

Decision Making Under Uncertainty: Social Choice and Manipulation *

Nicholas Mattei

Department of Computer Science
University of Kentucky
Lexington, KY 40506, USA

My research seeks insight into the complexity of computational reasoning under uncertain information. I focus on preference aggregation and social choice. Insights in these areas have broader impacts in the areas of complexity theory, autonomous agents, and uncertainty in artificial intelligence.

Motivation: Planning and reasoning in nondeterministic settings is something that people take for granted every day. We do not know for certain that each small action we choose will succeed or fail, if the actions we choose will lead us to catastrophic consequences or land us safely on the other side of the street. The ability to reason in a domain where actions are not guaranteed to succeed is something that humans do fairly well and machines do not.

The field of social choice allows us a rich set of domains and problems within which we can work. A central question of social choice is: how do we aggregate a (possibly) contradictory set of individual preferences and/or observations into an appropriate global decision? We focus on the question of *manipulation* of social choice functions when the individual agents' preferences are represented as probability distributions rather than a set of deterministic preferences. This notion of uncertainty has been introduced hesitantly, if at all, in the existing literature. We wish to fill this gap.

Background: The field of preference aggregation manipulation stems from that of social choice. Building on the work of Arrow [1], the Gibbard–Satterthwaite Theorem shows that any aggregation system, meeting a set of simple fairness conditions, can be manipulated by non-truthful voting [7, 12]. This was extended again by the Duggan–Schwartz Theorem to an even larger set of aggregation methods [4]. These results tell us that we cannot devise a “good” preference aggregation scheme that is immune to manipulation. This implies that groups can never come to provably fair, non-manipulated agreements. However, in the early 1990s, Bartholdi et al. proposed the idea of protecting the aggregation schemes through computational complexity [2]. The idea, much like cryptography, is: if it is difficult to compute a manipulation scheme then it is unlikely that there will be manipulation. The ComSoc community seeks to classify aggregation systems in terms of their susceptibility to manipulation.

There is a rich literature on the computational complexity of elections [6], and on the worst-case complexity of manip-

ulation of voting mechanisms by voters [3]. However, all of these studies are built around the assumption of perfect information on the part of the manipulator. In fact, there has been little work on stochastic models of elections. The closest notion is that of a *possible winner*, a notion intrinsic to reasoning under uncertainty introduced and studied by Konczak and Lang [9]; this works answer complexity questions when voters are defined by their (possibly incomplete) preference profiles over a set of outcomes. The question of complexity of manipulation under stochastic assumptions remains open.

Sports tournaments represent another domain where it is natural to express winners and losers in terms of probabilities of outcomes. We use this as a motivating example to study the related problems of manipulation of elimination tournaments and round-robin tournaments. Additionally, results for tournament rules can be mapped to their corresponding voting rules. Under the cup rule, also known as a knock-out tournament, a winner is decided through repeated head-to-head comparisons with pairings organized in a tree. This rule has been significantly studied with respect to manipulation through control of the ordering of the head-to-head pairing. This form of manipulation, considered in the deterministic case by Lang et al. [10] (for the problem of sequential majority voting) and in the stochastic case by Hazon et al. [8], applies to setting the seeding, or ordering, of the cup. Russell and Walsh [11] recently introduced the notion of manipulation through team influence and we extend this work into uncertain domains.

My Work: I seek to gain a deeper understanding of the impact that uncertainty has on social choice and preference aggregation domains. Much of the existing literature has focused on situations where complete information is present. While this assumption is appropriate for investigations about worst case scenarios, I feel applications that include the presence of uncertain information is an important area of study. Many actors have access to accurate probabilistic information regarding agents' preferences. I study and classify what, if any, effect this knowledge has on an agent's ability to reason effectively about manipulation schemes.

Completed Work: We proposed a novel and flexible model to represent majority voting with uncertain information. Given a set of voters, their individual preferences are represented as probability distributions over a set of issues, their prices for changing their preferences, and a budget, we

*This work is supported by NSF EAGER grant CCF-1049360

classified the complexity of finding efficient bribery schemes [5]. We developed three different ways in which an outside actor can channel money to the individual voters. We also established three different criteria to evaluate the outcome of a vote. We showed that, depending on the particular combination of evaluation and bribery models chosen, the nine problems range in complexity from polynomial time to NP-complete. This difference reveals that modeling choices can have significant effects on the complexity of calculating efficient bribery schemes.

