Research Statement - Mohamad Latifian

My research lies at the intersection of theoretical computer science, artificial intelligence, and economics, focusing on algorithmic problems that involve agents with diverse preferences who are jointly affected by a common decision. I study collective decision-making mechanisms under informational, strategic, and computational constraints, aiming to understand the trade-offs between efficiency, fairness, and incentives. My work develops theoretical frameworks and algorithms for preference aggregation and resource allocation, motivated by applications such as voting, participatory budgeting, matching markets, as well as decision-making problems arising in multi-agent systems. In this statement, I outline my research contributions and describe my medium- and long-term research agenda.

Collective Decision-Making under Limited Information

A fundamental question in social choice theory and multi-agent systems is how to elicit and aggregate individual agents' preferences to make a prudent collective decision. Such problems arise in a wide range of settings, including political elections, committee selection, resource allocation among automated agents, and the aggregation of outputs produced by multiple algorithms or learning systems. There are two broad approaches to studying collective decision-making rules. A classical line of work is axiomatic, characterizing mechanisms based on the properties they satisfy. A complementary line of work views elicited preferences as partial information about underlying cardinal utilities and aims to optimize social objectives based on these values. While I find both perspectives important and have contributed to each [6], my research primarily focuses on the latter approach, which is commonly studied through the distortion framework.

Distortion. A recurring difficulty in many social choice problems is that the mechanism cannot directly optimize the objective it ultimately seeks to optimize. While social welfare is naturally defined in terms of agents' underlying cardinal utilities, these utilities are rarely fully observable, and the algorithm must instead rely on limited preference information elicited from the agents. This lack of information creates an inherent gap between the outcome that is optimal with full information and what can be achieved by any feasible decision rule. The notion of distortion in social choice was introduced by Procaccia and Rosenschein [18] to formalize this gap by comparing the social welfare achieved by a rule using limited preference information to that of an optimal outcome with full access to agents' utilities. My work uses this framework to characterize both achievable guarantees and fundamental limits across a range of collective decision-making settings, including voting [13, 14, 7], committee selection [5], matching [17, 12], and participatory budgeting [4].

I have proposed and analyzed mechanisms based on both classical elicitation formats such as rankings and approvals[5, 7], and more expressive yet practical approaches, including eliciting intensities within rankings[16, 1] and multi-round adaptive elicitation[8, 4]. These results contribute to a broader understanding of how limited communication can be traded off against efficiency guarantees. Beyond individual domains, I have also investigated fundamental limits of distortion. In particular, I study settings in which mechanisms must optimize multiple objectives simultaneously[14], as well as settings where mechanisms are constrained to be explainable or to use restricted forms of randomization[9].

More broadly, distortion provides a principled lens for analyzing robustness in AI and multi-agent systems, where algorithms must aggregate preferences, predictions, or recommendations under incomplete information. From this perspective, distortion captures a form of misalignment induced by informational constraints, making it directly relevant to the theoretical foundations of AI-mediated decision-making.

Fair Division of Resources

Another central component of my research concerns fair division, where the goal is to allocate resources among agents with heterogeneous preferences in a way that satisfies strong fairness guarantees. I have worked on problems involving both indivisible items and divisible resources (cake cutting), focusing on the algorithmic limits of envy-based fairness notions.

For indivisible goods, a major open problem is the existence of envy-freeness up to any item (EFX) allocations. Despite extensive effort, it is still unknown whether EFX allocations always exist even for additive valuations. My work contributes to this line of research by studying approximate and relaxed variants of EFX. In [11], we provide a simple polynomial-time algorithm achieving the best-known approximation guarantees for EFX, which remain state of the art. In the same work, we introduce envy-freeness up to a random item (EFR), a natural relaxation of EFX. While we design efficient algorithms with strong approximation guarantees for EFR, the existence of EFR allocations remains open in general, mirroring the unresolved status of EFX even for special cases.

I have also studied extensions of fair division models that capture additional structure and more realistic settings. In [3], we consider settings where items have both subjective utilities and exogenous market values, and characterize the limits of achieving fairness with respect to both. In [10], we study an informed online model in which items arrive sequentially and must be allocated irrevocably, and introduce temporal envy-freeness up to one item (TEF1) as a fairness notion that can be maintained over time despite strong online constraints.

