

Mechanism Design with Uncertainty

Taiki Todo

Kyushu University, Japan
todo@inf.kyushu-u.ac.jp

Abstract

My research is summarized as *mechanism design with uncertainty*. Traditional mechanism design focuses on static environments where all the (possibly probabilistic) information about the agents are observable by the mechanism designer. In practice, however, it is possible that the set of participating agents and/or some of their actions are not observable a priori. We therefore focused on various kinds of uncertainty in mechanism design and developed/analyzed several market mechanisms that incentivise agents to behave in a sincere way.

1 Introduction

The rapid development of deep neural network and analysis technique for petabytes of data will make autonomous agents on vehicles, including automated driving cars and delivery drones, more practical in near future. The researches on multi-agent systems and their coordination is, therefore, increasing its importance, in order to prevent accidents and unexpected disasters in such environment where human and agents are interacting. In particular, analyzing what are their objectives, how they behaves when they are selfish, and what should be the result of such selfish behaviors are quite demanding. For these analysis, the fundamental theory of multi-agent systems, especially game theory and market design, will become much more important.

My research is summarized as *mechanism design with uncertainty*. Traditionally, the theory of mechanism design have been developed, in the literature of micro-economics, for targeting static environments where all the (possibly probabilistic) information about the agents, such as the set of participating agents and their action spaces, are observable by the mechanism designer. However, in practice, it is very natural to consider the situations where the set of participating agents and/or their action spaces are not observable a priori. We therefore focused on various kinds of uncertainty in mechanism design and developed/analyzed several market mechanisms that incentivise agents to behave in an expected/sincere way.

2 Identity Uncertainty in Resource Allocation

Resource allocation, including auction and matching, is one of the most promising application fields of the theory of mechanism design. Under some anonymous environment, such as the Internet, each agent might pretend to be multiple agents by using different email addresses. Such an attack has been called *false-name manipulations* [Yokoo *et al.*, 2004]. Designing and analyzing resource allocation mechanisms that are resistant to false-name manipulations (false-name-proof mechanisms in short), has been a big challenge in the literature.

One of the promising approach for designing and analyzing mechanisms is to give a characterization of them by a set of desiderata. In single-item auction, it is well-known that the allocation rule of a truthful auction mechanism must be monotonic [Myerson, 1981]. In combinatorial auction (and for more general multi-dimensional settings), a property called weak monotonicity is known to characterize allocation rules of truthful mechanisms [Bikhchandani *et al.*, 2006]. Most of our works consider such a characterization approach for false-name-proof mechanisms.

In our AAMAS-09 paper [Todo *et al.*, 2009], we gave a complete characterization of the allocation rules of false-name-proof combinatorial auction mechanisms by a property called sub-additivity. It requires that the winning value for a bundle of goods when no false-name manipulation is conducted must be smaller than or equal to the sum of the winning values for each single good when a false-name manipulation is conducted. By utilizing this results, we pointed out the failure of two existing combinatorial auction mechanisms that had been presented in past AAMAS and believed to be false-name-proof.

In real electronic markets, each bidder arrives and departs over time. Thus, such a mechanism that must make decisions dynamically without knowledge of the future is called an *online mechanism*. In our AAMAS-12 paper [Todo *et al.*, 2012], we formalize false-name manipulations in online mechanisms and identify a simple property that characterizes the allocation rules of false-name-proof online auction mechanisms. In an online mechanism, it is very unlikely that the mechanism designer knows the number of bidders beforehand or can verify the identity of all of them. Thus, a bidder can easily submit multiple bids using different identifiers (e.g., different e-mail addresses). To the best of our knowl-

edge, this is the first work on false-name-proof online mechanisms. Furthermore, we develop a new false-name-proof online auction mechanism for k identical items.

Matching has many real life applications, such as school choice and hospital residency assignment. In the framework of matching, even if each agent does not want more than one object, it is not obvious whether strategy-proof mechanisms such as top-trading-cycles (TTC) and deferred acceptance (DA) necessarily satisfy false-name-proofness. For example, a student may represent him/herself differently to one school than to another school, by using fake accounts, based on what these schools consider important. Furthermore, people applying for political asylum under multiple identities is a recognized problem (here the countries take the roles of schools). In our AAMAS-13 paper [Todo and Conitzer, 2013], we show that DA and TTC are both false-name-proof in a weaker sense. Furthermore, DA also satisfies a stronger version of false-name-proofness, while TTC fails to satisfy it when a natural assumption does not hold.

