

A SIMULATION ANALYSIS OF LARGE CONTESTS WITH THRESHOLDING AGENTS

Wen Shen
Rohan Achar
Cristina V. Lopes

Department of Informatics
University of California, Irvine
Irvine, CA 92697, USA

ABSTRACT

Running contests has been an effective way to solicit efforts from a large pool of participants. Existing research mostly focuses on small contests that typically consist of two or several perfectly rational agents. In practice, however, agents are often founded in complex environments that involve large numbers of players, and they usually use thresholding policies to make decisions. Despite the fact, there is a surprising lack of understanding of how contest factors influence their outcomes. Here, we present the first simulation analysis on how parameters of the contest success function, the population dynamics, and the agents' cutoff policies influence the outcomes of the contests with thresholding agents that are non-cooperative. Experimental results demonstrate that stakeholders can design (approximately) optimal contests to satisfy both their interests and the agents' by choosing a relatively low bias factor. Our work brings new insights into how to design proper competitions to coordinate thresholding agents.

1 INTRODUCTION

In many real-world situations, stakeholders often hold contests to solicit efforts from a large pool of participants by distributing rewards to the winners. Typical examples include sports (Fort and Winfree 2009; Malueg and Yates 2010), advertising competitions (Gordon and Hartmann 2016), crowdsourcing (Shen, Feng, and Lopes 2019), and diffusion of innovations (Baye and Hoppe 2003). These activities are all adversarial settings where the outcomes are typically uncertain, but they do depend on the efforts committed by a player and those by her adversaries. In particular, the outcomes are determined by a contest success function that takes both the player's efforts and her adversaries' (Skaperdas 1996; Corchón and Dahm 2010). Taken all players' efforts as the input, the contest success function quantifies each contest player's probability of winning the contest (Skaperdas 1996). To design efficient contests, it is crucial to understand how the design of contest success functions influences the contest outcomes (Jia, Skaperdas, and Vaidya 2013; Clark and Riis 1998; Blavatsky 2010).

Previous literature on contest success functions has focused mainly on small contests that typically have two or a few perfectly rational agents due to tractability concerns (Jia, Skaperdas, and Vaidya 2013; Münster 2009; Runkel 2006). Most of the research restricts to theoretical analysis with the solution concepts of Nash Equilibrium and its variants (Corchón and Dahm 2010; Jia, Skaperdas, and Vaidya 2013; Runkel 2006). This line of research has substantially improved our knowledge of how to design optimal contests when agents have perfect information and perfect rationality.

However, studies on human behavior show that humans usually use thresholding policies to make decisions in real-world environments such as sports (Dimant and Deutscher 2015), crowdsourcing contests (Easley and Ghosh 2015), online shopping (Ohannessian, Roozbehani, Materassi, and Dahleh 2014; Shen, Crandall, Yan, and Lopes 2018; Shen 2019), smart grids (Almahmoud, Crandall, Elbassioni, Nguyen, and Roozbehani 2018) and social networks (Shen, Feng, and Lopes 2019). On the one hand, agents use

thresholding policies because they often have informational uncertainties, cognitive limitations or computational limits (Simon 1997). On the other hand, under certain conditions, thresholding policies are optimal strategies and thus represent rational behavior (O'hannessian, Roozbehani, Materassi, and Dahleh 2014). That is, thresholding agents are rational agents when they have optimal (perfect) thresholds. In this sense, the thresholding behavioral model is more generalized than the rational behavioral model because it extends the rational behavioral model to the scenarios when agents do not necessarily possess perfect information or reasoning abilities. Therefore, it is essential and demanding to understand how parameters of success functions and agents' behavior influence the performance of the contests with thresholding agents.

In this paper, we present the first simulation analysis to quantify how parameters of contest success functions, the population dynamics, and agents' thresholding behavior influence the contest outcomes. In doing so, we first developed a simulation platform based on real-world contest data. We then conducted four groups of experiments to study how the bias factor, the number of total rewards, the population dynamics of the agents, and the coefficient affected the performance of the contests. Experimental results demonstrate that there are "win-win" situations for both the stakeholders and the agents in contests with thresholding agents. Specifically, it is feasible for the stakeholders to design (approximately) optimal contests that can achieve good results for both parties by choosing a relatively low bias factor. Our results complement the state-of-the-art by offering a new perspective on how to design contests to coordinate agents with thresholding behaviors. In particular, our study provides a systematic approach for stakeholders to simulate systems that consists of thresholding agents as well as methods to evaluate the performance of the systems brought by different incentive mechanisms.

2 CONTESTS WITH THRESHOLDING AGENTS

This section first describes the decision models that players use to determine the number of efforts invested in the contests. It then introduces the contest success functions that generate each player's probability of success.

2.1 Agents' Decision Models

We consider a contest in which agents follow thresholding policies to decide how much efforts they should invest in the contests. Let I be the set of agents in a contest, an agent $i \in I$ will exert positive efforts $e_i \in R_{\geq 0}$ only if her (expected) utility $U(e_i) \in R$ is greater than or equal to zero. Here, R is the set of real numbers. The (expected) utility is the difference between the (expected) rewards $r(e_i) \in R_{\geq 0}$ that the agent will receive and the costs $c_i \in R_{\geq 0}$ she spends. Therefore, agent i 's (expected) utility U_{e_i} is calculated as follows:

$$U(e_i) = r(e_i) - c_i \quad (1)$$

where $r(i)$ is determined by the total rewards $M \in R_{\geq 0}$ and the contest success function P implemented by the stakeholders, and $c_i = \delta_i \cdot e_i$. Here, $\delta_i \in R_{\geq 0}$ refers to a private coefficient that determines the marginal cost that agent i incurs when she invests an extra unit of efforts. Thus, agent i 's (expected) utility can be computed by:

$$U(e_i) = r(e_i) - \delta_i \cdot e_i \quad (2)$$

Note that the thresholding agent model should not be confused with the threshold model of collective behavior (Granovetter 1978). The threshold model assumed that each agent's behavior depends on the number of other agents already adopting that behavior while the thresholding agent model only assumed that individual agents have their own private thresholds to trigger their behaviors.

