Proceedings of the Fifth Conference on
Human Computation and Crowdsourcing
(HCOMP 2017)

Effective Prize Structure for Simple Crowdsourcing Contests with Participation Costs

David Sarne, Michael Lepioshkin
Department of Computer Science
Bar Ilan University, Israel
sarned@cs.biu.ac.il,mishamikel@gmail.com

Abstract

This paper studies the use of a multi-prize compensation scheme for "simple" contests where participation is costly and the quality of participants' contributions is a priori uncertain at the time they make their decision related to participating in the contest. The equilibrium analysis provided enables demonstrating not only that a multi-prize structure is often beneficial but also that in some cases the principal's expected profit is maximized when offering a second prize greater than the first prize. This may seem somehow counter-intuitive especially given that the principal's profit is only influenced by the quality of the best submission rather than the aggregate of submissions. Special emphasis is placed on the case where the contestants are a priori homogeneous which is often the case in real-life, whenever the contestants are basically a priori alike and the quality of their submissions is determined subjectively by some referee. Here, we manage to prove that a multi-prize structure is dominated by a winner-takes-all scheme, suggesting that the benefit in the multi-prize contest scheme fully derives from the heterogeneity between prospective contestants. Finally, we show that there is a class of settings where the use of the multi-prize crowdsourcing contest model enables achieving the performance of the fully cooperative model (which is an upper bound for the performance in any type of contest), and that for settings of this class the optimal prize allocation can be extracted through a set of linear equations.

Introduction

A "Crowdsourcing Contest" is an important and nowadays highly popular crowdsourcing mechanism aiming to solicit effort of the crowd in solving problems (DiPalantino and Vojnovic 2009; Chawla, Hartline, and Sivan 2012; Liu et al. 2014; Vojnovi 2016). The contest typically specifies a well defined task the requester is interested in accomplishing and a reward scheme according to which contributors will be compensated for their efforts, with a strong correlation between the quality of the contribution one submits and the expected prize awarded.

There are numerous examples for crowdsourcing contests, perhaps the most famous is the Netflix Prize Challenge which was an open contest, with a one million USD prize, soliciting innovative movie-preference prediction algorithms (www.netflixprize.com/). Other examples include platforms such as Taskcn (www.taskcn.com/), TopCoder (www.topcoder.com) and Kaggle (www.kaggle.com). All three allow requesters to solicit contributions through contests with monetary prizes. The use of crowdsourcing contests is not limited to firms seeking technological solutions—these are nowadays broadly used in Q&A websites, allowing users to post questions and typically awarding "best" responses through points and reputation scores, by government agencies such as in Challenge.gov platform, with the goal of hosting challenges and prize contests to solve mission-centric problems and by non-profit organizations that seek major breakthroughs for the benefit of humanity, such as X Prize (www.xprize.org) and the Hult Prize (www.hultprize.org).

The key feature of contest-based crowdsourcing is its all-pay nature: online workers invest (costly and irreversible) efforts in producing solutions to a task, and only those whose solutions are selected as the winning ones are awarded a prize (Vojnovi 2016; DiPalantino and Vojnovic 2009; Chawla, Hartline, and Sivan 2012). As such, and much like with any crowdsourcing mechanism aiming to increase effort and participation of workers (Difallah et al. 2014; Elmalech et al. 2016; Faradani, Hartmann, and Ipeirotis 2011), the choice of the payments (prizes in our case) to be awarded to workers is fundamental.

In this paper we study multi-prize allocation for crowdsourcing contests where participants' strategy space is captured entirely by their decision of whether to participate or not to participate in the contest (Ghosh and Kleinberg 2016; Levy, Sarne, and Rochlin 2017). This kind of contest is applicable in any setting where participants do not know ahead of time or have no influence or control over the quality of their contributions. For example, consider the case of having to decide whether to agree to be nominated to an award. Obviously at the time of nomination a candidate has no influence whatsoever over her past achievements, typically accumulated over decades, though the recommendation to be made by the award committee is based entirely on those.

Prior work provides various other examples for contest settings of this kind, differing primarily in whether or not participants know the quality of their contributions in the contest at the time of making their decision. In this work we focus in contests where the quality of contribution is a priori uncertain. This is typically the case whenever submission quality is being determined subjectively by external judges based on some attributes, tastes and factors that are not fully disclosed. For example, in award nominations there is much uncertainty concerning how much the committee will be impressed by any specific achievement and what weight will be given to different details in one's CV. This is especially true when the identity of the committee members is not known, which is often the case in real life.