Cake cutting has been recognized as a fundamental model in fair division. In our AAMAS-15 paper [Tsuruta *et al.*, 2015], we first show that no randomized false-name-proof cake cutting mechanism simultaneously satisfies ex-post envy-freeness and Pareto efficiency. We then propose a new randomized mechanism that is optimal in terms of worst-case loss among those that satisfy false-name-proofness, ex-post envy-freeness, and a new weaker efficiency property. However, it reduces the amount of allocations for an agent exponentially with respect to the number of agents. To overcome this negative result, we also provide another new cake cutting mechanism that satisfies a weaker notion of false-name-proofness, as well as ex-post envy freeness and Pareto efficiency.

When the set of agents are uncertain, it is quite hard for a mechanism designer to observe which group of agents are in a cooperation relation. In such a situation, an extended notion of fairness make sense; no group of agents envies any group of agents. In our IJCAI-11 paper [Todo *et al.*, 2011b] and COCOON-19 paper [Todo *et al.*, 2019], we consider such a group envy in combinatorial auction. In particular, we investigate the worst case efficiency loss of combinatorial auction mechanisms that satisfy such a group fairness property in the framework of approximate mechanism design.

Furthermore, when the set of agents are uncertain, a mechanism must be defined for any number of agents. It is then desirable that the more bidders participate, the more a mechanism designer earns. Such a property is called *revenue monotonicity*. In our IAT-10 paper [Todo *et al.*, 2010], a characterization of truthful and revenue monotonic combinatorial auction mechanisms is revealed, combined with the characterization of truthful mechanisms.

3 Identity Uncertainty in Social Choice

Voting is another well-studied problem in the literature of mechanism design. Conitzer [Conitzer, 2008] initiated false-name-proof mechanism design for voting. He considered general voting situation where no geometrical structure among candidates exists. My research therefore therefore fo-

cused on more specific situations where voters' preferences are assumed to be single-peaked.

A well-known median voter scheme [Moulin, 1980] has been known to the only class of voting mechanisms that satisfy a certain set of desirable properties when voters' preferences are assumed to be single-peaked. In our AAMAS-11 paper [Todo *et al.*, 2011a], we showed that the class of voting mechanisms satisfying desirable properties drastically shrinks when false-name manipulations are possible. Intuitively, such a voting mechanism has to have a pre-determined alternative and returns it as a solution if and only if it is one of the optimal for voters; otherwise the mechanism can return a different solution. Such a voting mechanisms are also known as *target rules* in the literature.

When the candidates have a more complex structure, which can typically be represented as a discrete graph, and voters' preferences are single-peaked on it, a natural question is whether there exists a voting mechanism that is false-name-proof and Pareto efficient. Our AAMAS-19 paper [Nehama *et al.*, 2019] found a class of graphs on which such desirable voting mechanisms exist. Furthermore, our ECAI-20 paper [Todo *et al.*, 2020] investigates whether well-known parameterized graph structures, such as grid and cycle, guarantees the existence of such desirable voting mechanisms. Our PRIMA-19 paper [Okada *et al.*, 2019] established an automated mechanism design framework to verify under which graph structure a false-name-proof and Pareto efficient voting mechanism exists.

In our AAI-16 paper [Sonoda *et al.*, 2016], we considered an extended facility location model where a mechanism selects a pair of locations on the interval and an agent's utility is determined by the location of the better one (typically the closer to her ideal point). We first characterize possible outcomes of mechanisms that satisfy false-name-proofness, as well as some mild conditions. By extending the result, we completely characterize the class of false-name-proof mechanisms when locating two facilities on a circle. We then clarify the approximation ratios of the false-name-proof mechanisms on a line metric for the social and maximum costs.

When the number of voters is not observable a priori and multiple choices must be made continuously while agents dynamically arrive/leave, a social choice function needs to give each agent an incentive to sincerely report her existence. In our AAMAS-18 paper [Wada *et al.*, 2018], we investigate facility location models with variable and dynamic populations. For a static, i.e., one-shot, variable population model, we provide a necessary and sufficient condition for a social choice function to satisfy participation, as well as truthfulness, anonymity, and Pareto efficiency. The condition is given as a further restriction on the median voter schemes. For a dynamic model, we propose an online social choice function, which is optimal for the total sum of the distances between the choices in the previous and current periods, among any Pareto efficient functions.

4 Endowment Uncertainty

Barter exchange is one of the most appealing problem of mechanism design, due to its application to US Nation-Wide

kidney exchange program. In practice, agents' initial endowments are private, and the mechanism designer could not force them to bring all their endowments to the market.