In a contest, agent i exerts positive efforts $E_i > 0$ only if her expected utility $U(e_i) \geq 0$. Otherwise, she would be better off by not participating the contest since it would be irrational for her to exert efforts that bring no good or even harm to her. In practice, agent i 's efforts e_i are often bounded by her maximum ability due to various factors such as time (Giampietro, Bukkens, and Pimentel 1993), psychological

constraints (Hunter and Schmidt 1983) and working environments (Iverson and Zatzick 2011; Niemelä, Rautio, Hannula, and Reijula 2002). Therefore, we have $e_i \leq \hat{e}_i$ where \hat{e}_i is the maximum amount of efforts that agent i can invest in a contest.

2.2 Contest Success Function

A contest success function (CSF) determines the probability of winning for each player as a function of all the players' effort (Skaperdas 1996). In general, there are two types of contest success functions: the *ratio* form and the *difference* form (Hirshleifer 1989). The former determines the probability of winning according to the ratio of the respective efforts committed by each player. In the difference form, however, the probability of winning is computed as a function of the difference in the efforts that each player has exerted. In this paper, we focus on the ratio form of contest success functions because it can be naturally scaled up to large contests that involve millions of players while extending the difference form to these contests is usually non-trivial and difficult (Jia, Skaperdas, and Vaidya 2013).

The ratio-form CSF usually determines the probability of winning as a function of the ratio of the efforts that each player has committed. In a contest, agent i 's probability of winning is computed by:

$$p(e_i) = \begin{cases} \frac{f(e_i)}{\sum_{j \in I} f(e_j)} & \text{if } \sum_{j \in I} f(e_j) > 0 \\ \frac{1}{|I|} & \text{otherwise} \end{cases} \quad (3)$$

where $f(\cdot)$ is a non-negative, strictly increasing function, and $|I|$ is the number of players in the contest. The most widely used function for $f(\cdot)$ is $f(e_i) = e_i^\mu$ ($0 \leq \mu \leq 1$) due to Tullock (1980).

In our simulation analysis, we will concentrate on the Tullock function because of its broad applications in real-world scenarios such as technology race (Baye and Hoppe 2003), sports (Dietl, Franck, and Lang 2008), crowdsourcing (Shen, Feng, and Lopes 2019), lobbying and rent-seeking (Tullock 1980). Agent i 's (expected) rewards $r(e_i)$ for investing e_i efforts is calculated by:

$$r(e_i) = \begin{cases} \frac{(e_i)^\mu}{\sum_{j \in I} e_j^\mu} \cdot M & \text{if } \sum_{j \in I} e_j^\mu > 0 \\ \frac{1}{|I|} \cdot M & \text{otherwise} \end{cases} \quad (4)$$

where $0 \leq \mu \leq 1$ is the bias factor. The bias factor μ is often interpreted as the ‘‘noise’’ of a contest. It quantifies the effect that an increment in an agent's efforts on her probability of winning (Jia, Skaperdas, and Vaidya 2013). A contest that has low μ can be viewed as a lowly discriminating contest in which payers with different levels of efforts may have a similar degree of chance to win. In contrast, a contest that has high μ is a highly discriminating contest that favors players with high efforts. Despite the substantial understanding of the contest outcomes when agents are perfectly rational, little has been known about the contest dynamics with thresholding agents. In this paper, we present the first exploratory study to bridge the gap.

3 SIMULATING CONTESTS

In this section, we introduce the simulation platform that we have used in the experiments. Simulating large contests often involves updating data objects frequently because agents in the contest have their private information as well as different decision-making mechanisms. Therefore, it is desirable to simulate the contests in a modular approach.

We build our simulation platform based on a simulation framework called *SpaceTime* (Valadares, Lopes, Achar, and Bowman 2016; Lopes, Achar, and Valadares 2017; Achar and Lopes 2019). *SpaceTime* was initially developed for urban simulation in distributed environments (Shen, Achar, and Lopes 2018). In our paper, we extend it to large-scale contest simulations for the first time. The *SpaceTime* framework utilizes a programming language model called *Global Object Tracker* (GOT) that manages the version of objects

in a way that is similar to the popular version control tool *Git* (Spinellis 2012). Objects in SpaceTime are GOT nodes that consists of a dataframe, an application node, and the data operation component. The dataframe keeps track of both the object revision history and a snapshot of the current working data. The application node defines the logical process of the designated procedures. The data operation component allows the application node to manipulate the dataframe in the four ways: read, modify, add and delete.

In our simulation platform (see Fig. 1), there are two types of roles: the contest manager and the agents. The contest manager determines the contest success function and the number of the total rewards. The agents invest efforts and compete with each other for the rewards that are determined by the contest success function. Each agent is represented as a GOT node. In each node, the dataframe stores the agent’s private information including the maximum ability \hat{e}_i , the cost coefficient δ_i , the efforts to commit in the contest e_i and the expected utility $U(e_i)$. The decision-making mechanism is the application component that calculates the agent’s expected utility according to her private information and the public information obtained from the contest manager.