More recently, I have turned my attention to the cake-cutting problem. In this problem, the goal is to divide a divisible resource (e.g., cake, land, time) among agents with diverse preference over different parts of that resource in a fair manner. Here, we know that envy-free allocations always exist but it is not clear whether we can find them in polynomial time. Building on recent progress for four agents [15], I am currently working on extending these techniques to the case of five agents with a carefully designed framework which has the longer-term potential of scaling to larger numbers of agents. This work is currently being prepared for submission to ACM EC.

Road Map

Distortion. Building on my prior work on distortion, I plan to pursue several directions that aim to deepen our understanding of trade-offs between information and efficiency. One direction concerns the role of randomization: while randomized mechanisms can achieve substantially lower distortion than deterministic ones, it remains unclear what level of randomness is actually required to obtain these improvements. I aim to characterize the Pareto frontier between randomness complexity and achievable distortion.

A second direction focuses on graph-based objectives in metric distortion. In committee selection and assignment problems where agents and items are embedded in a metric space, prior work has largely focused on utilitarian or egalitarian objectives. However, these settings naturally support richer objectives defined over the induced metric structure, including graph cost functions such as minimum spanning trees or Hamiltonian paths. The study of distortion with respect to such objectives is largely unexplored and opens the door to several interesting problems.

A third direction concerns distortion in continuous metric spaces, where the set of alternatives is a compact and convex domain rather than a finite set. In this setting, the mechanism must adaptively elicit information from agents through comparison or preference queries, to select a point in the domain that minimizes the social cost. I am interested in understanding the trade-offs between the number of queries, the structure of the metric space, and the achievable distortion guarantees. Beyond its theoretical appeal, this model is practically relevant to problems such as parameter tuning and decision aggregation in AI systems, where optimal choices lie in continuous spaces and direct utility elicitation is infeasible.

Holistic approach to social choice. Many problems studied in social choice, such as single-winner voting, committee selection, and matching, share common structure both in how preferences are elicited and in how decision rules are designed. A long-term goal of my research is to develop a unifying model that captures these settings within a single framework and allows results to transfer across problems. One approach I am exploring is modeling a social choice problem as a computational circuit. In this view, the mechanism produces an ordered outcome (e.g., a ranking), which is then processed by an agent-specific module that maps this outcome to a utility value using simple operations such as max, sum, or min. Classic settings arise as special cases, for example, single-winner voting corresponds to taking the utility of the top-ranked alternative, while k-committee selection corresponds to taking the maximum utility over the top k positions. This approach enables the systematic study of rules for different classes of social choice problems, defined by constraints on the underlying circuits, and opens the door to designing mechanisms with provable guarantees that apply uniformly across multiple settings.

Fair division. Two longstanding open problems will be central to my future work in fair division: the existence of envy-freeness up to any item (EFX) allocations for indivisible goods, and the computational complexity of finding envy-free cake divisions. Both problems have remained open for several years and require a systematic approach. For indivisible goods, I plan to further investigate structured classes of instances in which EFX allocations can be shown to exist, such as settings that restrict the number of items each agent values positively or the number of agents who value a given item. In parallel, I will study broader existence guarantees for envy-freeness up to a random item (EFR), with the goal of using these results to better understand the boundary between achievable and unattainable forms of EFX. Progress on these questions can shed more light on characteristics of instances in which finding EFX allocations is more challenging which in turn might lead to counterexamples or existential results for EFX in more general classes of instances.

A complementary direction explores the connection between indivisible fair division and cake cutting. Known existence results for EFX with two and three agents closely mirror techniques used in envy-free cake cutting, suggesting a deeper structural relationship between the two models. I am currently investigating whether recent advances in envy-free cake cutting for four agents can be adapted to establish the existence of EFX allocations for four agents. This direction is technically challenging due to the combinatorial complexity of the four-agent case, but it offers a promising path toward resolving one of the most central open problems in the area.

