Contests and Other Topics in Multi-Agent Systems

Doctoral Consortium

Abheek Ghosh
University of Oxford
abheek.ghosh@cs.ox.ac.uk

ABSTRACT

Contests are games where agents compete by making costly and irreversible investments to win valuable prizes. They model diverse scenarios ranging from competition among Bitcoin miners to crowdsourcing. My work has touched upon the following topics in contests theory: (i) design of contests to get a moderate output from many agents rather than a very high output from a few; (ii) design of contests to get higher output from an underrepresented group of agents; (iii) existence, computational complexity, and price of anarchy of equilibria in a model where agents participate in several simultaneous contests; (iv) convergence of best-response dynamics in contests. In addition to the above, my ongoing work focuses on topics in contest theory like learning dynamics in contests and analysis of contests where groups of agents (and not just individual agents) compete to get an outcome that affects all of them.

More broadly, I have also worked on the following topics: (i) improved, near-optimal algorithms for restless multi-armed bandits with applications to healthcare; (ii) analysis of coalition formation dynamics for deliberation.

KEYWORDS

Contest Theory; Mechanism Design; Learning Dynamics; Computational Complexity; Multi-Armed Bandits

ACM Reference Format:

Abheek Ghosh. 2023. Contests and Other Topics in Multi-Agent Systems: Doctoral Consortium. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

CONTESTS

Contests are games where agents invest effort and produce output to win valuable prizes, the investments are costly and irreversible, and the prizes are allocated based on the outputs. Contests model many practical problems, from crowdsourcing to labor markets. Due to their generality and wide use, several models have been studied in the literature [16].

Diversity in Contests

One of the primary motivations for organizing contests is to elicit effort from participants. For example, olympiads and hackathons encourage students to put effort and learn about specific subjects and technologies. In these contests, the designer wants to select a prize allocation mechanism to incentivize the agents to produce higher output, and the dominant paradigm is to incentivize the agents to produce a higher total output [13]. In such models, the designer values the contribution from each agent equally, i.e., the marginal output produced by any agent contributes the same marginal value to the designer's objective.

In many practical applications, the designer may want to elicit an adequate output from several agents instead of a very high output from a few agents. For example, in a crowdsourcing task, such as a survey, it may be more valuable to get many contributors to give adequate responses than to get a few people to submit perfect responses. In [4], we designed contests to optimize such objectives. These contests can get good participation from many agents rather than a very high output from a few, or they can target a specific segment of agents based on their output level, like incentivizing agents to produce moderate output and giving less importance to agents producing very high or very low output.

In other applications, it may be important to elicit higher participation from a protected group. Many competitions and hackathons are organized to elicit engagement from underrepresented groups, e.g., hackathons to get women interested in AI [12]. Contests are widely used for crowdsourcing, which is also an important source of training data for machine learning algorithms; in this context, eliciting input from disadvantaged groups is particularly important, as it helps the algorithms learn decision-making rules that reflect the opinions and preferences of such groups. Our work provides a better understanding of how to encourage contributors from such a target group [6]. Our intended follow-up work is to relax certain assumptions made in our models, e.g., study some general prize structures and non-linear cost functions.

Simultaneous Contests

In many applications, including the ones mentioned earlier, agents may have to choose among several contests, e.g., in crowdsourcing, an agent has several choices of projects. Simultaneous contests—where multiple contests run in parallel and the agents have to strategically allocate their limited resources across the contests—are natural extensions of the standard single-contest models.

In [7], we study a model where contests allocate prizes based on equal sharing: each contest awards its prize to all agents who satisfy some contest-specific criterion, and the value of this prize to a winner decreases as the number of winners increases. The agents produce outputs for a set of activities, and the winning criteria of the contests are based on these outputs. We show the existence, computational complexity, and price of anarchy of equilibria in these games. In particular, the complexity of computing a mixed Nash equilibrium in a version of these games is complete for the class PPAD∩PLS [1, 9].

Despite previous attempts, simultaneous contests are less well understood than their single-contest counterparts. We intend to make progress by assuming a large number of agents, where each

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

agent has minimal influence on the outcome. Simultaneous contests also lead to new mechanism design problems not present in a single contest. For example, consider the blockchain platform Cardano, which uses crowdsourcing to get quality reviews for proposals to make changes to the platform [2]. Cardano incentivizes reviewers by providing rewards for quality reviews. This leads to a game of simultaneous contests. In addition to maximizing the number and the quality of the reviews the proposals get, we want to have good coverage of reviews across the proposals.