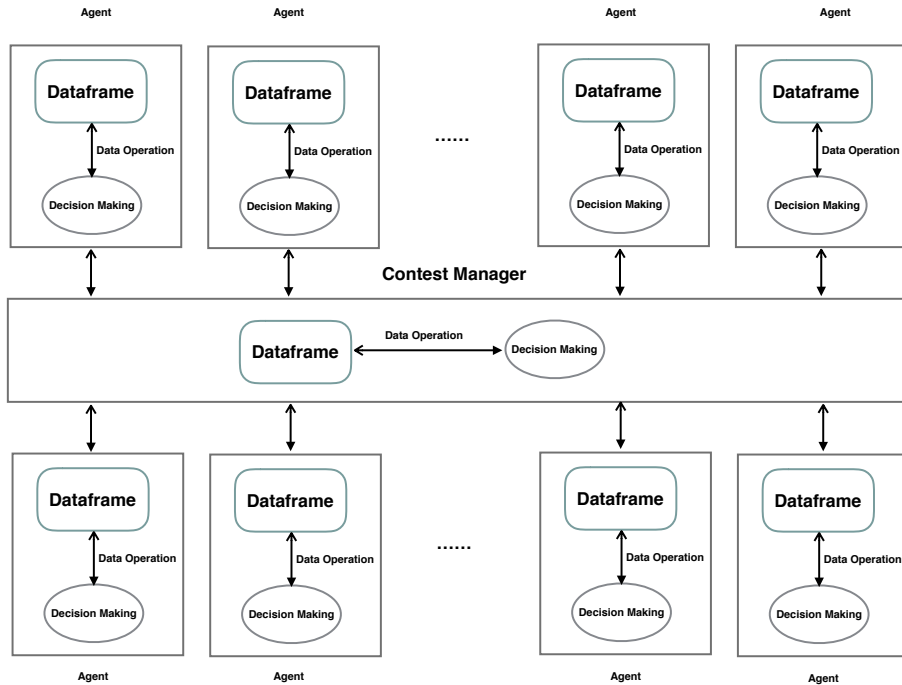


Figure 1: An architectural view of the simulation platform.

The contest manager is also a GOT node. The dataframe stores the contest information including the bias factor μ , and the number of total rewards M . The decision-making mechanism computes the probability of success for each agent given her invested efforts. The contest manager can request information (i.e. the number of efforts e_i) from the agents. Similarly, the agents can access information (i.e. the rewards for each player) from the contest manager. As a result, the majority of the data is stored locally and data communication between the agents and the contest manager reduces to the minimum level. The modular design makes the simulation platform highly expandable and configurable for large contests that have thousands of or even millions of players. In what follows, we discuss the role of contest manager and the individual agent in turn.

- **Contest Manager:** The contest manger is a representative of the stakeholder. It serves as the simulation engine.
 - *Dataframe:* The dataframe records initial parameters of the contest.

- *Data operator*: The contest manager set the parameters of the contests, update the records of the efforts, and the probability of winning, the rewards for each participating agent.
- *Decision making*: The goal of the stakeholder is to maximize the aggregated efforts of all the agents by selecting the optimal parameters for the contest. The contest manager can observe individual agents' decisions.
- **Agent**: Each agent is an actor in the simulation.
 - *Dataframe*: The dataframe stores the private threshold (the cost coefficient in our simulation analysis) and the expected rewards of the agent.
 - *Data operator*: The agent updates the expected rewards and the utilities according to the calculated probability of winning.
 - *Decision making*: The agent decides whether to contribute and how much efforts to invest according to the thresholding policies.

4 EXPERIMENTAL SETUP

After introducing the dataset used in the simulations, we describe the methods that we took to investigate how parameters of the contest success function influence the contest outcomes in contests with thresholding agents.

4.1 Dataset

We used the Kaggle ranked user data (Felipe Salvatore 2019) in our experiments. The dataset was obtained by a data crawling from the Kaggle competitions. The original dataset contains 4,767 rows. Each row represents a ranked player. For each row, there are six columns: the register date, the current points, the current ranking, the highest ranking, the country and the continent. In our work, we removed three irrelevant columns: the register date, the country, and the continent. That is, we only considered the following three columns: the current points, the current ranking, and the highest ranking.

The original dataset did not include all the necessary data fields for the experiments. Therefore, we performed data processing before conducting the experiments. We first normalized the players' current points by dividing each value by the maximum value observed in the data samples. We then took each player's normalized points as her efforts e_i (i.e., $e_i = i$'s normalized points). We estimated each agent's maximum ability \hat{e}_i according to the following equation:

$$\hat{e}_i = e_i \cdot \frac{i's \text{ Current Ranking}}{i's \text{ Highest Ranking}} \quad (5)$$

In our experiments, agent i 's maximum ability \hat{e}_i served as the upper bound of her efforts e_i exerted in a contest.

To calculate the expected utility, an agent needs to know her cost coefficient. Unfortunately, the original dataset did not provide such information. We estimated each agent's cost coefficient δ_i by assuming her utility was zero. This is without of loss generality because although agent i 's utility is usually her private information, it can be integrated into the cost coefficient δ_i . According to Eq. 2, we have

$$\delta_i = \frac{r(e_i)}{e_i} = \frac{4767 - i's \text{ Current Ranking}}{4767 \cdot e_i} \quad (6)$$

where e_i was estimated as the normalized points number. In our simulation, the cost coefficient served as the threshold of a player. If a player has a high cost coefficient, then the player needs to exert significant efforts in order to compete with others and receive the rewards. When the cost coefficient is sufficiently high, the player would not benefit from contributing to the tasks.

4.2 Methods

We identified six main performance metrics to quantify how the parameters of the contest success function affect the performance of the contests. The metrics include:

- the total efforts exerted by all the players (the higher, the better);
- the average efforts exerted by each participating player (the higher, the better);
- the total utility earned by all the players (the higher, the better);
- the maximum efforts exerted by a player in the contest (the higher, the better);
- the average utility earned by each participating player (the higher, the better);
- the maximum utility earned by a player in the contest (the higher, the better).

We also measured the standard deviation of the efforts (the lower, the better) and the standard deviation of the utility (the lower, the better) to quantify the variation of the efforts and utility, respectively. Note that all the numbers were normalized by dividing by the maximum value.

