Modeling Human Interactions: Facets of Algorithmic Game Theory and Computational Social Choice

(Doctoral Consortium)

Omer Lev
School of Computer Science and Engineering, Hebrew University
Jerusalem, 91904, Israel
omerl@cs.huji.ac.il

Introduction

Modeling human activities and interactions — especially when participants exhibit rational behavior — has long been a research goal in economics, mathematics (specifically, game theory), and computer science. In computer science, this was initially framed in the classical sense of artificial intelligence (AI), in which attempts were made to mimic human behavior. However, in the past 20 years, especially since the advent of the World Wide Web and the explosion in the number of people interacting using various computing systems, research in the field has expanded dramatically. In particular, there is a significant and growing interest in the types of interactions commonly seen on the web:

- **Auctions:** Auction-like mechanisms are at the heart of the ad-generated profit model powering most websites, as the major ad-selling networks perform virtual auctions to sell ad space. Furthermore, auctions model many human interactions that are non-commercial: for example, all-pay auctions, in which all participants pay their bid, regardless of who wins, can be used to model a job market, in which all job-seekers are, in effect, bidding for a job, and putting in effort to get it, regardless of whether or not they eventually succeed.

- **Voting:** When several people, or agents, need to decide on what option to pursue, some mechanism is needed to decide among their respective preferences. Such elections, whether they involve a large or small number of participants, occur all around us, outside political contexts, like when choosing a restaurant for dinner. While it may seem straightforward to simply choose an election method, different voting rules have different properties and potentially different outcomes; in addition, the choice is complicated by the proven lack of a strategy-proof voting mechanism.

- **Crowdsourcing:** When multiple people are working together towards a shared goal, various problems arise: how to incentivize massive participation, how to decide among differing options, and how to manage groups and coalitions of participants to work together. Such issues may involve auctions (particularly, all-pay auctions), voting, and coalitional game theory.

All these human (and human-machine) interactions open a wide field of research, encompassing both theoretical aspects and empirical ones. On the theoretical side, examining the properties of existing mechanisms, exploring extensions and complications of mechanisms, and understanding what are the optimal mechanisms to choose for various tasks are a few key research directions. The properties are, in many cases, similar across the various interactions; we seek to understand how susceptible these mechanisms are to manipulations (participants being untruthful), and what are the equilibria points of these mechanisms. Furthermore, we wish to examine what the characteristics of these equilibria are, and how good they are for participants (or for everyone) vs. more optimal states.

Enhancing this analysis is the empirical side, examining situations more akin to "real-life", which may be harder to solve mathematically. This is done both by running experiments using people, and by running simulations (particularly on complex models, for which computation may be worst-case intractable). Both are alternative paths to examine the quality — and soundness — of the theoretical models, as well as examining some of their properties.

Research To-Date & Ongoing

Complexity of Manipulating Voting Rules

While the Gibbard-Satterthwaite theorem shows that all voting rules are susceptible to manipulation, the computational social choice community has shown that many voting rules are computationally hard to manipulate. However, NP-hardness is a "worst-case" result, indicating that some cases are very hard to compute. Some research has been done on trying to find a more "common-case" alternative, but another branch of research tries to find polynomial algorithms that reach good approximations of complex voting rules.

While scoring-rules have some very good approximations using greedy algorithms, the Maximin voting rule turns out to only have a 2-approximation using a greedy approximation algorithm. We have shown that it is not the best approximation: we proved a 1.5-approximation of Maximin

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

(due to the special case of 3 voters, it is technically a $\frac{3}{5}$-approximation) using a more sophisticated algorithm.

Iterative Voting Processes

When voters have little information on others' preferences, they cannot manipulate and reach a Nash equilibrium. However, certain voting scenarios (small groups; several web-based voting systems) allow voters to change their minds and change their votes, enabling them to strategize and modify their vote according to the choices made by the other participants. A basic question in such scenarios is whether or not this process ever ends and converges to a Nash equilibrium. Previous work established that when voting using plurality and a linear-ordered tie-breaking rule, the iterative process indeed converges when participants use a myopic best-response strategy. In further work, we showed the critical importance of the tie-breaking rule, establishing that for any scoring rule (and, in addition, maximin) there is a tie-breaking rule that would prevent its convergence. Furthermore, we established that there are voting rules for which no convergence is possible (Borda), and showed that the veto voting rule converges when using a linear-ordered tie-breaking rule.