The choice of participation in contests of the above type is not trivial, as participation typically incurs a cost. This cost may be substantial as in the case of the reputational loss incurred when being nominated for a prestigious award and not winning. Indeed, it is not uncommon for people to choose not to become nominated if they think the chance of winning is small. In general the effect of the participation cost depends on the magnitude of the prize awarded and the number of prospective contestants, and even a small cost can lead to a mixed participation strategy whenever the prize is moderate.

Contributions. In this paper we study the usefulness of prize division (i.e., dividing the grand prize into several subprizes) in the above archetypal contest model. Our analysis is theoretical, and assumes participants are fully rational and aim to maximize their expected profit. The contributions of the paper are threefold. First, the equilibrium analysis provided enables demonstrating that, somehow counterintuitively, despite the fact that the contest organizer is interested only in the best performance obtained (e.g., getting the most deserving candidate for the award), allocating the prizes budget in the form of several prizes is beneficial, as far as the organizer's expected profit is concerned. As we discuss in more detail in the related work section, the idea of prize-splitting per-se is not new - it has been researched in literature (Ghosh and Hummel 2015; Kaplan et al. 2002; Moldovanu and Sela 2001) and used in practice.[1] Still, it has been shown to be effective mostly in contest models that assume the contest organizer benefits from the sum of yields (hence the increase in benefit there is quite intuitive) (Ghosh and Kleinberg 2016; Luo, Tan, and Xia 2014; Koutsopoulos 2013) or in models where contributors strategize on the effort they exert (Moldovanu and Sela 2001; Archak and Sundararajan 2009), and in various others it has been proved ineffective (Chawla, Hartline, and Sivan 2012; Liu et al. 2014).

Second, we show that it is possible (setting-dependent) that the optimal allocation of prizes is such that the j-th ($j > 1$) prize is greater than the first prize. Meaning that the winner envies others that actually performed worse in the contest. To the best of our knowledge, this phenomena has not been demonstrated in any prior contest model studied. The increase in the organizer's profit derives from various dynamics the prize division generates and the paper illustrates some of them. Furthermore, we prove that the organizer's benefit from splitting the prize is fully attributed to the heterogeneity of contestants since in the case of a priori homogeneous contestants the winner-takes-all scheme is the dominating one.

Finally, we show that in some settings, through prize division the contest organizer can actually reach a profit equivalent to the fully cooperative case, which is an upper bound to the expected profit in any contest. Furthermore, in such case the optimal (expected-profit-maximizing) prize allocation can be extracted using a set of linear equations.

[1] E.g., TopCoder contests typically offer two prizes, where the second prize is half the amount of the first prize.

The Model

The model considers a contest organizer (or principal, denoted "manager" onwards) and a set $A = \{A_1, ..., A_k\}$ of contributors (denoted "agents" onwards) that can potentially participate in (i.e., "contribute" to) the contest. The contest is considered "simple" in the sense that the quality of the contribution of a participating agent is beyond its control—it can only decide whether to participate or not participate in the contest (Ghosh and Kleinberg 2016; Levy, Sarne, and Rochlin 2017). The agents are heterogeneous in terms of their competence and cost of taking part in the contest - the quality of the contribution made by agent A_i if participating (or its "performance", for short) is a random variable characterized by a probability distribution function $f_i(x)$ and its cost of taking part in the contest is denoted c_i.

Agents are assumed to be fully rational and self-interested. Their participation in the contest thus depends on the compensations (in the form of prizes) offered by the manager. We use M to denote the prize budget available to the manager. The manager can divide this amount into any set of prizes $\{M_1, ..., M_n\}$ ($n \leq k$) such that M_i is the prize to be awarded to the i-ranked participating agent ($\sum M_i = M$), where ranking is determined according to quality of contributions made (with no requirement to exceed any performance threshold in order to become eligible). The model assumes that a prize can be awarded only to agents participating in the contest, i.e., if the number of participating agents is $k' < n$ then only prizes $M_1, .., M_{k'}$ will be awarded. Similar to most prior work (Chawla, Hartline, and Sivan 2012; DiPalantino and Vojnovic 2009; Ghosh and Kleinberg 2016) our model assumes full and symmetric information. Specifically, in our case it is assumed that the manager and all the agents are familiar with the prize allocation $\{M_1, ..., M_n\}$ and the individual distributions $f_1(x), ..., f_k(x)$ and costs $c_1, ..., c_k$.[2]

