Balancing Fairness, Efficiency and Strategy-Proofness in Voting and Facility Location Problems

Doctoral Consortium

Alexander Lam
UNSW Sydney
alexander.lam1@unsw.edu.au

ABSTRACT

In the field of computational social choice, researchers seek mechanisms that fulfil the notion of *strategy-proofness*, so that it is optimal for agents to simply report their truthful preferences. This notion can be very difficult to achieve, and mechanisms that satisfy this strict constraint often sacrifice ideal properties such as fairness and efficiency. For example, a surjective voting rule satisfying strategy-proofness must be dictatorial, which may be considered unfair and wasteful to the voters, as it only takes into account one voter's preference. Strategy-proof facility location mechanisms are also known to be sub-optimal in terms of fairness, and in some scenarios, efficiency. Focussing on these two areas, we question if strategy-proofness is too strict of a constraint, and to what extent are mechanisms satisfying weaker variations of the property more fair and efficient.

KEYWORDS

Fairness; Efficiency; Strategy-proofness; Voting; Facility Location

ACM Reference Format:

Alexander Lam. 2021. Balancing Fairness, Efficiency and Strategy-Proofness in Voting and Facility Location Problems: Doctoral Consortium. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 2 pages.

1 INTRODUCTION

The field of voting has long existed throughout history, originating as a means to elect public officials based off the citizens' preferences. In a standard voting scenario, each agent submits its preference ordering over the candidates, and a voting rule/social choice function is used to select a single election winner. A classical example of a voting rule is the Borda count, which gives each candidate a score based on their position in the voters' preference orders [5]. A widely studied notion in voting theory is that of *agent strategic behaviour*: agents may misreport their preferences to achieve an outcome that they prefer over the outcome that would be selected if they reported truthfully. A *strategy-proof* voting rule, in which truth-telling is optimal for agents, may be used to prevent such strategic behaviour. However, it is well known that a surjective, strategy-proof voting rule must be dictatorial [9]. A dictatorial voting rule may be considered unfair and wasteful, as we are only considering one voter's preference, hence there is a trade-off between fairness, efficiency and strategy-proofness in voting.

It is worth noting that some voting rules are more manipulable than others, suggesting that we can find voting rules that are difficult to manipulate, and 'fair' and 'efficient' in the non-dictatorial sense. This has led researchers to find alternate approaches to analysing agent strategic behaviour in voting. The computational complexity of computing a strategic manipulation can be considered as a barrier to manipulation. For example,the authors of [4] give sufficient conditions for the manipulation problem to be solvable in polynomial time. It has also been shown that the manipulation problem is NP-Hard for certain rules, such as the second-order Copeland rule, and the Nanson and Baldwin's rules [14]. The hardness of computing a manipulation under partial or zero information has also been considered; certain rules such as the Borda rule and the Copeland rule are *resistant* to manipulations in this setting [7].

In the classic facility location problem, we are tasked with placing a facility to serve a set of agents on a line. The objective is to minimize some monotonic function of the agents' costs, defined individually as each agent's distance from the nearest facility. Although it is intuitive to exemplify this problem geographically, such as by placing a market stall or a warehouse along a street, the problem can also be applied to voting with singled-peaked preferences. By expressing each agent's position on the political spectrum as their location on the line, we can select a politician (facility) whose policies accurately reflect the agents' opinions. Like in the voting setting, an agent may misreport their location to attain a better facility location outcome, so we may seek a strategy-proof facility placement when agent strategic behaviour is a concern.

Research on facility location mechanisms satisfying strategy-proofness was initiated by Moulin in 1980, who showed that placing the facility at the median agent is strategy-proof and optimal for total cost [12]. This concept was popularized by Procaccia and Tennenholtz, who investigated strategy-proof mechanisms for approximating the utilitarian objective of total cost and the egalitarian objective of maximum cost [15]. The compromise of fairness and efficiency resulting from strategy-proofness is more apparent in the 2-facility location problem. In this context, any deterministic strategy-proof mechanism must have an approximation ratio of at least 2 for maximum cost and at least $n - 2$ for total cost [8]. Although the field of voting has a myriad of research on the manipulability of social choice functions, there has been little work in facility location of the same vein. Disregarding strategy-proofness, we also find that facility location mechanisms balancing fairness and efficiency have been largely unexplored.

The fairness and efficiency sacrifice resulting from the constraint of strategy-proofness has been examined in other fields of social choice. In school choice, there is a trade-off between efficiency and strategy-proofness. When examining NYC school matching

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

data, it was found that 1.9 percent of eighth graders (1500 students) could be matched to more preferred schools without harming other students if the restriction of strategy-proofness was relaxed [1]. Consequently, there has been relevant research on weaker, approximate notions of strategy-proofness that are satisfied by efficient mechanisms. For example, *strategy-proofness in the large* (SP-L) only requires truth-telling to be approximately optimal in a large enough market [2]. Another example is partial strategy-proofness, in which truth-telling is required to be the dominant strategy for agents with sufficiently different utilities for any two different objects [11].

The Nash welfare objective function, defined as maximizing the product of agent utilities, has been used in many areas of the social choice literature to find a reasonable compromise between fairness and efficiency. When fairly allocating indivisible goods, the allocation maximizing Nash welfare satisfies envy-freeness up to one good and is Pareto optimal [6]. In a variation of the participatory budgeting model, the optimal Nash solution is ex-ante efficient and satisfies certain fairness guarantees [3]. However, in these contexts, this solution balancing fairness and efficiency fails to meet the constraint of strategy-proofness.