Best-Response and Learning

Contests are studied using either complete or incomplete information models. In a complete information model, we assume that each agent exactly knows her own ability and the ability of every other agent, and the agents play rationally to reach a Nash equilibrium. On the other hand, in an incomplete information model, an agent does not exactly know the abilities of the other agents but has distributional knowledge; and based on this information, the agents reach a Bayes–Nash equilibrium. Both these models make strong assumptions about the information available to the agents, whether exactly knowing the abilities of the other agents or exactly knowing the distribution. These models also assume fully rational behavior. Such assumptions may not hold in practice.

Learning dynamics, e.g., best response and fictitious play, are alternative techniques for analyzing strategic games. These models make weaker assumptions regarding the cognitive ability of the agents and the information available to them. In our upcoming work [10], we study the convergence of best-response dynamics in Tullock [15] contests. For non-homogeneous agents, we show that best response dynamics may not converge, but for homogeneous agents, we show very fast convergence: the agents reach an ϵ-approximate equilibrium in $\Theta(\log\log(1/\epsilon))$ steps for two agents and in $\Theta(\log(1/\epsilon))$ steps for three or more agents.

We are working on a few others problems related to learning in contests. Fictitious play is known to converge for all-pay and Tullock contests with two agents, but for more agents, convergence has not been (dis)proven. We are also looking at learning in models where agents only see the outcome of the contest (allocation of the prizes) but not the actions of the agents.

Some Other Directions in Contest Theory

Fairness. Scholarships and funding are generally allocated by mechanisms that naturally induce competition among the applicants. In these scenarios, it is crucial to allocate the resources fairly and to understand the effect of a given allocation scheme on the behavior of the agents. Although several design objectives have been considered for contests (e.g., [3, 4]), fairness has largely remained unexplored. We plan to introduce well-motivated fairness objectives into standard models like Tullock [15] and rank-order allocation [13].

Repeated Contests. Contests are generally modeled as one-stage simultaneous games, and the recommended prize structure is to allocate the entire budget to the first prize (for natural models, e.g., [13]). The intuition is that a larger first prize incentivizes the stronger agents to put in more effort, which outweighs the opposite effect on the weaker agents due to smaller prizes for lower ranks. But, intuitively, there may be second-order effects of such design choices in the longer run. The weaker agents may get demotivated, which may change the distribution of the types of agents in future contests. This may decrease competition and lead to lower output. Our goal is to model such effects using repeated contests and study the design of contests that achieve long-term higher output.

Group Contests. Consider an election where it is costly for an agent to vote, and an agent gets a weight in the decision-making proportional to his costly investment. A similar situation occurs in proof-of-stake cryptocurrencies, where an agent gets voting rights in proportion to the amount of his stake that gets blocked, which leads to opportunity costs. An agent i would want to bear a higher cost and larger voting weight if his vote can affect the outcome, which depends upon how the agents who have a similar preference to i are doing as a group in comparison to other groups of agents. We intend to study these games by modeling them as contests.

OTHER TOPICS

Restless Multi-Armed Bandits (RMABs). An RMAB [17] is a generalization of a multi-armed bandit where multiple arms can be pulled simultaneously and the arms that are not pulled can also change their states. The planning problem for RMABs is PSPACE-hard [14]. The most popularly used algorithm, called the Whittle index policy [17], is a heuristic based on a Lagrangian relaxation that has asymptomatic optimality guarantees under certain conditions. These conditions required for the asymptotic optimality of Whittle are hard to verify for most practical problems, but practitioners still use it hoping that the bad instances do not occur naturally. In [11], we show that this is not true. We provide simple and natural examples where Whittle fails. Second, we study a mean-field based approximation algorithm and provide non-asymptotic high-probability approximation guarantees (which implies asymptomatic optimality), and these results hold unconditionally, unlike Whittle. For two applications, maternity healthcare intervention and tuberculosis monitoring, we show that our algorithm performs better than alternatives in simulations based on real-life data.