Among the researches on mechanism design for barter exchange, Sönmez [Sönmez, 1999] showed that there exists no exchange mechanism that is truthful, individually rational and Pareto efficient when some agent have more than one object, the preference domain is rich enough, and some natural assumptions are satisfied. We therefore considers the cases where some of these requirement does not hold.

In our AAI-14 paper [Todo *et al.*, 2014], we considered various restrictions on the preference domain. We still focused on strict preferences and assumed that all the objects in the market are treated symmetrically, but three domain restrictions are introduced. Furthermore, we consider the case where agents' endowments are private and the mechanism designer cannot observe which objects are owned by which agents. As fundamental results, we first show that the revelation principle holds under a natural assumption and that truthfulness and Pareto efficiency are incompatible even under the smallest domain. We then proposed, for each domain, a class of exchange mechanisms that is truthful and individually rational.

In our another AAI-14 paper [Sonoda *et al.*, 2014], we allows preferences to be indifferent over some bundles of objects, i.e., the preferences are no longer assumed to be strict. The purpose of this paper is to discuss the compatibility of truthfulness, individual rationality, and Pareto efficiency when preferences may contain indifferences. We proposed a domain of agents' preferences called top-only preferences and showed that under the domain, there is an exchange mechanism that is truthful, individually rational, and Pareto efficient. We also find a class of preference domains called m -chotomous preferences, under which the three properties cannot be simultaneously satisfied by any exchange mechanism for any $m \geq 3$.

In our IJCAI-15 paper [Sun *et al.*, 2015], we study the exchange of indivisible objects where agents' possible preferences over the objects are strict and share a common structure among all of them, which represents certain level of asymmetry among objects. A typical example of such an exchange model is a re-scheduling of tasks over several processors, since all task owners are naturally assumed to prefer that their tasks are assigned to fast processors rather than slow ones. We focus on designing exchange mechanisms that simultaneously satisfy strategy-proofness, individual rationality, and Pareto efficiency. We first provide a general impossibility result for agents' preferences that are determined in an additive manner, and then show an existence of such an exchange rule for further restricted preferences. We finally find that for the restricted case, a previously known equivalence between the single-valuedness of the strict core and the existence of such an exchange rule does not carry over.

When each agent has more than one object, he or she will face an optimization problem of deciding which subset of objects to bring to the market. In our JAIR paper [Fujita *et al.*, 2018], we showed that for each agent, finding one of the best, i.e., the utility-maximizing, subset of objects to bring to the market is NP-hard under an TTC-based exchange

mechanism. This means that, even though the exchange mechanism is not truthful, computationally bounded agents such as real humans would prevent from trying to cheat the mechanism. Furthermore, the TTC-based mechanism is core-selecting when preferences are conditionally lexicographic.

5 Future Directions

The general research question in my future directions is "for a given problem and for any given uncertainty, what is a necessary and sufficient condition for the optimal mechanisms to be manipulation-resistant, e.g., false-name-proof?" If the optimal mechanism cannot be manipulation-resistant, I will then investigate approximation mechanisms and manipulation complexity. Besides the uncertainties mentioned above, I will also investigate on the following two uncertainties: (1) uncertainty on agents' preferences, and (2) uncertainty on outcome spaces. Also I will apply the developed theory to practical domains.

(1) Uncertainty on Agents' Preferences: If a mechanism runs only one-time, as the traditional theory of mechanism design assumes, the problem is quite simple. In practice, however, a mechanism needs to used more than one time, as agents usually arrives and departs over time. For such dynamic environments, it is also natural to consider that agents' preferences are also dynamically varying. For instance, in the first fiscal year of a project, a PI would like to purchase many equipment, and from the second fiscal year, he/she would prefer using the funding to hire academic staffs. Under such dynamic preferences, is the optimal mechanism still manipulation-resistant? If not, can we find approximation mechanisms? What if we consider different problems, such as auction, voting, and barter exchange?

(2) Uncertainty on Outcome Spaces: The set of possible outcomes may also vary according to agents' preferences. For example, for students-school matching problems, if many students apply for a certain school, the government (or some other agency) may also offer additional budget for the school, in order to maximize the students' welfare. If such a flexible support is available, can we appropriately incentivise students to report their sincere preferences?

Application to Practical Domains: As a first step, we are considering to apply some matching algorithm to car-sharing. More precisely, consider assigning the usage right of electric vehicles to potential users by some matching algorithm. In practice, population and arrival time of such potential users are uncertain. We will then evaluate the performance of such algorithms in practice, as well as the ratio of users' participation, e.g., the percentage of registered users to submit the queries for using electric vehicles.