The goal of each agent is to maximize its expected profit, defined as the expected prize it is being awarded minus its participation cost whenever participating. The manager is assumed to be interested only in the highest quality contribution among those participating agents and has no value for lesser contributions (Chawla, Hartline, and Sivan 2012; Ghosh and Kleinberg 2016; Liu et al. 2014; Moldovanu and

[2] Alternatively we can assume that each agent is of a specific type t, where type t is characterized by a distribution $f_t(x)$ and cost c_t, and the manager and agents are familiar with the type distribution. This will require very minor modifications in the analysis and all the main results will still hold.

Sela 2001; Archak and Sundararajan 2009). Assuming the qualities of contributions can be expressed in terms of their monetary value to the manager, the manager's goal is to allocate the prize budget M such that the expected highest contribution achieved minus the expected overall prize payments made is maximized. For simplicity we take the expected highest contribution to be of zero quality in case all agents opt not to participate in the contest though the extension of the analysis to incorporate some pre-set fallback quality v_0 is trivial.

Equilibrium Analysis

We provide a formal equilibrium analysis to the crowdsourcing contest model described above under the multi-prize scheme. To illustrate findings, we use a tractable synthetic setting that simplifies calculations yet enables demonstrating the main solution characteristics in a clean manner (i.e., eliminating the need to isolate external phenomena that are commonly present in simulated or real-life applications).

We use $\{P, \neg P\}$ to denote the set of actions available to each agent, corresponding to participating and not participating in the contest, respectively. A strategy can thus be represented by the probability p ($0 \leq p \leq 1$) of participating in the contest. A strategy profile is a vector $S = (p_1, p_2, ..., p_k)$ such that p_i is the probability agent A_i will participate in the contest.

We use $G_i(x, l)$ to denote the probability that agent A_i will end up ranked l in the contest if participating and its eventual contribution quality turns to be x. The value $G_i(x, l)$ obtains is given by:

$$G_i(x,l) = \sum_{A' \subseteq A_{-i} \wedge |A'|=l-1} \left(\prod_{A_j \in A'} p_j(1-F_j(x)) \prod_{A_j \in A_{-i} \setminus A'} (p_j F_j(x) + (1-p_j)) \right) \quad (1)$$

where A_{-i} is the set of all agents other than A_i and $F_j(x)$ is the cumulative distribution functions of $f_j(x)$. Ending up ranked l happens if exactly $l-1$ agents perform better. Therefore, the calculation iterates over all possible subsets of A_{-i} with size $l-1$ and multiplies the probability of having all agents of that set perform better than x by the probability that all remaining agents (i.e., in $A_{-i} \setminus A'$) perform worse than x or do not participate at all.

This leads to the calculation of the expected profit of any agent A_i that participates in the contest, denoted $B_i(P)$, which is a simple integration over all possible performance values y minus the cost of participation:

$$B_i(P) = \int_{y=-\infty}^{\infty} \sum_{1 \leq j \leq n} (M_j G_i(y,j)) f_i(y) dy - c_i \quad (2)$$

If not participating, the expected profit is $B_i(\neg P) = 0$.

An equilibrium is therefore a solution $(p_1, ..., p_k)$ such that all agents are using their best-response strategy given the strategies of the others, i.e.: (a) $B_i(P) \leq B_i(\neg P)$ for every agent A_i that uses $p_i = 0$; (b) $B_i(P) \geq B_i(\neg P)$ for every agent A_i that uses $p_i = 1$; and (c) $B_i(P) = B_i(\neg P)$ for every agent A_i that uses $0 < p_i < 1$. It is possible that a given setting will have more than a single equilibrium solution (i.e., multi-equilibria), though the question of which of those will be used is beyond the scope of the current paper.