We studied how the bias factor μ influenced the contest outcomes by changing μ from 0.0 to 10. with an increment of 0.01. To investigate how the growth rates of the total rewards affect contest performance, we let the total rewards M grow linearly with different growth rates from 0.0 to 1.0 with an increasing step of 0.05.

We performed another two groups of experiments to study if the population of players or coefficients had any effect on the performance of the contests. To investigate how the number of participating players affected the contest outcomes, we varied the number of population from 0.25 to 5 times of the original population (i.e., 4767) with an increasing step of 0.25. If the selected population x was smaller than 4767, we randomly selected the samples from the current population. If the selected population x was larger than 4767, we first added $x/4767$ times of the current population, and then selected the remainder ($x \bmod 4767$) randomly. In the second group, we varied the cost coefficient of each player from 0 to 10 with an increment of 0.1 to study how agents' thresholding behavior influence the contest performance.

We ran each group of experiments for 100 times and reported the averaged numbers of each metrics in our paper. All the simulations were conducted on the same 3.7GHz 6-core Linux machine with 32GB RAM. It took approximately six hours to ran all the simulations.

5 RESULTS

Based on the data points obtained from the numerical simulations, we highlight four main observations. After presents the results, we summarize our findings.

5.1 Bias Factor

Observation 1 *As the bias factor increases, agents' efforts and utility typically increase significantly to the peak and then decreases gradually. Contests with a moderately low bias factor typically perform the best.*

This trend was observed in all the metrics: the total efforts, the average efforts, the maximum efforts, the standard deviation of efforts, the total utility, the average utility, the maximum utility, and the standard deviation of utility. Figure 2 demonstrates that as the bias factor μ increased from 0.1 to 1.0, the performance of the contests (i.e., the total efforts, the average efforts, the maximum efforts) first increased significantly to the peak and then decreased gradually. Contests with a moderately low bias factor (e.g. $0.1 \leq \mu \leq 0.4$) performed the best among the contests with different bias factors. An explanation for this phenomenon is that the majority of the participating agents are players with low ability, while contests with moderately low bias factors favor agents with relatively low abilities.

The number of efforts invested by each agent dispersed widely: while the maximum efforts could reach above 3.5, the average efforts were below 0.02 (see Figures 2c and 2b). A possible explanation is that the

agents’ coefficients and maximum ability varied substantially, resulting in a wide range of variations on the effort values (see Figure 8d). Agents’ committed efforts varied the most when the noise factor was low (e.g. $\mu = 0.1$).

Figure 3 shows that with the bias increasing, agents’ total utility, average utility, and maximum utility also experienced the same trend as the efforts. They also climbed to the peak when the bias factor was relatively low (e.g. $0.1 \leq \mu \leq 0.4$).

Observation 1 indicates that it is possible to implement a contest success function that serves the interests of both the stakeholders and the participating agents in contests with thresholding agents. That is, there can be a “win-win” situation when the stakeholders select a moderately low bias factor. This finding is surprisingly different from contests (with perfectly rational agents) in which the stakeholders’ interests and the players’ are usually conflicting.

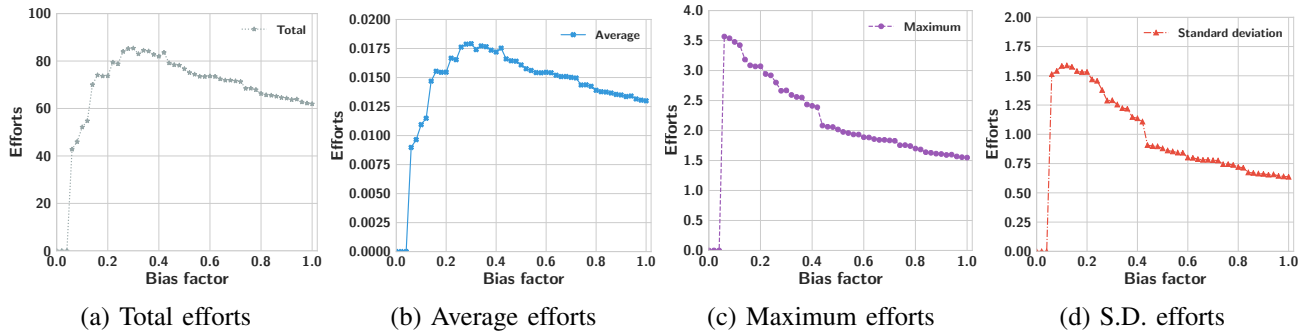


Figure 2: A comparison of four effort metrics by varying the bias factor from 0.0 to 1.0.

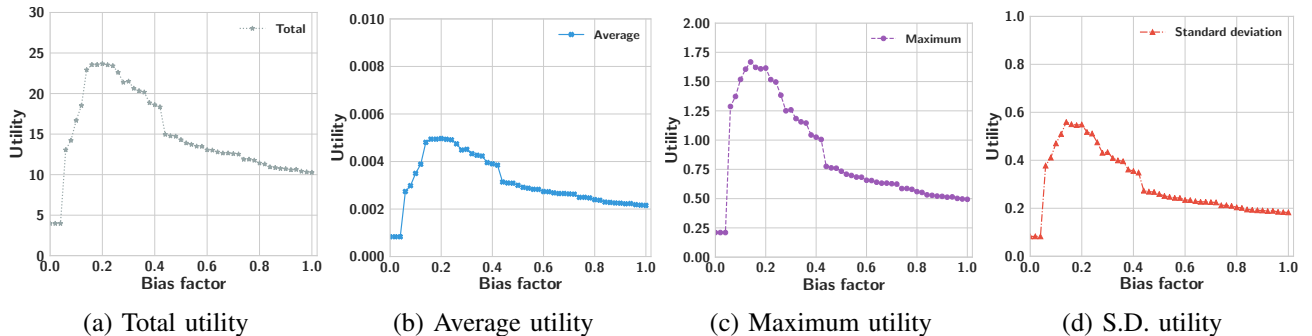


Figure 3: A comparison of four utility metrics by varying the bias factor from 0.0 to 1.0.