Deliberative Coalition Formation. We study a dynamics for deliberative coalition formation where agents dynamically form coalitions around proposals that they prefer over the status quo [8]. The deliberation works through compromise moves, where agents from several (current) coalitions come together to form a larger coalition to support a (perhaps new) proposal, possibly leaving some of the agents from their previous coalitions, who do not support the new proposal, behind. A deliberation succeeds if it terminates by identifying a proposal with the largest possible support. It has been previously shown that compromise moves, involving a small number of coalitions in each move, may lead to successful deliberation in polynomially many steps [8]. Our work [5] resolved several open problems; we showed that for a successful deliberation, the following may happen: (1) agents from a very large number (exponential in relevant parameters) of coalitions have to simultaneously join forces to form a larger coalition; (2) the sequence of compromise moves may be exponentially long; (3) even for a sequence of compromise moves that are polynomial in length, it may be NP-hard to make some of the moves in the sequence.

ACKNOWLEDGMENTS

I thank my advisors, Edith Elkind and Paul Goldberg, for their guidance and support. I also thank Milind Tambe, Dheeraj Nagaraj, and Manish Jain for an internship opportunity at Google and for introducing me to problems related to multi-armed bandits. My Ph.D. is funded by the Clarendon Fund Scholarship and the SKP Scholarship.

REFERENCES

[1] Yakov Babichenko and Aviad Rubinstein. 2021. Settling the Complexity of Nash Equilibrium in Congestion games. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 1426–1437.

[2] Georgios Birmpas, Lyudmila Kovalchuk, Philip Lazos, and Roman Oliynykov. 2022. Parallel Contests for Crowdsourcing Reviews: Existence and Quality of Equilibria. ArXiv preprint arXiv:2202.04064.

[3] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. 2019. Optimal crowdsourcing contests. *Games and Economic Behavior (GEB)* 113 (2019), 80–96.

[4] Edith Elkind, Abheek Ghosh, and Paul Goldberg. 2021. Contest Design with Threshold Objectives. In *Web and Internet Economics (WINE)*. 554. arXiv:2109.03179.

[5] Edith Elkind, Abheek Ghosh, and Paul Goldberg. 2022. Complexity of Deliberative Coalition Formation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 4975–4982. arXiv:2202.12594.

[6] Edith Elkind, Abheek Ghosh, and Paul Goldberg. 2022. Contests to Incentivize a Target Group. In *International Joint Conference in Artificial Intelligence (IJCAI)*. 279–285. arXiv:2204.14051.

[7] Edith Elkind, Abheek Ghosh, and Paul Goldberg. 2022. Simultaneous Contests with Equal Sharing Allocation of Prizes: Computational Complexity and Price of Anarchy. In *Symposium in Algorithmic Game Theory (SAGT)*. 133–150. arXiv:2207.08151.

[8] Edith Elkind, Davide Grossi, Ehud Shapiro, and Nimrod Talmon. 2021. United for Change: Deliberative Coalition Formation to Change the Status Quo. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5339–5346.

[9] John Fearnley, Paul Goldberg, Alexandros Hollender, and Rahul Savani. 2022. The complexity of gradient descent: CLS= PPAD∩PLS. *J. ACM* 70, 1 (2022), 1–74.

[10] Abheek Ghosh and Paul Goldberg. 2023. Best-Response Dynamics in Tullock Contests. In submission.

[11] Abheek Ghosh, Dheeraj Nagaraj, Manish Jain, and Milind Tambe. 2023. Indexability is Not Enough for Whittle: Improved, Near-Optimal Algorithms for Restless Bandits. In *Proceedings of the 2023 International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*. arXiv:2211.00112.

[12] Alice LaPlante. 2021. Microsoft News. Bridging the STEM divide: AI hackathon helps young women excel in computer science. https://bit.ly/microsoft-hackathon. Accessed: 2022-09-19.

[13] Benny Moldovanu and Aner Sela. 2001. The optimal allocation of prizes in contests. *American Economic Review (AER)* 91, 3 (2001), 542–558.

[14] Christos H Papadimitriou and John N Tsitsiklis. 1999. The Complexity of Optimal Queuing Network Control. *Mathematics of Operations Research* 24, 2 (1999), 293–305.

[15] Gordon Tullock. 1980. Efficient Rent-Seeking. In *Toward a Theory of the Rent Seeking Society*, James Buchanan, Robert Tollison, and Gordon Tullock (Eds.). Texas A & M University, College Station, Texas, USA, 131–146.

[16] Milan Vojnović. 2016. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press, Cambridge, UK.

[17] Peter Whittle. 1988. Restless bandits: Activity allocation in a changing world. *Journal of applied probability* 25, A (1988), 287–298.