Figure 1 illustrates the changes in the equilibrium obtained as a function of the portion allocated for second prize out of the total prize budget M, denoted α, in a two-prize scheme (i.e., $\alpha = M_2/M$). Note that $\alpha = 0$ corresponds to winner-takes-all whereas $\alpha = 1$ to allocating the entire prize budget to the second prize. It is based on a setting of two agents - a "strong" and a "weak" one. The strong agent differs from the weak one in the sense that its underlying distribution function and participation cost are more favorable. Specifically, the strong agent uses a uniform distribution function over the interval $(0, 2)$ and its participation cost is 0.1. The weak agent uses a uniform distribution function over the interval $(0, 1)$ and its participation cost is 0.23. The total prize budget is $M = 0.4$. As can be seen in the figure, for $\alpha < 0.66$ the only equilibrium is the one where only the strong agent participates in the contest. The weak agent has no incentive to deviate from not participating, as even with a second prize of $2M/3$ its expected profit is negative, due to its substantial participation cost. For $\alpha > 0.66$ having both agents participate is in equilibrium. Here, neither of the agents have an incentive to deviate to not participating, as they win the first prize with probability 0.75 (for the strong agent) and 0.25 (for the weak agent) and the second prize with probability 0.75 and 0.25, respectively, which means they always end up with a positive profit. Interestingly, within the interval $(0.75, 1)$ we find two additional equilibria. The first is where neither participate and the second is when they both mix between participating and not participating, ending up with a zero profit (where the participation probability of the weaker player increases and the participation probability of the stronger player decreases as α increases). This calls for some intuitive explanation, especially given the fact that these equilibria co-exist alongside an equilibrium according to which both agents participate. We begin with the coexistence of the equilibrium according to which both agents participate along the one where neither participate. This can happen only when the second prize offered is greater than the first prize. Here, each agent actually benefit from the participation of the other, as this brings in the option to gain the second (greater) prize. Without the other agent in the contest, the agent can only win the first prize, which is less than the participation cost for

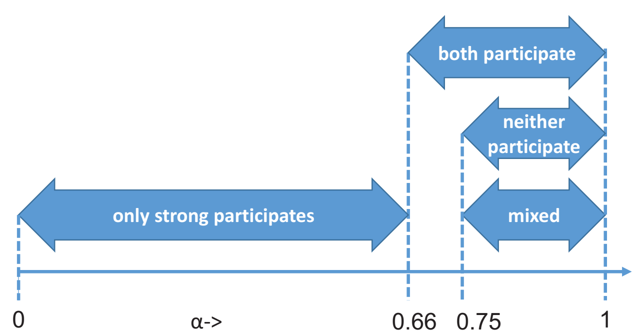

Figure 1: Different equilibria that hold for different prize allocations. See main text for the parameters of the setting used.

both. The coexistence of the mixed equilibrium along the pure one according to which both participate is even more counter-intuitive. If both agents find it optimal to participate when the other participates (according to the pure equilibrium) then why should a solution according to which they both mix be stable? The answer is once again based on the contribution of the second prize to the individual profit—the agent actually loses from the fact that the other agent is mixing between participating and not participating, as it is now less likely that it will win the second prize, which enables the indifference between participating and not participating for the agent itself.

We now turn to calculating the manager's expected profit from the contest, denoted $B^{manager}$. For this purpose we first calculate the probability that the best performance obtained in the contest is lower than or equal to x, denoted $F^M(x)$:

$$F^M(x) = \prod_{A_j \in A} (p_j \cdot F_j(x) + (1 - p_j)) \qquad (3)$$

Consequently, the expected profit of the manager is:

$$B^{manager} = -M + \int_{y=-\infty}^{\infty} y \frac{d(F^M(y))}{dy} dy \qquad (4)$$
$$+ \sum_{(A' \subset A) \wedge (0 \leq |A'| \leq n-1)} \left((\prod_{A_j \in A'} p_j)(\prod_{A_j \in A \setminus A'} (1 - p_j)) \sum_{j=|A'|+1}^{n} M_j \right)$$

where the second term is the expected quality of the first-ranked contribution and the third term is the expected unawarded prize (iterating over all cases where the subset of agents participating in the contest is smaller than the number of prizes offered n).