5.2 Total Rewards

Observation 2 *Increasing the amount of rewards according the total efforts has little impact on the contest outcomes in contests with different bias factors.*

Figures 4 and 5 shows that the total efforts, the average efforts, the total utility, and the average utility rarely changed when the reward growth parameter increased. That is, when the total amount of rewards grows linearly in the total efforts, the contest outcomes remained unchanged. Figures 4d and 5d confirmed that the variances of agents’ individual efforts and utility kept almost unaltered. Among contests with all the four bias factors, the one with the lowest bias factor (i.e., $\mu = 0.1$) performed the best while the one with the highest bias factor ($\mu = 0.4$) performed the worst. When it came to the maximum efforts and maximum utility, however, the results were quite the opposite: the contest with the highest bias factor

($\mu = 0.4$) performed the best while the one with the lowest bias factor ($\mu = 0.1$) performed the worst. Two reasons contribute to this trend. First, the majority of the players are the ones with low abilities. Second, when the bias factor μ is low, the contest favors the players with low skills. As a result, the contest can solicit efforts from more players. When the bias factor is high, the contest favors the players with high abilities. Because of it, the contest becomes much more appealing to top performers.

Observation 2 indicates that it is more beneficial for the stakeholders to choose a low bias factor when the total amount of rewards grows linearly in the total efforts. However, if the stakeholders' goal is to solicit the maximum individual efforts, a high bias factor is more desirable.

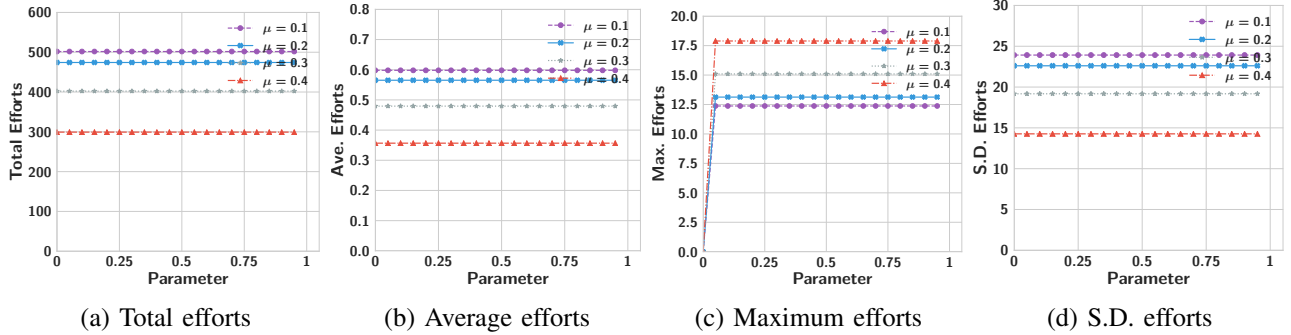


Figure 4: A comparison of four effort metrics by varying the reward growth parameter from 0.0 to 1.0 for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

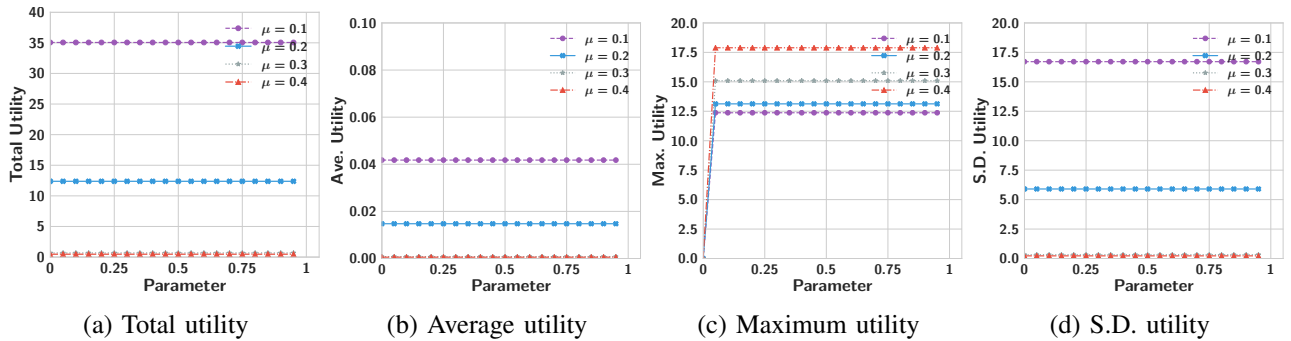


Figure 5: A comparison of four utility metrics by varying the reward growth parameter from 0.0 to 1.0 for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

5.3 Population Dynamics

Observation 3 *As the number of participating agents increases, agents' total efforts and maximum efforts fluctuate, while agents' utility typically decreases.*

Figure 6 shows that the total efforts fluctuated heavily at around 3 for contests with all the four bias factors (i.e. $\mu = 0.1, 0.2, 0.3$, and 0.4), and the maximum efforts changed back and forth between 0.3 and 2.8. Agents' individual efforts changed drastically (see Figure 6d). However, the average efforts first dropped significantly to a low level (0.0002) and then varied moderately. A reason for this trend is that the increase in the player population changed the degree of competition in the contests and subsequently changed the contest dynamics. As a result, it was possible that some agents would be better off not to participate (i.e. their expected utility becomes negative) under one contest dynamics while they would have

higher incentives to exert more efforts under another contest dynamics It is worth noting that stakeholders can usually obtain a positive amount of efforts from the agents. Stakeholders’ ability to solicit efforts is independent of the bias factors and the population of the players.

