Selective Information Disclosure in Contests

Extended Abstract

Priel Levy, David Sarne and Yonatan Aumann
Bar Ilan University, Israel
priel.levy@live.biu.ac.il,sarned@cs.biu.ac.il,aumann@cs.biu.ac.il

ACM Reference Format:
Priel Levy, David Sarne and Yonatan Aumann. 2019. Selective Information Disclosure in Contests. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

Contests are important mechanisms to elicit work (/effort/ideas) from crowds. While contests have been used throughout history (e.g. the British government's 1714 Longitude Prize), they have gained popularity in the current Internet era, and, in particular, in the context of crowdsourcing [2, 7, 14, 41, 59, 60]. Well known examples include the Netflix prize (netflixprize.com), Darpa challenges [3, 57] and the Hult prize (hultprize.org), as well as various public platforms that allow requesters to solicit contributions through contests with monetary prizes, such as taskCN (www.taskcn.com), TopCoder (www.topcoder.com) and Kaggle (www.kaggle.com). As such, the study and analysis of contests have become prominent in mechanism design and multi-agent systems literature [6, 14, 23, 36–39, 43, 54]. These include both the analysis and determination of optimal strategies - for the contestants, and methods for the design of effective contests - for the contest's organizer. In this work we concentrate on the latter issue - that of contest design.

Effective contest design has been studied extensively in literature, but most work has focused on how to design the payoffs structure [2, 7, 14, 23, 43, 45]. That is, how many and what prizes should be awarded, and to which contestants. When designing a contest, however, the organizer has freedom to structure all aspects of the contest protocol, not only the payoffs, and this entire structure determines the contestants behavior. In particular, the *information* available to the contestants during the contest has a dramatic impact on their behavior, hence by controlling this information the contest designer can promote its goals.

There are various types of information that can potentially affect prospective contestants' decisions in a contest: information about the protocol/mechanism, the way it is executed and winner determination, information about own competence which is often uncertain, information about the other contestants – their number, costs of participation (or cost of exerting effort), competence, etc., and information about the actions of other contestants and their resulting performance so far. Potentially, the contest designer may try and control the disclosure of any of the above to its benefit (though some are at times not under its control).

Recent literature has acknowledged the importance of information design in contests, studying various issues arising from asymmetric information and information disclosure (see following section). Yet, the information considered in the models studied relates to the inherent characteristics of contestants rather than their actions. Furthermore, and perhaps more importantly, in the models used all information is disclosed at the beginning of the contest (or prior to the contest), thus lacking any temporal aspect.

In this research, we study information design in contests, focusing on *temporal* information disclosure, wherein information on the actions of contestants (participation and the amount of effort expended) is disclosed *during* the course of the contest. Importantly, the temporal disclosure of information turns the contest mechanism from a pure parallel game to a semi-sequential one where contestants make their decisions in real-time while new information unfolds, resulting in a plethora of new hybrid designs.

2 RELATED WORK

Much theoretical work has been devoted to the design of an optimal contest that best serves the organizer's objective function, typically by assuming a specific structure and studying its equilibrium under different assumptions. Common structures are one-stage, where contestants compete simultaneously [12, 19, 23, 52] or multiple rounds consisting of series of contests (most known as a tournament) [9, 10, 26, 27, 46, 51]. Common assumptions regarding the underlying contest model include: the number of contestants [5, 22], their level of heterogeneity [56], asymmetry regarding their information [8, 40, 55], their tendency towards risk [2], rewards allocation [2, 6, 11, 21, 45], and the organizer's objective function (e.g., maximizing overall performance, best performance) [35, 46].

