-sided" manipulations where a misreporting agent (aka the *accomplice*) and the manipulator (or the beneficiary) are on different sides of the market. We characterized optimal and sub-optimal manipulation strategies when agents can employ one or a coalition of accomplices to strategically change the outcome of the celebrated Deferred Acceptance algorithm (DA) [Gale and Shapley, 1962]. These characterizations immediately imply polynomial-time algorithms for computing an optimal strategy for coalitions of size two [Hosseini *et al.*, 2021c; Bendline and Hosseini, 2019], and give rise to strategies that benefit all agents of one side [Hosseini *et al.*, 2022c]. A rather *surprising* consequence is that an optimal strategy preserves fairness (i.e. justified envy-freeness), while strategies that involve misreporting agents on both sides may no longer guarantee fairness, and thus, result in unstable allocations.

4 Alignment with Human Values

Due to the simplistic nature of fairness axioms (and their approximate counterparts), they often fail to capture the intricacies of the human perception of fairness in complex real-world problems. It is not yet clear how the algorithmic outcomes driven by these axioms are perceived by individuals and the society. A human-centered approach to fairness calls for establishing a 'hierarchy of fairness axioms' and designing algorithmic solutions that are aligned with human and societal values, among several other natural criteria such as explainability [Hosseini, 2024].

In a recent work, we conducted a series of studies with human subjects, to investigate how individuals perceive two conceptually different relaxations of envy-freeness. In particular, our results demonstrated that an *epistemic relaxation* (as discussed in Section 2) that is based on information asymmetry [Hosseini *et al.*, 2020] is perceived to be fairer compared to other approximations (e.g. EF1) that are rooted in counterfactual reasoning [Hosseini *et al.*, 2023b]. These results demonstrate the influence of human cognition, societal values, and individual biases in human judgments, and further call for axioms based on 'distributional preferences' [Herreiner and Puppe, 2009] that are aligned with human values.

5 Concluding Remarks

Algorithmic fairness is fundamental in a variety of application domains that involve distribution of resources, tasks, or mixtures thereof. It demands the cross-fertilization of ideas between AI, computation, and economics to establish the foundation of socially-responsible collective decision-making. The grand vision is promoting fairness principles in (re)designing new AI and algorithmic systems that are provably fair and are aligned with human and societal values; one that can be suited for any complex domain involving a variety of diverse and often conflicting preferences.

Fulfilling this vision involves overcoming numerous milestones throughout this exciting journey. For example, with algorithmic solutions we need to take into account how individuals interact with algorithms, and whether outcomes (or procedures) are explainable. En route to achieving these goals, we additionally need platforms that facilitate human interactions with allocating algorithms [Ferris and Hosseini, 2020; Bao and Hosseini, 2023] and systems that provide explainable recommendations to users, planners, and the society at large [Suryanarayana *et al.*, 2022].

Acknowledgments

I am grateful to all my mentors for their unwavering support. I am thankful to the students in my group (FAIR Lab) and my amazing collaborators who have been vital part of this research journey. The reported research has been generously supported by NSF Awards IIS-2052488 and IIS-2107173, and an NSF Early Career Award IIS-2144413.

References

[Abdulkadiroğlu and Sönmez, 1998] Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

[Amanatidis *et al.*, 2017] 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.

[Amanatidis *et al.*, 2023] 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*, page 103965, 2023.

[Aziz and Shah, 2021] 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, 2021.

[Aziz *et al.*, 2015] Haris Aziz, Toby Walsh, and Lirong Xia. Possible and necessary allocations via sequential mechanisms. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 468–474, 2015.

[Aziz *et al.*, 2017] Haris Aziz, Gerhard Rauchecker, Guido Schryen, and Toby Walsh. Algorithms for max-min share fair allocation of indivisible chores. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

[Bao and Hosseini, 2023] Yeting Bao and Hadi Hosseini. Mind the gap: The illusion of skill acquisition in computational thinking. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, pages 778–784. ACM, 2023.

[Barman *et al.*, 2019] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. Fair division of indivisible goods among strategic agents. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1811–1813, 2019.

[Bendline and Hosseini, 2019] Theodora Bendline and Hadi Hosseini. Partners in crime: Manipulating the deferred acceptance algorithm through an accomplice. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, pages 9917–9918, 2019.

[Bertsimas *et al.*, 2011] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011.

[Bogomolnaia and Moulin, 2004] Anna Bogomolnaia and Hervé Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

[Borodin *et al.*, 2022] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. Distortion in voting with top-t preferences. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 116–122, 7 2022. Main Track.

[Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.

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

[Caragiannis *et al.*, 2019] 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.

[Dickerson *et al.*, 2014] John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. Price of fairness in kidney exchange. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 1013–1020, 2014.

[Ferris and Hosseini, 2020] James Ferris and Hadi Hosseini. Matchu: An interactive matching platform. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, pages 13606–13607, 2020.

[Foley, 1967] DC Foley. Resource Allocation and the Public Sector. *Yale Economic Essays*, 7(1):73–76, 1967.

[Gale and Shapley, 1962] David Gale and Lloyd S Shapley. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[Garg and Taki, 2020] Jugal Garg and Setareh Taki. An improved approximation algorithm for maximin shares. *Economics and Computation*, pages 379–380, 2020.

[Ghodsi *et al.*, 2018] Mohammad Ghodsi, MohammadTaghi HajiAghayi, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. Fair allocation of indivisible goods: Improvements and generalizations. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 539–556, 2018.