Figure 7 demonstrates that agents’ utility first dropped substantially and then kept at a similar level as the number of participating agents increased. This observation was expected because with the player population growing, the competition among the agents became more fierce. However, the total amount of rewards M remained unchanged. After the population had increased to a sufficiently large number (e.g. 1.5 times of the original population), the contest became less appealing to the many players because the competition was so stiff that these players could hardly make profits from participating in it.

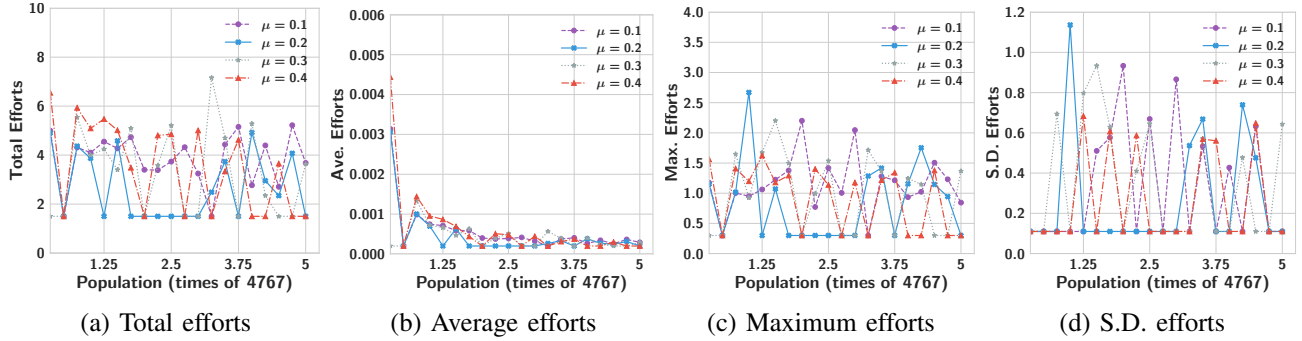


Figure 6: A comparison of four effort metrics by varying the population from 0.25 to 5 times of the original value for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

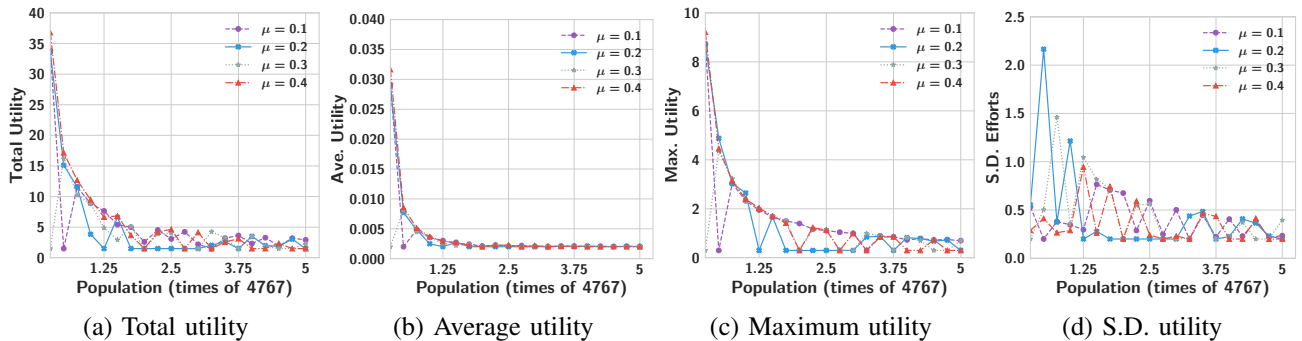


Figure 7: A comparison of four utility metrics by varying the population from 0.25 to 5 times of the original value for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

Observation 3 demonstrates that the population of participating agents has a tremendous impact on the contest outcomes. Despite it, stakeholders can usually have a guarantee that they receive a positive amount of efforts in contests with thresholding agents.

5.4 Coefficients

Observation 4 As agents’ coefficients increase, both their efforts and utility drop sharply to the lowest level and then kept steady.

Figures 8 and 9 demonstrate that agents’ efforts and utility decreased abruptly to the lowest level and then kept steady. This trend was observed in all the metrics for contests with all the four different bias factors. The underline reason for this trend is that the marginal cost for an agent to invest an additional unit

of efforts grows significantly as the coefficient increases. When the coefficient exceeds a level (e.g. 0.1), it becomes unprofitable for agents to participate in the contests. As a result, the majority of the participants chose to exert no efforts.

Observation 4 indicates that agents' coefficients have a significant impact on the performance of the contest outcomes before they become sufficiently large. When the coefficients go above the level, they do not affect the performance of the contest outcomes as most of the agents commit no efforts.

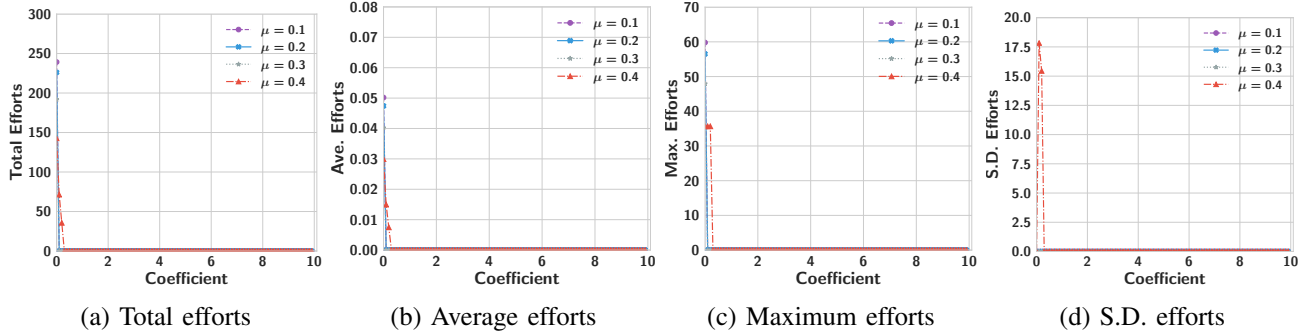


Figure 8: A comparison of four effort metrics by varying the coefficient from 0.1 to 10 times of the original value for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