Acknowledgments

This work was partially supported by JSPS Kakenhi Grant Numbers JP17H04695 and JP20H00587. The author thanks all the co-authors, collaborators, and discussants.

References

- [Bikhchandani *et al.*, 2006] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [Conitzer, 2008] Vincent Conitzer. Anonymity-proof voting rules. In *the Proceedings of the Fourth International Workshop on Internet and Network Economics (WINE-08)*, pages 295–306, 2008.
- [Fujita *et al.*, 2018] Etsushi Fujita, Julien Lesca, Akihisa Sonoda, Taiki Todo, and Makoto Yokoo. A complexity approach for core-selecting exchange under conditionally lexicographic preferences. *Journal of Artificial Intelligence Research (JAIR)*, 63:515–555, 2018.
- [Moulin, 1980] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [Myerson, 1981] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [Nehama *et al.*, 2019] Ilan Nehama, Taiki Todo, and Makoto Yokoo. Manipulations-resistant facility location mechanisms for ZV-line graphs. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-19)*, pages 1452–1460, 2019.
- [Okada *et al.*, 2019] Nodoka Okada, Taiki Todo, and Makoto Yokoo. Sat-based automated mechanism design for false-name-proof facility location. In *Proceedings of the 22nd International Conference on Principles and Practice of Multi-Agent Systems (PRIMA-19)*, pages 321–337, 2019.
- [Sönmez, 1999] Tayfun Sönmez. Strategy-proofness and essentially single-valued cores. *Econometrica*, 67(3):677–689, 1999.
- [Sonoda *et al.*, 2014] Akihisa Sonoda, Etsushi Fujita, Taiki Todo, and Makoto Yokoo. Two case studies for trading multiple indivisible goods with indifferences. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 791–797, 2014.
- [Sonoda *et al.*, 2016] Akihisa Sonoda, Taiki Todo, and Makoto Yokoo. False-name-proof locations of two facilities: Economic and algorithmic approaches. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI-16)*, pages 615–621, 2016.
- [Sun *et al.*, 2015] Zhaohong Sun, Hideaki Hata, Taiki Todo, and Makoto Yokoo. Exchange of indivisible objects with asymmetry. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 97–103, 2015.
- [Todo and Conitzer, 2013] Taiki Todo and Vincent Conitzer. False-name-proof matching. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-13)*, pages 311–318, 2013.
- [Todo *et al.*, 2009] Taiki Todo, Atsushi Iwasaki, Makoto Yokoo, and Yuko Sakurai. Characterizing false-name-proof allocation rules in combinatorial auctions. In *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-09)*, pages 265–272, 2009.
- [Todo *et al.*, 2010] Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo. Characterization of revenue monotonicity in combinatorial auctions. In *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT-10)*, pages 383–390, 2010.
- [Todo *et al.*, 2011a] Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo. False-name-proof mechanism design without money. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-11)*, pages 651–658, 2011.
- [Todo *et al.*, 2011b] Taiki Todo, Runcong Li, Xuemei Hu, Takayuki Mouri, Atsushi Iwasaki, and Makoto Yokoo. Generalizing envy-freeness toward group of agents. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 386–392, 2011.
- [Todo *et al.*, 2012] Taiki Todo, Takayuki Mouri, Atsushi Iwasaki, and Makoto Yokoo. False-name-proofness in on-line mechanisms. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12)*, pages 753–762, 2012.
- [Todo *et al.*, 2014] Taiki Todo, Haixin Sun, and Makoto Yokoo. Strategyproof exchange with multiple private endowments. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI-14)*, pages 805–811, 2014.
- [Todo *et al.*, 2019] Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo. Competitive auctions and envy-freeness for group of agents. In *Proceedings of the 25th International Computing and Combinatorics Conference (COCOON-19)*, pages 541–553, 2019.
- [Todo *et al.*, 2020] Taiki Todo, Nodoka Okada, and Makoto Yokoo. False-name-proof facility location on discrete structures. To appear in *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-20)*, 2020.
- [Tsuruta *et al.*, 2015] Shunsuke Tsuruta, Masaaki Oka, Taiki Todo, Yuko Sakurai, and Makoto Yokoo. Fairness and false-name manipulations in randomized cake cutting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-15)*, pages 909–917, 2015.
- [Wada *et al.*, 2018] Yuho Wada, Tomohiro Ono, Taiki Todo, and Makoto Yokoo. Facility location with variable and dynamic populations. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-18)*, pages 336–344, 2018.
- [Yokoo *et al.*, 2004] Makoto Yokoo, Yuko Sakurai, and Shigeo Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.