The study of information design in contests is very recent and encompasses various aspects of information providing. For example, Einy et al. [16] study how changes in the information available to the players of a symmetric common-value Tullock contest with incomplete information (captured by players' uncertainty about their common value and common cost) affect their equilibrium payoffs and their incentives to exert effort. Morath and Munster [47] study incentives for information acquisition ahead of a contest, focusing on the effect of whether the decision to acquire information is observable or not. Denter et al. [13] studied information policies of competitors in a contest who can decide and commit to acquire relevant information about their rivals or disclose their own private information to their rivals, before the contest. Kovenock et al. [32] study the incentives to share private information (e.g., signals related to the value of winning in the contest) ahead of the contest. Fu et al. [18] study the use of players' costly confidence-of-winning signaling (which may disclose a player's private information about his strength) prior to a contest. Ponce [50] analyzes the extent to which an organizer can influence players' behavior by revealing information about the players capabilities (which are a priori uncertain) in

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

binary action contests. Kramm [33] studies effective information structures for transmitting to players before the contest in a model where an individual's success in the contest will depend on the mixture of effort investments in different tasks while the weights given to the different tasks are not fully clear. Dubey [15] studies the effect of the value of the prize awarded over the preference of having the players know the ability of their rivals, assuming each player knows her own ability. He finds that to inspire performance, it is often better for the contest organizer neither to reveal all nor to conceal all, but to follow a middle path of partial revelation. Common to all the above works is that the information they consider relates to the types of contestants, i.e., their competence, cost of participation, etc., rather than their actions. Furthermore, and perhaps more importantly, all information is disclosed at the beginning of the contest (or prior to the contest), thus lacking any temporal aspect. As such, decisions of prospective contestants are made in parallel, as in most contest literature [12, 24, 42, 45, 46]. Models where contestants' decisions are made sequentially, taking into account information related to the decisions (and consequently the results) of others were mostly studied with two-period models [1, 34, 44, 48], and rarely with more than two periods [20, 25, 39]. Moreover, in models of the latter type the process is fully-sequential, i.e., there is only one contestant at a time (according to some pre-defined ordering) deciding on the extent of its participation in the contest and contestants become aware of the performance of all those who engaged before them (e.g., in rhythmic and artistic gymnastics).

Work on temporal information design in contests is very limited, typically by the number of contestants used and the assumptions made. For example, Epstein and Mealem [17] study the equilibrium in a two-stage two-players contest model in which the informed player declares its type (or does not declare) in the first stage and in the second stage the two players play according to the information available to them. Gurtler et al. [28] studies sabotage activity in a tournament (i.e., whenever a player invests in reducing the effectiveness of a rival's effort), demonstrating that by concealing intermediate information about players' performances the incidence of sabotage is mitigated. The most relevant work in the context of this paper is the work of Hinnosaar [30] which has been carried out in parallel to ours.

More broadly, we note that selective information disclosure has been extensively studied in recent years in the field of psychology and behavioral economics [31, 58] and in multi-agent literature [49, 53]. For example, in advice-giving settings, where the advisor's interests may conflict with the interests of the people who receive the advice, selective information disclosure can improve the advisor's expected benefit [4]. Or, in adversarial setting, where comparison shopping agents can influence human users not to query additional agents by using selective information disclosure [29]. Nevertheless, despite the rich literature on selective information disclosure, no previous work has addressed hybrid designs of information disclosure in contests of our type.

3 THE GENERAL CONTEST MODEL

The basic contest setting considers a contest organizer and a set $A = \{A_1, ..., A_k\}$ of $k > 1$ potential contestants (denoted "agents" onwards). The agents are heterogeneous, such that different agents may differ in their types, where type may define various parameters, e.g., competence, cost of engagement in the contest and knowledge. The agents need to decide if and to what extent to engage in the contest (e.g., how much effort to exert). All agents and the organizer are fully-rational and self-interested. Agents' extent of participation in the contest depends thus on the compensation offered by the organizer. To elicit participation and effort the organizer offers a prize to be awarded to the highest-ranked agent, where the ranking is a function of the different efforts/performance exerted by the participating agents (prize allotment may also involve a stochastic element). We also assume, like many prior works [7, 14, 39] that the organizer and the agents are familiar with the prize offered and the agents' type distributions. The goal of the organizer running the contest is to maximize some objective function that takes as an input the efforts/performance exerted by the agents and the prizes awarded. The goal of each agent is to maximize its own expected profit, defined as the value of the prize it receives minus the cost incurred while engaging in the contest.