Empirical Voting Analysis

Due to the nature of voting games, they have an extremely large number of Nash equilibria, despite many of them being "unnatural" (e.g., if everyone has the same least-favored candidate, it is still a Nash equilibrium if everyone voted for it), easily reaching hundreds of thousands even with single digit number of voters and candidates. We showed that if we add a very small preference for truthfulness (i.e., if a voter cannot influence the result, he defaults to voting for his truthful preferences), the number of equilibria in plurality voting drops dramatically (in some cases, games have no equilibrium at all). This, we believe, better resembles "real-life" behavior in similar circumstances.

Furthermore, we showed that truthfulness incentive influences the remaining equilibria so that they have certain desirable characteristics: many equilibria result in the truthful winner being selected; many equilibria result in a Condorcet winner; and generally, the average score of the winning candidates is significantly high, with many games not having any equilibria that results in a low-ranking candidate.

We also expanded our model to examine Bayes-Nash equilibrium, and showed that many equilibria have a small support of only 2 or 3 candidates. This led us to hypothesize — and prove — a proposition, showing that for a small enough truthfulness incentive, for every 2 candidates where one does not pareto dominate the other, there is a Bayes-Nash equilibrium with the support of just these 2 candidates.

Collusions in All-Pay Auctions

As detailed above, all-pay auctions model a wide variety of situations, both online (crowdsourcing participation) and offline ("employee-of-the-month" competitions). While collusions among bidders has been analyzed for several types of auctions, it has not been dealt with in regard to all-pay auctions. We examined its properties, showing how mergers (when bidders' cooperation is publicly known) slightly increase social welfare, but the additional profit goes, in expectation, solely to the auctioneer. In contrast, collusions, while often reducing social welfare, promote a more egalitarian division, as the expected profit of the colluders is positive (without collusion, it is 0). We proved that having several colluders (each as small as possible) is good (up to a certain amount) for the auctioneer, as it gives it a higher "minimum price", which grows in significance as the number of bidders increases and the expected bid drops.

Future Work & Planned Research

Naturally, there are many potential paths to choose from, and as work progresses, directions may open that do not seem interesting or feasible from the current vantage point. However, some possible interesting future work can be suggested now: examining scenarios that try to take an abstract model and add to it some of the complexities that might help make the model more applicable to real-life situations.[1]

Some future research lies in expanding our previous work. For example, although we examined the effects of the truthfulness incentive in plurality, we have only started to expand our research to further voting systems, and the ability to have an overall view of all the Nash equilibria in a given voting game (with the truthfulness incentive), has, in our view, great potential to present questions that we have yet to raise, enabling a more holistic view of the solution concepts in elections. Also, our work on iterative voting should be expanded to examine more strategies (besides best-response), more voting systems, and extended beyond myopic analysis.

Furthermore, due to their particular peculiarities, research on all-pay auctions tended to ignore many common human behaviors, such as collusions and mergers. While we have begun work in this area , there is still much to do. It is needed, in our view, to build on our work and research more realistic models, which have, for example, a probabilistic relation between the effort levels and winning, such as the Tulloc contest function. We have only briefly touched the middle-term, where agents have partial information regarding the coalitions of other bidders. A more detailed analysis of the middle-term dynamics is required in order to have a complete picture of the impact of bidder collaboration in all-pay auctions. Additionaly, while much analysis assumes a constant number of bidders, which is common knowledge, this is not realistic in many anonymous online settings.

In cooperative game theory, there has long been a focus on achieving a "grand coalition" in which all parties eventually cooperate. However, in many cases (e.g., political settings in multi-party parliamentary democracies) a grand-coalition is not just unachievable — it is not even desirable by the various participants ex-ante. A better model of coalitions in needed, including a better, pragmatic concept of stability. The current binary definition, of whether a core exists or not, is not helpful in the many situations where there is no core. Therefore, a concept of stability that ranks coalitions, so a hierarchy of coalitions exists, is needed. Using, in addition, a simplified concept of preferences among coalitions (perhaps a directed graph, indicating the change in the level of satisfaction for party A when party B joins a coalition in which it is a member), may yield interesting and useful results.

[1]In a sense, many game theory advances stem from a similar drive — voting games, auctions, and coalitions are all specific families of problems of the general family of normal-form games. They were separated by expanding the models of these problems, inspired by some real world complexities.