There has been much work on developing weaker notions of strategy-proofness and analysing fair and efficient mechanisms which meet these notions. However, in voting, there remain open questions regarding agent strategic behaviour under partial information, and there is much room for research on balancing fairness, efficiency and strategy-proofness in the facility location game. Our research aims to address these open questions and provide further insights into mechanisms which find a balance between our specified ideal properties.

2 CONTRIBUTION AND FUTURE WORK

We approach our voting and facility location problems from algorithmic and mechanism design perspectives. Our aim is to introduce and analyse weaker forms of strategy-proofness, and design fair, efficient and tractable mechanisms which satisfy these properties.

2.1 Obvious Manipulability of Voting Rules

In many real-life scenarios, agents may be *cognitively-limited* and can lack the contingent reasoning required to properly compute a profitable manipulation [10]. It can be assumed they only know the possible range of outcomes that could arise from each report, possibly by trial and error, or by observing previous mechanism users. We say that a mechanism satisfies *not obvious manipulability* if an agent cannot improve either their best case or worst case outcome by some untruthful report [16]. This concept is much weaker than strategy-proofness, and implies the assumption that agents lack information regarding other agents' reports. We apply this concept to the field of voting, as in reality, ballot information is often withheld from other agents.

Our working paper identifies several classes of voting rules that satisfy not obvious manipulability, and categorizes conditions under which certain rules are obviously manipulable. One of our insights is that certain rules are obviously manipulable when the number of alternatives is relatively large as compared to the number of voters. In contrast to the Gibbard-Satterthwaite theorem, many of the rules we examined are not obviously manipulable. This reflects the relatively easier satisfiability of the notion and the effect of the zero information assumption of not obvious manipulability. We also present algorithmic results for computing obvious manipulations and report on experiments.

2.2 Nash Welfare in the Facility Location Problem

In the facility location problem, the facility is typically located to minimize some monotonic function of the agents' costs, defined as their distance from the nearest facility. In his book "Fair Division and Collective Welfare", Moulin describes an alternate approach, in which the agent costs are converted to utilities, which are maximized [13]. Specifically, he mentions maximizing the Nash collective welfare, a well-studied objective function that provides a compromise between egalitarian and utilitarian measures [6].

We primarily examine the mechanism that places a single facility such that the Nash welfare is maximized. In our working paper, we give a polynomial-time approximation algorithm for this facility placement, and show that a fairness-efficiency compromise is achieved by proving approximation ratio bounds for both optimal egalitarian and utilitarian social welfare objectives. We also demonstrate that this mechanism satisfies certain fairness properties, and aim to analyse the agent strategic behaviour that may occur. When restricting to strategy-proof mechanisms, we find that the results are more negative, as we prove that no mechanism can achieve a constant approximation ratio for optimal Nash welfare. Finally, we give a strategy-proof mechanism with a bounded approximation ratio for the optimal Nash welfare.

REFERENCES

[1] A. Abdulkadiroğlu, P. Pathak, and A. Roth. 2009. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review* 99, 5 (2009), 1954–78.
[2] E. M. Azevedo and E. Budish. 2018. Strategy-proofness in the Large. *The Review of Economic Studies* 86, 1 (2018), 81–116.
[3] H. Aziz, A. Bogomolnaia, and H. Moulin. 2019. Fair Mixing: the Case of Dichotomous Preferences. In *Proc. of the 20th ACM-EC Conference*. 753–781.
[4] J. Bartholdi, C. Tovey, and M. Trick. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6, 3 (1989), 227–241.
[5] D. Black. 1986. Borda, Condorcet and Laplace. In *The Theory of Committees and Elections*. Chapter 18, 156–162.
[6] I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. 2016. The Unreasonable Fairness of Maximum Nash Welfare. In *Proc. of the 17th ACM-EC Conference*. 305—322.
[7] V. Conitzer, T. Walsh, and L. Xia. 2011. Dominating Manipulations in Voting with Partial Information. In *Proc. of the 25th AAAI Conference*.
[8] D. Fotakis and C. Tzamos. 2014. On the Power of Deterministic Mechanisms for Facility Location Games. *ACM Transactions on Economics and Computation* 2, 4, Article 15 (2014).
[9] A. Gibbard. 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* (1973), 587–601.
[10] S. Li. 2017. Obviously Strategy-Proof Mechanisms. *American Economic Review* (2017), 3257–3287.
[11] T. Mennle and S. Seuken. 2021. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory* 191 (2021), 105144.
[12] H. Moulin. 1980. On Strategy-Proofness and Single Peakedness. *Public Choice* 35, 4 (1980), 437–455.
[13] H. Moulin. 2003. *Fair Division and Collective Welfare*. MIT Press.
[14] N. Narodytska, T. Walsh, and L. Xia. 2011. Manipulation of Nanson's and Baldwin's Rules. In *Proc. of the 25th IJCAI*.
[15] A. D. Procaccia and M. Tennenholtz. 2013. Approximate Mechanism Design without Money. *ACM Transactions on Economics and Computation* 1, 4, Article 18 (2013).
[16] P. Troyan and T. Morrill. 2020. Obvious manipulations. *Journal of Economic Theory* 185 (2020).