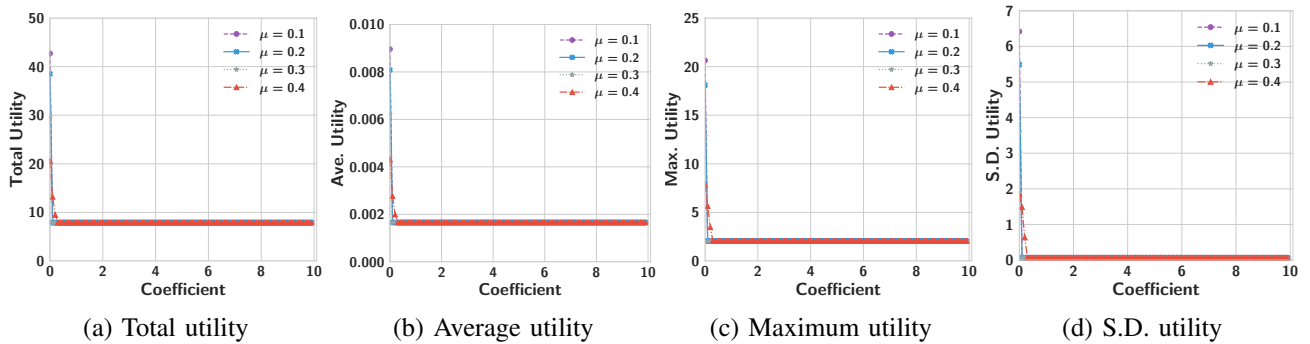


Figure 9: A comparison of four utility metrics by varying the coefficient from 0.1 to 10 times of the original value for contests with four bias factor (0.1, 0.2, 0.3, and 0.4).

5.5 Discussion

The four observations tell a similar story: in contests with thresholding agents, the contest outcomes are jointly influenced by the bias factor, the population dynamics, and each agent's coefficient; it is feasible for the stakeholders to select a bias factor that is (approximately) optimal for both the stakeholders and the agents. To this end, it is usually more desirable for the stakeholders to set the bias factor to be moderately low (e.g. $0.1 \leq \mu \leq 0.4$). This finding complements previous research on contests with perfectly rational agents that the stakeholders of the contests and the agents are often in a conflicting situation. While illuminating, our simulation analysis is by no means an exhaustive one due to the huge space of agents' private information.

Our study is intended to open a new line of research that can deepen our understanding of how to optimally coordinate thresholding agents. In doing so, we must build high-fidelity simulations based on real-world data, and develop new methods to efficiently learn and infer agents' decision models based on

their actual behavior. Addressing these challenges will augment humans' ability to coordinate systems with large populations of self-interested agents better.

6 CONCLUSION

In this paper, we describe a simulation study that quantifies how parameters of a contest success function and the agents influence the contest outcomes when non-cooperative agents use thresholding policies. We performed a series of experiments on a simulation platform that was built on top of a framework called SpaceTime. Experimental results demonstrate that although contest outcomes are jointly influenced by the bias factor, the population dynamics, and agents' coefficients, it is feasible for a stakeholder to design an (approximately) optimal contest that serves both the stakeholder's interests as well as the agents' by choosing a moderately low bias factor when the agents use thresholding policies. Our research sheds light on how to design proper competitions to coordinate thresholding agents' behavior for desirable outcomes.

Our work opens several exciting avenues for future research. In this paper, the thresholding agents do not adapt their coefficients dynamically based on experience. It would be interesting to study the scenarios when agents learn to adjust their coefficients. Our work is focused on one-stage contests. In practice, however, players need to go through multi-stage contests before they can win (Fu and Lu 2012). It would be worth investigating how to design optimal contests with thresholding agents. Another fruitful area is to develop new methods to select the bias factor for a given contest automatically. It is also interesting to investigate how to design incentive mechanisms to encourage thresholding agents to participate in the contests.

ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation through grant no. CCF-1526593.

REFERENCES

- Achar, R., and C. V. Lopes. 2019. "Got: Git, but for Objects". *arXiv preprint arXiv:1904.06584*.
- Almahmoud, Z., J. Crandall, K. Elbassioni, T. T. Nguyen, and M. Roozbehani. 2018. "Dynamic Pricing in Smart Grids Under Thresholding Policies". *IEEE Transactions on Smart Grid* 10(3):3415–3429.
- Baye, M. R., and H. C. Hoppe. 2003. "The Strategic Equivalence of Rent-Seeking, Innovation, and Patent-Race Games". *Games and Economic Behavior* 44(2):217–226.
- Blavatsky, P. R. 2010. "Contest Success Function with the Possibility of A Draw: Axiomatization". *Journal of Mathematical Economics* 46(2):267–276.
- Clark, D. J., and C. Riis. 1998. "Contest Success Functions: An Extension". *Economic Theory* 11(1):201–204.
- Corchón, L., and M. Dahm. 2010. "Foundations for Contest Success Functions". *Economic Theory* 43(1):81–98.
- Dietl, H. M., E. Franck, and M. Lang. 2008. "Overinvestment in Team Sports Leagues: A Contest Theory Model". *Scottish Journal of Political Economy* 55(3):353–368.
- Dimant, E., and C. Deutscher. 2015. "The Economics of Corruption in Sports: The Special Case of Doping". *Edmond J. Safra Working Papers* (55).
- Easley, D., and A. Ghosh. 2015. "Behavioral Mechanism Design: Optimal Crowdsourcing Contracts and Prospect Theory". In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 679–696. Portland, Oregon: Association for Computing Machinery.
- Felipe Salvatore 2019. "Ranked Users Kaggle Data". <https://www.kaggle.com/felsal/ranked-users-kaggle-data/>. accessed 04th February 2019.
- Fort, R., and J. Winfree. 2009. "Sports Really Are Different: The Contest Success Function and the Supply of Talent". *Review of Industrial Organization* 34(1):69–80.
- Fu, Q., and J. Lu. 2012. "The Optimal Multi-Stage Contest". *Economic Theory* 51(2):351–382.
- Giampietro, M., S. G. Bukkens, and D. Pimentel. 1993. "Labor Productivity: A Biophysical Definition and Assessment". *Human Ecology* 21(3):229–260.
- Gordon, B. R., and W. R. Hartmann. 2016. "Advertising Competition in Presidential Elections". *Quantitative Marketing and Economics* 14(1):1–40.
- Granovetter, M. 1978. "Threshold Models of Collective Behavior". *American Journal of Sociology* 83(6):1420–1443.

