Optimal Contest Design for Multi-Agent Systems - Extended Abstract

Priel Levy
Department of Computer Science, Bar Ilan University, Israel
priel.levy@live.biu.ac.il

Abstract

Contests have become a highly popular crowdsourcing mechanism aiming to solicit effort of the crowd in solving well defined problems, and as such are extensively studied within the framework of contest design. In this summary, I give a brief background for my Ph.D thesis research and explain the work I have accomplished so far, under the supervision of Prof. David Sarne.

Introduction

Contests have been used since the dawn of man as a mechanism for inducing individual efforts. In recent years, contests are used not only as a means for determining the best contestant, but also for generating value. For example DARPA is offering Grand Challenges to promote the development of cutting-edge technologies, firms run contests to come up with new products (such as the LEGO Ideas contests (Schlagwein and Bjorn-Andersen 2014)) and not-for-profit organizations are organizing contests for transformative solutions that benefit mankind (as with the Hult Prize).[1]

The ubiquitousness of contests, has led to a huge body of literature on contest design studying how utility-maximizing contestants' willingness to make high-quality submissions in a contest is affected by parameters such as prize structure, contest structure, contestants' rationality (people or fully rational agents), contestants' heterogeneity and number of contestants (DiPalantino and Vojnovic 2009; Cavallo and Jain 2013; Liu et al. 2014). Alongside the many theoretical contributions made to the field, there is much experimental research aiming at empirically investigating individual behavior in different crowdsourcing contests (Dechenaux, Kovenock, and Sheremeta 2014). One common finding in this literature is that people tend to exert more effort in contests compared to the theoretical equilibrium-based predictions, leading to better expected contributions overall (Sheremeta 2013). The reasons proposed for this phenomenon are diverse, the most common reason is that subjects derive a non-monetary utility from winning in addition to monetary incentives (Sheremeta 2010). Other reasons suggested that subjects are prone to mistakes and systematic biases (Chowdhury, Sheremeta, and Turocy 2014).

[1] (www.hultprize.org)

My thesis research aims to contribute to contest theories through the theoretical analysis of and comprehensive experimentation with a type of contest that has gained momentum in research only very recently (Ghosh and Kleinberg 2016). In this type of contest potential contestants only strategize on whether to participate or not in the contest, as they have no control over the quality of their submission if participating. This model is a special case of strategic effort models, where instead of adjusting their efforts, contestants are restricted to a binary choice between the maximum possible effort (i.e., participating) and zero effort (i.e., not participating). The goal of the principal is to come up with a contest design that will maximize the expected best quality (performance-wise) among the elicited submissions.

Indeed, there is a wide spectrum of contests in real-life, where participants cannot strategize over the quality of their submissions and are limited to choosing whether or not to participate in the contest. This holds whenever contestants are being evaluated subjectively by anonymous judges, based on the taste of the public (in case of a voting open to the public), based on criteria that are not fully disclosed to them, or simply based on luck or uncertain environment parameters (e.g., weather). For example, at the time a department chair solicits submissions from faculty members for excellence in research prize, prospective candidates cannot influence anymore the quality of the results they submit (as the research has already been carried out and published) and the way their results will be evaluated by the department chair compared to others' is somehow uncertain. Here, the contestants' strategy is limited to participating or not participating in the contest, which becomes non-trivial whenever participation incurs some cost (e.g., sending a CV, emotional suffer and reputational loss in case of not winning).

Research Hypothesis and Goals

The main hypothesis of this research is that different variants of the above archetypal contest model may be associated with different preference of the type of contest to be used, even when the difference between these variances are minor. Hence, the goal of this research is to provide an equilibrium analysis for a wide spectrum of such contest model variants and to reason about the preference between them (from the principal's and contestants' point of view). These should lead to optimal contest design for fully rational agents.

In addition, we hypothesize that different optimal contest designs will require adjustments whenever contestants are human. This coincides with various literature on people's bounded rationality as well as value theories. Therefore, we aim to adjust the models to be developed, in a way that takes into consideration user modeling and bounded-rationality theories, to enable contests that perform better with people.

General Model

We consider a setting of a contest organizer and a set $A = \{A_1, ..., A_k\}$ of $k > 1$ potential contestants. Each contestant A_i can either participate in the contest, or opt to avoid participating in the contest. Most of the works in the field of contest design deal with effort-based contests, where the effort expended by contestants in case of participating determines their performance in the contest and consequently their probability of wining. There are also works, although not many where each contestant only strategizes about participation and its performance in the contest is determined according to some known probability distribution function. Yet, we do not limit ourselves to a specific model.