Figure 2 illustrates the effect of splitting the prize budget over the manager's expected profit (top graph) as a function of the portion of the second prize out of the total prize budget (α) in the setting used for Figure 1 above. The two bottom graphs in the figure decompose the expected profit into the expected performance (left) and the expected prize payment (right). The equilibrium used is the one that maximizes the manager's expected profit (out of those detailed in Figure 1) for the corresponding α value, though in this example any different selection criterion does not change the observations that follow, qualitatively. Here, the increase in the second prize (at the expense of the first prize) results in an increase in the manager's expected profit, as long as the equilibrium is to have only the strong agent participate in the contest (as the expected performance does not change, while the prize payment decreases). At $\alpha = 0.66$ the equilibrium changes to having both agents participate. For any setting where this equilibrium holds the manager's expected profit is fixed, as both the expected performance and the payment made (total of first and second prize) do not change. For $\alpha > 0.75$ we obtain three equilibria, as depicted in Figure 1. Among these, the one that results in the maximum expected profit for the manager in the interval $(0.75 - 0.87)$ is the mixed one. Here, the increase in α results in an increase both in the expected performance and in the expected prize awarded—both phenomena explained by the fact that the increase in

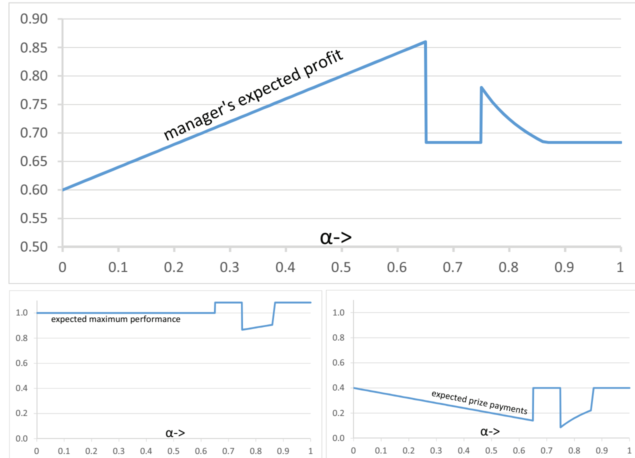

Figure 2: The manager's expected profit (top), the expected maximum performance achieved (bottom left) and the expected prize payment (bottom right) as a function of the portion of the second prize out of the total prize budget (α). See main text for the parameters of the setting used.

the participation probability of the weaker agent is substantially greater than decrease in the participation probability of the stronger agent. Still, the sum of the two components decreases as α increases, as observed in the top graph. Finally, for $\alpha > 0.87$ the equilibrium that maximizes the manager's expected profit is the one where both agents participate.

An interesting observation made based on Figure 2 is that it is possible that the manager's expected profit is maximized when the second prize is greater than the first prize. This, as mentioned earlier, is a highly counter-intuitive phenomena that has not been reported to hold in prior contest-theory literature. In this example, it is maximized for $\alpha = 0.66$. The improvement in the manager's expected profit derives from the fact that the prize division enables saving some portion of the prize without having any of the agents change its participation decision (and therefore we obtain the same performance level for a lesser expense). This, however, is not the sole source of improvement that can be achieved through prize splitting in our model, and not the only dynamic explaining the phenomenon according to which the manager's expected profit is maximized for $\alpha > 0.5$. A different (and perhaps more intuitive) dynamic that leads to the phenomenon is the ability to push more agents to participate in the contest through offering (possibly even greater) second prize. This is illustrated in Figure 3 which depicts the manager's expected profit for a 2-agents setting where the strong agent uses a uniform distribution function over the interval $(0, 3.3)$ and its participation cost is 0.165 whereas the weak agent uses a uniform distribution function over the interval $(0, 2.3)$ and its participation cost is 0.22. Here, as α increases the equilibrium changes from having only the strong agent participate to having both agents not participating, having both participating and finally having both not participating. The maximum expected profit is achieved at $0.67 \leq \alpha \leq 0.78$ when both agents participate in the con-

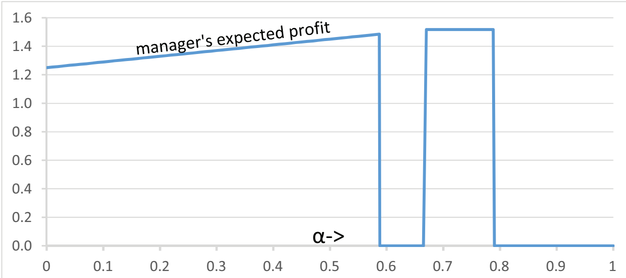

Figure 3: The manager's expected profit as a function of the portion of the second prize out of the total prize budget (α). Here the improvement is attributed to the ability to push the weak agent to participate through prize splitting. See main text for the parameters of the setting used.