We also proposed a novel and flexible model to represent sports tournaments and competitions. Given a set of teams and their probabilities of each possible win, along with prices for decreasing their competitive output (purposefully losing or underperforming in a match), we classified the complexity of finding efficient bribery schemes for three common types of sports tournaments. We formulated and provided a complexity analysis for these three tournaments over five different problem variants. The evaluation complexity of these problems range from polynomial time to NP^{PP} . Early results show that in some cases the added uncertainty increases the complexity of manipulating sports tournaments while in some cases it does not. While this increase in complexity is not uniform across all tournament types, the change shows strong evidence that reasoning in domains with uncertain data leads to an increase in reasoning complexity.

The primary results at this point in the work are: (1) novel models to represent uncertain information in preference aggregation domains, (2) complexity analysis for these models, and (3) a still unclear picture as to the effects of uncertainty on preference aggregation schemes. We have shown that while under some voting and aggregation rules complexity remains unchanged with relaxed assumptions in other domains complexity increases. The initial work on majority voting domains was presented at ADT-09 [5] and a longer version has been submitted to a journal. The initial work on sports tournaments has been presented at a Dagstuhl Seminar on Social Choice and is being prepared for submission to conference.

Future Work: Our initial work ignored the complexity of interactions between agents. We are extending our model to take into account voting in the context of social networks and interdependent agents. This change will add more modeling fidelity to our existing models of voting in stochastic domains; we will classify these models in terms of the complexity of manipulation. We plan to test heuristics and approximation schemes for the NP-complete and harder problems that we have classified during our early work.

We will also look at different ways that time can affect the process by extending our model to allow for influence over time. We will model and classify the effects of limiting the amount of contributions in a given time period, or by looking at the gradual corrosion of influence over time (agents who do not stay influenced). We will also study the scenarios where the manipulator must learn a model of an agent by watching their response to input on earlier votes, in order to predict their later responses.

In addition to our theoretical work we are undertaking empirical studies of voting. To facilitate an empirical study we

are constructing a large dataset of “elections” from real preference data. We have several hundred thousand individual elections with tens to tens of thousands of voters. We expect that analysis of this dataset will provide valuable insights into the nature of elections. Many studies that seek to classify forms and types of elections suffer from a dearth of real world data. We expect to address this gap with our current study.

This work has left us with an unclear picture of the effects of uncertainty on voting rules. We are looking for some unifying factors within the analyses of deterministic and probabilistic voting models. This would allow us to make firm statements as to the change in complexity under uncertainty.

References

- [1] K. J. Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 1951.
- [2] J. Bartholdi, C. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.
- [3] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.
- [4] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-satterthwaite generalized. *Social Choice and Welfare*, 17:85–93, 2000.
- [5] G. Erdélyi, H. Fernau, J. Goldsmith, N. Mattei, D. Raible, and J. Rothe. The complexity of probabilistic lobbying. In *1st Int. Conf. on Algorithmic Decision Theory (ADT 2009)*, pages 86–97. Springer-Verlag, 2009.
- [6] P. Faliszewski, L. Hemaspaandra, and E. Hemaspaandra. How hard is bribery in elections? *Journal of AI Research*, 35:485–532, 2009.
- [7] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41:587–601, 1973.
- [8] N. Hazan, P. Dunne, S. Kraus, and M. Wooldridge. How to rig elections and competitions. In *Proc. of the 2nd Int. Workshop on Computational Social Choice (COMSOC 2008)*, 2008.
- [9] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *IJCAI’05: Proceedings of the 19th Int. Joint Conf. on Artificial Intelligence: Workshop on Advances in Preference Handling*, 2005.
- [10] J. Lang, M. Pini, F. Rossi, K. Venable, and T. Walsh. Winner determination in sequential majority voting. In *IJCAI’07: Proceedings of the 20th Int. Joint Conf. on Artificial Intelligence*, pages 1372–1377, 2007.
- [11] T. Russell and T. Walsh. Manipulating tournaments in cup and round robin competitions. In *1st Int. Conf. on Algorithmic Decision Theory (ADT 2009)*, pages 26–37. Springer-Verlag, 2009.
- [12] M. Satterthwaite. Strategy-proofness and Arrow’s Conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–216, 1975.