Strategic agents. A central long-term goal of my research is to develop algorithmic frameworks that remain robust when agents behave strategically. Much of my work so far studies efficiency and fairness under truthful reporting, and a natural next step is to understand how these bounds degrade, or can be preserved, once incentives are taken into account. One direction I plan to pursue is the study of distortion in the presence of strategic agents, including both analyzing the distortion of mechanisms at equilibrium and capturing the price that agents pay (in terms of utility) for their truthfulness. The ultimate goal would be to design rules that have both low distortion and low price of truthfulness for the agents.

In parallel, I am interested in extending my work on fair division to settings where fairness guarantees hold in equilibrium. This direction is directly inspired by the recent work of Amanatidis et al. [2], which studies equilibrium notions of fairness in allocation mechanisms and demonstrates that strong fairness properties can emerge at equilibrium despite strategic behavior. Building on this framework, and on my prior results on approximate EFX and the introduction of EFR, I plan to design allocation mechanisms whose equilibrium outcomes satisfy envy-based fairness notions such as EF1 and EFX. A particular focus will be on settings where EFX allocations are known to exist, including bi-valued instances and the case of three agents, with the goal of understanding whether EFX can be guaranteed at equilibrium in these canonical domains. More broadly, I aim to characterize the computational and informational tradeoffs between incentives, efficiency, and fairness, and to develop polynomial-time mechanisms that reconcile these competing objectives.

My long-term goal is to advance the theoretical foundations of collective decision-making by understanding how informational, strategic, and computational constraints shape what can be achieved algorithmically. I aim to pursue this agenda through a combination of foundational results, algorithmic design, and the development of unifying frameworks that connect social choice, fair division, and mechanism design.

References

[1] Mehrad Abbaszadeh, Ali Ansarifar, Mohamad Latifian, and Masoud Seddighin. Metric distortion with preference intensities. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence (AAAI) (Forthcoming)*.

[2] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness. In *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, 2021.

[3] Siddharth Barman, Soroush Ebadian, Mohamad Latifian, and Nisarg Shah. Fair division with market values. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, pages 13589–13596, 2025.

[4] Mark Bedaywi, Bailey Flanigan, Mohamad Latifian, and Nisarg Shah. The distortion of public-spirited participatory budgeting. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, pages 13605–13613, 2025.

[5] Allan Borodin, Daniel Halpern, Mohamad Latifian, and Nisarg Shah. Distortion in voting with top-t preferences. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 116–122, 2022.

[6] Martin Bullinger, Edith Elkind, and Mohamad Latifian. Towards fair and efficient public transportation: A bus stop model. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 427–435, 2025.

[7] Soroush Ebadian, Aris Filos-Ratsikas, Mohamad Latifian, and Nisarg Shah. Computational aspects of distortion. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1–14, 2023.

[8] Soroush Ebadian, Mohamad Latifian, and Nisarg Shah. The distortion of approval voting with runoff. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1752–1760, 2023.

[9] Soroush Ebadian, Aris Filos-Ratsikas, Mohamad Latifian, and Nisarg Shah. Explainable and efficient randomized voting rules. *Proceedings of 37th Advances in Neural Information Processing Systems (NeurIPS)*, 36, 2024.

[10] Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. Temporal fair division of indivisible items. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 676–685, 2025.

[11] 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 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5355–5362, 2021.

[12] Aris Filos-Ratsikas, Vasilis Gkatzelis, Mohamad Latifian, Emma Rewinski, and Alexandros A Voudouris. Optimal metric distortion for matching on the line. In *Proceedings of 34rd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1–25, 2025.

[13] Mohammad Ghodsi, Mohamad Latifian, and Masoud Seddighin. On the distortion value of the elections with abstention. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1981–1988, 2019.

[14] Vasilis Gkatzelis, Mohamad Latifian, and Nisarg Shah. Best of both distortion worlds. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, pages 738–758, 2023.

[15] Alexandros Hollender and Aviad Rubinstein. Envy-free cake-cutting for four agents. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 113–122. IEEE, 2023.

[16] Anson Kahng, Mohamad Latifian, and Nisarg Shah. Voting with preference intensities. In *Proceedings of 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5697–5704, 2023.

[17] Mohamad Latifian and Alexandros A Voudouris. The distortion of threshold approval matching. *Autonomous Agents and Multi-Agent Systems*, 39(2):1–21, 2025.

[18] Ariel D Procaccia and Jeffrey S Rosenschein. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, pages 317–331. Springer, 2006.