The goal of the organizer is to maximize the expected maximum performance obtained by contestants in a contest it runs. In order to encourage participation in the contest the organizer offers a set of prizes $M_1, ..., M_n > 0, (0 < n \leq k)$ to the contestants ranked firsts (performance-wise) in the contest. The goal of each contestant is to maximize its own expected profit, defined as the expected prize awarded to it minus the cost incurred if participating in the contest.

Preliminary Results

In the past year, we have focused on a model where contestants only strategize about participation and the performance is determined according to some known probability distribution function. Our preliminary results have focused in the comparison of a parallel and sequential contest models whenever performance is beyond the influence of the contestant and yet a priori uncertain at the time of making the participation decision. In the parallel contest, each contestant's participation decision takes part in parallel to the others', where in the sequential contest each contestant in its turn (according to some pre-defined order) gets to see the results of its predecessors (whether participated, and if so also their performance) and then decides whether to participate in the contest. First, we have provided a comparative game-theoretic based solution to these two variants of the model, enabling a characterization of the equilibrium strategies in each. Special emphasis was placed on the case where the contestants are a priori homogeneous which is often the case in contests where ranking is mostly influenced by some probabilistic factors (e.g., luck) or whenever contestants are evaluated subjectively by a referee whose taste cannot be a priori predicted. Here, several (somehow counter-intuitive) properties of the equilibrium were proved, in particular for the sequential contest, leading to a comprehensive characterization of the principal's preference between the two. These results have were published in IJCAI-2017 (Levy, Sarne, and Rochlin 2017). Second, we have conducted experiments with people (using Amazon Mechanical Turk) for studying whether the phenomenon of over participation in simple contests is due to people competitiveness or perhaps it can be attributed to other factors that hold also in non-competitive settings. Our unique experimental design, enabled an important insight that is absent in prior work - it is not the competitive nature of the interaction that accounts for the excessive effort exerted, as speculated in prior work, but rather some other factor that holds also in non-competitive similar decision settings, most probably people's tendency towards risk. This insight is supported by direct comparison of decisions made in contests and in equivalent decision situations from which the competitive aspect is absent, and is of great importance to mechanism designers. These results will be presented in AAAI-2018 (Levy and Sarne 2018).

Research Method and Plan

In the upcoming future I would like to advance my understanding of optimal contest design, either theoretically (i.e., when contestants are fully rational) principally using concepts of search theory and game theory or experimentally when contestants are human (relying on the extensive literature on people's behaviors found in experimental economics and psychology). I plan to vary the different model assumptions, generating variants of the model that can be applicable to different real-life problems, both for effort-based contests, and simple contests (i.e., the one detailed in this summary).

References

Cavallo, R., and Jain, S. 2013. Winner-take-all crowdsourcing contests with stochastic production. In *Proc. of HCOMP*, 34–41.

Chowdhury, S. M.; Sheremeta, R. M.; and Turocy, T. L. 2014. Overbidding and overspreading in rent-seeking experiments: Cost structure and prize allocation rules. *Games and Economic Behavior* 87:224 – 238.

Dechenaux, E.; Kovenock, D.; and Sheremeta, R. 2014. A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics* 1–61.

DiPalantino, D., and Vojnovic, M. 2009. Crowdsourcing and all-pay auctions. In *Proc. of ACM-EC*, 119–128.

Ghosh, A., and Kleinberg, R. 2016. Optimal contest design for simple agents. *ACM Transactions on Economic and Computation* 4(4):22:1–22:41.

Levy, P., and Sarne, D. 2018. Understanding over participation in simple contests. In *Proc. of AAAI*. (to appear).

Levy, P.; Sarne, D.; and Rochlin, I. 2017. Contest design with uncertain performance and costly participation. In *Proc. of IJCAI*, 302–309.

Liu, T. X.; Yang, J.; Adamic, L. A.; and Chen, Y. 2014. Crowdsourcing with all-pay auctions: A field experiment on taskcn. *Management Science* 60(8):2020–2037.

Schlagwein, D., and Bjorn-Andersen, N. 2014. Organizational learning with crowdsourcing: The revelatory case of lego. *J. of the Association of Information Systems* 15(11):754–778.

Sheremeta, R. M. 2010. Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior* 68(2):731 – 747.

Sheremeta, R. 2013. Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys* 27(3):491–514.