The above underlying contest model encompasses almost any contest model found in prior work and can be augmented to support temporal information design. The contest organizer may have access to information not available to the agents. This information may either be known to it in advance - e.g., a priori knowledge of the agents' types - or become available during the course of the contest - e.g., the agents' actual performance levels. To further its goals, the organizer may decide to temporally disclose some of this latter information to the agents. In principle, such selective information disclosure can take many forms. For concreteness, in this paper we focus on a case where the contest's design and rules, as well as the agents' a priori types are common knowledge and the only information that becomes available to the organizer, which may not be known to the agents, is the actual actions of the agents during the course of the contest. The organizer has the freedom to choose if and when to disclose this information, and commits in advance, to its information disclosure scheme. Furthermore, we assume the information disclosed by the organizer becomes publicly available and each information disclosure act involves the disclosure of the actions and performance of some set of agents S. Such disclosure will take place immediately after all agents in S completed their participation in the contest. Finally, we assume the organizer is fully reliable, in the sense that it reveals only the truth.

4 DISCUSSION AND CONCLUSIONS

The research reported in this paper touches on the essence of running contests. The idea in using a contest-based mechanism rather than simply paying people to take part in the process is that the strategic behavior of contestants reveal additional information about themselves. A proper temporal information design has the potential to further leverage this advantage. We therefore expect this research to have significant impact on the study and implementation of contests as a mechanism for eliciting effort, especially with the recent increased interest in crowdsourcing contests.

ACKNOWLEDGMENTS

This research was partially supported by the ISRAEL SCIENCE FOUNDATION grant No. 1162/17.

REFERENCES

[1] J. Atsu Amegashie. 2000. Some results on rent-seeking contests with shortlisting. *Public Choice* 105, 3-4 (2000), 245–253.

[2] Nikoly Archak and Arun Sundararajan. 2009. Optimal design of crowdsourcing contests. *ICIS 2009 proceedings* (2009), 200.

[3] Christopher G. Atkeson et al. 2018. *What Happened at the DARPA Robotics Challenge Finals.* Springer International Publishing, Cham, 667–684.

[4] Amos Azaria et al. 2012. Giving advice to people in path selection problems. In *Proc. of AAMAS*. 459–466.

[5] Kevin J. Boudreau, Nicola Lacetera, and Karim R. Lakhani. 2011. Incentives and Problem Uncertainty in Innovation Contests: An Empirical Analysis. *Management Science* 57, 5 (2011), 843–863.

[6] Ruggiero Cavallo and Shaili Jain. 2013. Winner-Take-All Crowdsourcing Contests with Stochastic Production. In *Proc. of HCOMP*. 34–41.

[7] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. 2012. Optimal crowdsourcing contests. In *Proc. of SODA*. 856–868.

[8] Yeon-Koo Che and Ian Gale. 2003. Optimal design of research contests. *American Economic Review* 93, 3 (2003), 646–671.

[9] Derek J Clark and Christian Riis. 1996. A multi-winner nested rent-seeking contest. *Public Choice* 87, 1-2 (1996), 177–184.

[10] Derek J. Clark and Christian Riis. 1998. Influence and the discretionary allocation of several prizes. *European Journal of Political Economy* 14, 4 (1998), 605–625.

[11] Chen Cohen, Todd Kaplan, and Aner Sela. 2008. Optimal rewards in contests. *RAND Journal of Economics* 39, 2 (2008), 434–451.

[12] Ani Dasgupta and Kofi O. Nti. 1998. Designing an optimal contest. *European Journal of Political Economy* 14, 4 (1998), 587 – 603.

[13] Philipp Denter, John Morgan, and Dana Sisak. 2014. "Where Ignorance is Bliss, 'tis Folly to be Wise": Transparency in Contests. (2014).