- Hirshleifer, J. 1989. "Conflict and Rent-Seeking Success Functions: Ratio VS. Difference Models of Relative Success". *Public Choice* 63(2):101–112.
- Hunter, J. E., and F. L. Schmidt. 1983. "Quantifying the Effects of Psychological Interventions on Employee Job Performance and Work-Force Productivity". *American Psychologist* 38(4):473.
- Iverson, R. D., and C. D. Zatzick. 2011. "The Effects of Downsizing on Labor Productivity: The Value of Showing Consideration for Employees' Morale and Welfare in High-Performance Work Systems". *Human Resource Management* 50(1):29–44.
- Jia, H., S. Skaperdas, and S. Vaidya. 2013. "Contest Functions: Theoretical Foundations and Issues in Estimation". *International Journal of Industrial Organization* 31(3):211–222.
- Lopes, C. V., R. Achar, and A. Valadares. 2017. "Predicate Collection Classes". *Journal of Object Technology* 16(2):3–1.
- Malueg, D. A., and A. J. Yates. 2010. "Testing Contest Theory: Evidence from Best-Of-Three Tennis Matches". *The Review of Economics and Statistics* 92(3):689–692.
- Münster, J. 2009. "Group Contest Success Functions". *Economic Theory* 41(2):345–357.
- Niemelä, R., S. Rautio, M. Hannula, and K. Reijula. 2002. "Work Environment Effects on Labor Productivity: An Intervention Study in A Storage Building". *American Journal of Industrial Medicine* 42(4):328–335.
- Ohannessian, M. I., M. Roozbehani, D. Materassi, and M. A. Dahleh. 2014. "Dynamic Estimation of the Price-Response of Deadline-Constrained Electric Loads under Threshold Policies". In *Proceedings of the 2014 American Control Conference*, 2798–2803. Portland, Oregon: Institute of Electrical and Electronics Engineers, Inc.
- Runkel, M. 2006. "Optimal Contest Design, Closeness and the Contest Success Function". *Public Choice* 129(1-2):217–231.
- Shen, W. 2019. *Beyond Nash Equilibrium: Mechanism Design with Thresholding Agents*. Ph. D. thesis, University of California, Irvine.
- Shen, W., R. Achar, and C. V. Lopes. 2018. "Toward Understanding the Impact of User Participation in Autonomous Ridesharing Systems". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 845–856. Gothenburg, Sweden: Institute of Electrical and Electronics Engineers, Inc.
- Shen, W., J. W. Crandall, K. Yan, and C. V. Lopes. 2018. "Information Design in Crowdfunding under Thresholding Policies". In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 632–640. Stockholm, Sweden: International Foundation for Autonomous Agents and Multiagent Systems.
- Shen, W., Y. Feng, and C. V. Lopes. 2019. "Multi-Winner Contests for Strategic Diffusion in Social Networks". In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 6154–6162. Honolulu, Hawaii: AAAI Press.
- Simon, H. A. 1997. *Models of Bounded Rationality: Empirically Grounded Economic Reason*, Volume 3. MIT Press.
- Skaperdas, S. 1996. "Contest Success Functions". *Economic Theory* 7(2):283–290.
- Spinellis, D. 2012. "Git". *IEEE Software* 29(3):100–101.
- Tullock, G. 1980. "Efficient Rent Seeking". In *Toward a Theory of Rent Seeking Society*, edited by J. Buchanan, G. Tullock, and R. Tollison, 97–112. Texas A&M University Press.
- Valadares, A., C. V. Lopes, R. Achar, and M. Bowman. 2016. "CADIS: Aspect-Oriented Architecture for Collaborative Modeling and Simulation". In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, 1024–1035. Washington, D.C.: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

WEN SHEN is a Ph.D. candidate in Information and Computer Science at the UC Irvine. He is interested in multi-agent systems and mechanism design. He holds a M.S. degree from Masdar Institute in Abu Dhabi, UAE and a B.E. degree from Northwestern Polytechnical University in Xi'an, China. His e-mail address is wen.shen@uci.edu.

ROHAN ACHAR is a PhD candidate in Software Engineering at UC Irvine. His research is on software engineering and programming languages. He received his M.S. at University of California, Irvine, and his B. Tech at National Institute of Technology Karnataka in India. His e-mail address is rachar@uci.edu.

CHRISTINA V. LOPES is a Professor of Informatics at UC Irvine. Her research focuses on software engineering for large-scale data and systems. She has a PhD from Northeastern University, and M.S. and B.S. degrees from Instituto Superior Tecnico in Portugal. Her email address is lopes@ics.uci.edu.