[Halpern *et al.*, 2020] 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.

[Herreiner and Puppe, 2009] Dorothea K Herreiner and Clemens D Puppe. Envy freeness in experimental fair division problems. *Theory and decision*, 67:65–100, 2009.

[Hosseini and Larson, 2019] Hadi Hosseini and Kate Larson. Multiple assignment problems under lexicographic preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 837–845, 2019.

[Hosseini and Pathak, 2024] Hadi Hosseini and Shraddha Pathak. Strategic aspects of stable matching markets: A survey. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, page forthcoming, 2024.

[Hosseini and Searns, 2021] Hadi Hosseini and Andrew Searns. Guaranteeing maximin shares: Some agents left behind. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 238–244. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.

[Hosseini *et al.*, 2015] Hadi Hosseini, Kate Larson, and Robin Cohen. Matching with dynamic ordinal preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, pages 936–943, 2015.

[Hosseini *et al.*, 2018] Hadi Hosseini, Kate Larson, and Robin Cohen. Investigating the characteristics of one-sided matching mechanisms under various preferences and risk attitudes. *Autonomous Agents and Multi-Agent Systems*, 32(4):534–567, 2018.

[Hosseini *et al.*, 2020] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, Jun Wang, and Lirong Xia. Fair division through information withholding. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*, pages 2014–2021, 2020.

[Hosseini *et al.*, 2021a] Hadi Hosseini, Vijay Menon, Nisarg Shah, and Sujoy Sikdar. Necessarily optimal one-sided matchings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5481–5488, 2021.

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

[Hosseini *et al.*, 2021c] Hadi Hosseini, Fatima Umar, and Rohit Vaish. Accomplice manipulation of the deferred acceptance algorithm. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 231–237. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.

[Hosseini *et al.*, 2022a] Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. Ordinal maximin share approximation for chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 597–605, 2022.

[Hosseini *et al.*, 2022b] Hadi Hosseini, Andrew Searns, and Erel Segal-Halevi. Ordinal maximin share approximation for goods. *Journal of Artificial Intelligence Research*, 74:353–391, 2022.

[Hosseini *et al.*, 2022c] Hadi Hosseini, Fatima Umar, and Rohit Vaish. Two for one & one for all: Two-sided manipulation in matching markets. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 321–327. International Joint Conferences on Artificial Intelligence Organization, 7 2022. Main Track.

[Hosseini *et al.*, 2023a] Hadi Hosseini, Zhiyi Huang, Ayumi Igarashi, and Nisarg Shah. Class fairness in online matching. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(5):5673–5680, Jun. 2023.

[Hosseini *et al.*, 2023b] Hadi Hosseini, Joshua Kavner, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Perceived fairness in envy-free allocations of indivisible goods. *arXiv preprint arXiv:2212.04574*, 2023.

[Hosseini *et al.*, 2023c] Hadi Hosseini, Aghaheybat Mammadov, and Tomasz Wąs. Fairly allocating goods and (Terrible) chores. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI-23)*, pages 2738–2746, 2023.

[Hosseini *et al.*, 2023d] Hadi Hosseini, Shivika Narang, and Tomasz Wąs. Fair distribution of delivery orders. *arXiv preprint arXiv:2305.00040*, 2023. to appear in IJCAI-24.

[Hosseini *et al.*, 2023e] Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. Graphical house allocation. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems*, pages 161–169, 2023.

[Hosseini *et al.*, 2023f] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Fairly dividing mixtures of goods and chores under lexicographic preferences. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems*, pages 152–160, 2023.

[Hosseini *et al.*, 2024a] Hadi Hosseini, Medha Kumar, and Sanjukta Roy. The degree of fairness in efficient house allocation. In *Proceedings of the Twenty-Seventh European Conference on Artificial Intelligence (ECAI-24)*, page (forthcming), 2024.

[Hosseini *et al.*, 2024b] Hadi Hosseini, Aghaheybat Mammadov, and Tomasz Wąs. Almost envy-freeness under weakly lexicographic preferences. *arXiv preprint arXiv:2404.19740*, 2024.

[Hosseini *et al.*, 2024c] Hadi Hosseini, Andrew McGregor, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. Tight approximations for graphical house allocation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, page 825–833. International Foundation for Autonomous Agents and Multiagent Systems, 2024.

[Hosseini, 2024] Hadi Hosseini. The fairness fair: Bringing human perception into collective decision-making. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(20):22624–22631, Mar. 2024.

[Kash *et al.*, 2014] Ian Kash, Ariel D Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.

[Kurokawa *et al.*, 2018] David Kurokawa, Ariel D Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, 65(2):1–27, 2018.

[Lipton *et al.*, 2004] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. *Electronic commerce*, pages 125–131, 2004.

[Roth, 1982] Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4):617–628, 1982.

[Searns and Hosseini, 2020] Andrew Searns and Hadi Hosseini. Fairness does not imply satisfaction (student abstract). In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*, volume 34, pages 13911–13912, 2020.

[Shapley and Scarf, 1974] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[Suryanarayana *et al.*, 2022] Sharadhi Alape Suryanarayana, David Sarne, and Sarit Kraus. Explainability in mechanism design: recent advances and the road ahead. In *European Conference on Multi-Agent Systems*, pages 364–382. Springer, 2022.

[Svensson, 1999] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.

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

[Zhou, 1990] Lin Zhou. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990.