[14] Dominic DiPalantino and Milan Vojnovic. 2009. Crowdsourcing and all-pay auctions. In *Proc. of ACM-EC*. 119–128.

[15] Pradeep Dubey. 2013. The role of information in contests. *Economics Letters* 120, 2 (2013), 160–163.

[16] Ezra Einy, Diego Moreno, and Benyamin Shitovitz. 2017. The value of public information in common-value Tullock contests. *Economic Theory* 63, 4 (2017), 925–942.

[17] Gil S Epstein and Yosef Mealem. 2013. Who gains from information asymmetry? *Theory and decision* 75, 3 (2013), 305–337.

[18] Qiang Fu, Oliver GÃijrtler, and Johannes MÃijnster. 2013. Communication and commitment in contests. *Journal of Economic Behavior & Organization* 95 (2013), 1 – 19.

[19] Qiang Fu and Jingfeng Lu. 2010. Contest design and optimal endogenous entry. *Economic Inquiry* 48, 1 (2010), 80–88.

[20] Qiang Fu and Jingfeng Lu. 2012. The optimal multi-stage contest. *Journal of Economic Theory* 51, 2 (2012), 351–382.

[21] Francis Galton. 1902. The Most Suitable Proportion Between the Value of First and Second Prizes. *Biometrika* 1, 4 (1902), 385–399.

[22] Stephen M. Garcia and Avishalom Tor. 2009. The N-Effect. *Psychological Science* 20, 7 (2009), 871–877.

[23] Arpita Ghosh and Robert Kleinberg. 2016. Optimal Contest Design for Simple Agents. *ACM Transactions on Economic and Computation* 4, 4, Article 22 (2016), 22:1-22:41 pages.

[24] Amihai Glazer and Refael Hassin. 1988. Optimal Contests. *Economic Inquiry* 26, 1 (1988), 133–143.

[25] Amihai Glazer and Refael Hassin. 2000. Sequential rent seeking. *Public Choice* 102, 3-4 (2000), 219–228.

[26] Mark Gradstein. 1998. Optimal contest design: volume and timing of rent seeking in contests. *European Journal of Political Economy* 14, 4 (1998), 575–585.

[27] Mark Gradstein and Kai A Konrad. 1999. Orchestrating Rent Seeking Contests. *Economic Journal* 109, 458 (1999), 536–545.

[28] Oliver Gurtler, Johannes Munster, and Petra Nieken. 2013. Information Policy in Tournaments with Sabotage. *The Scandinavian Journal of Economics* 115, 3 (2013), 932–966.

[29] Chen Hajaj, Noam Hazon, and David Sarne. 2014. Ordering Effects and Belief Adjustment in the Use of Comparison Shopping Agents. In *Proc. of AAAI*. 930–936.

[30] Toomas Hinnosaar. 2018. Optimal sequential contests. *American Economic Review* (2018). to appear.

[31] Sheena Iyengar. 2010. *The art of choosing.* Twelve.

[32] Dan Kovenock, Florian Morath, and Johannes MÃijnster. 2015. Information Sharing in Contests. *Journal of Economics & Management Strategy* 24, 3 (2015), 570–596.

[33] Michael Kramm. 2018. Information Design in Multi-Task Contests - Whom to Inform When the Importance of Tasks Is Uncertain. (2018). Working paper, Technical University Dortmund.

[34] Vijay Krishna and John Morgan. 1998. The winner-take-all principle in small tournaments. *Advances in applied microeconomics* 7 (1998), 61–74.

[35] Omer Lev, Maria Polukarov, Yoram Bachrach, and Jeffrey S. Rosenschein. 2013. Mergers and collusion in all-pay auctions and crowdsourcing contests. In *Proc. of AAMAS*. 675–682.

[36] Priel Levy. 2018. Optimal Contest Design for Multi-Agent Systems. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, 376–377.

[37] Priel Levy and David Sarne. 2018. Understanding Over Participation in Simple Contests. In *Proc. of AAAI*. 1571–1578.

[38] Priel Levy, David Sarne, and Yonatan Aumann. 2018. Tractable (Simple) Contests. In *Proc. of IJCAI*. 361–367.

[39] Priel Levy, David Sarne, and Igor Rochlin. 2017. Contest Design with Uncertain Performance and Costly Participation. In *Proc. of IJCAI*. 302–309.

[40] Bin Liu, Jingfeng Lu, Ruqu Wang, and Jun Zhang. 2013. Prize and Punishment: Optimal Contest Design with Incomplete Information. *National University of Signapore Typescript* (2013).

[41] Tracy Xiao Liu, Jiang Yang, Lada A. Adamic, and Yan Chen. 2014. Crowdsourcing with All-Pay Auctions: A Field Experiment on Taskcn. *Management Science* 60, 8 (2014), 2020–2037.

[42] Xuyuan Liu and Jingfeng Lu. 2014. The effort-maximizing contest with heterogeneous prizes. *Economics Letters* 125, 3 (2014), 422–425.

[43] Tie Luo, Salil S. Kanhere, Hwee-Pink Tan, Fan Wu, and Hongyi Wu. 2015. Crowdsourcing with Tullock contests: A new perspective. In *Proc. of INFOCOM*. 2515–2523.

[44] Alexander Matros. 2006. *Elimination tournaments where players have fixed resources.* Technical Report. University of Pittsburgh, Department of Economics.

[45] Benny Moldovanu and Aner Sela. 2001. The Optimal Allocation of Prizes in Contests. *American Economic Review* 91, 3 (2001), 542–558.

[46] Benny Moldovanu and Aner Sela. 2006. Contest architecture. *Journal of Economic Theory* 126, 1 (2006), 70–96.

[47] Florian Morath and Johannes Münster. 2013. Information acquisition in conflicts. *Economic Theory* 54, 1 (01 Sep 2013), 99–129.

[48] John Morgan. 2003. Sequential contests. *Public Choice* 116, 1-2 (2003), 1–18.

[49] Noam Peled, Sarit Kraus, et al. 2015. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems* 29, 1 (2015), 73–97.

[50] Alejandro Melo Ponce. 2018. Information Design in Contests. (2018). Working paper, Department of Economics Stony Brook University.

[51] Sherwin Rosen. 1986. Prizes and Incentives in Elimination Tournaments. *The American Economic Review* 76, 4 (1986), 701–715.

[52] Marco Runkel. 2006. Optimal contest design, closeness and the contest success function. *Public Choice* 129, 1 (01 Oct 2006), 217–231.

[53] David Sarne, Avshalom Elmalech, Barbara J. Grosz, and Moti Geva. 2011. Less is more: restructuring decisions to improve agent search. In *Proc. of AAMAS*. 431–438.

[54] David Sarne and Michael Lepioshkin. 2017. Effective Prize Structure for Simple Crowdsourcing Contests with Participation Costs. In *Proc. of HCOMP*. 167–176.

[55] Nirvikar Singh and Donald Wittman. 1988. Economic contests with incomplete information and optimal contest design. *Management Science* 34, 4 (1988), 528–540.

[56] Rudi Stracke. 2013. Contest design and heterogeneity. *Economics Letters* 121, 1 (2013), 4–7.

[57] V. S. Subrahmanian et al. 2016. The DARPA Twitter Bot Challenge. *IEEE Computer* 49, 6 (2016), 38–46.

[58] Richard H. Thaler and Cass R. Sunstein. 2008. Nudge: Improving decisions about health, wealth, and happiness. *Boston Yale University Press* (2008).

[59] Milan Vojnovic. 2016. *Contest Theory: Incentive Mechanisms and Ranking Methods.* Cambridge University Press. i–iv pages.

[60] Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. 2012. Crowdsourcing to Smartphones: Incentive Mechanism Design for Mobile Phone Sensing. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*. 173–184.