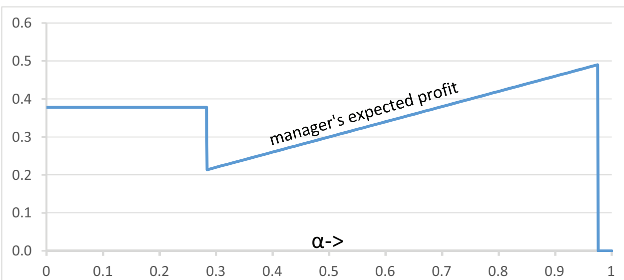

Figure 4: The manager's expected profit as a function of the portion of the second prize out of the total prize budget (α). Here the improvement is attributed to the ability to incentivize one of the agents to drop from the contest through prize splitting. See main text for the parameters of the setting used.

test.

Finally, we provide an example where the benefit from offering a greater second prize derives from the ability to deliberately incentivize one of the players to drop from the contest. This is illustrated in Figure 4. In this setting one of the agents' performance is drawn from a uniform distribution over the interval $(0, 1)$ and its participation cost is $c_1 = 0.01$ and the other agent uses a uniform distribution over $(0, 1.3)$ and its participation cost is $c_2 = 0.22$. For a relatively small second prize the equilibrium is based on having both agents participate. However, when the second prize is greater than 0.28 of the total prize budget the equilibrium changes to having only the first agent participate. Here, the expected profit of the manager increases as α increases, as it only pays the first prize, and reaches its peak at $\alpha = 0.97$ (i.e., allocating 97% of the total prizes to second prize and only 3% to first prize).

The above illustrated phenomena are not limited to cases of two agents. In fact, these occur to a greater extent in settings with more agents. Consider for example the following five-agents setting that uses a two-prize scheme. Out of the five agents, one (denoted "strong") uses a uniform distribution function over the interval $(0, 2)$ and its participation cost is 0.32. The others (denoted "weak") use a uniform distribution function over the interval $(0, 1)$ and their participation cost is 0.2. The total prize budget is $M = 1$. Figure 5 illustrates the changes in the equilibrium obtained in this setting as a function of α, the portion allocated for second prize out of the total prize budget M. Pure-strategy based equilibria hold along the the entire α interval, differing in the type of agents participating in the contest (see blue arrows). As α increases, we observe a transition from equilibria where only the strong and one weak agent participate in the contest (for $\alpha < 0.13$) to ones where the strong and two weak agents participate (for $0.13 \leq \alpha \leq 0.45$), to ones where the strong and three weak agents participate (for $0.45 \leq \alpha \leq 0.61$), to one where only the four weak agents participate (for $0.57 \leq \alpha$, overlapping the former ones within $0.57 \leq \alpha \leq 0.61$) and finally one where none of the agents participate (for $0.8 \leq \alpha$, overlapping the former one). A mixed equilibrium according to which th strong agent participates and all others mix holds within $\alpha < 0.59$ (where the participation probability p of the weak agents increases as α increases. Another mixed equilibrium, this time however where all agents mix, is obtained within the interval $0.57 \leq \alpha \leq 0.61$. In this equilibrium the participation probability p_s of the strong agent increases as α increases, whereas the participation probability p_w of the weak agents decreases as a function of α.

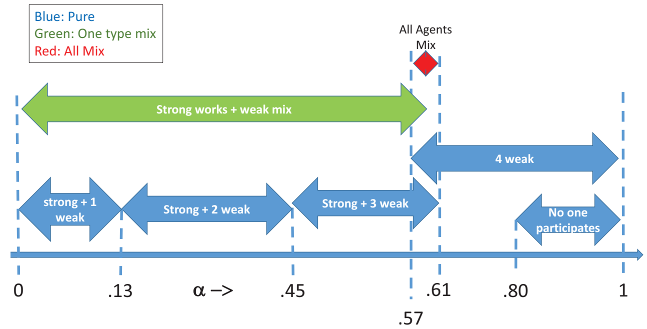

Figure 5: Different equilibria that hold for different prize allocations. See main text for setting parameters.

Figure 6 depicts the manager's expected profit as a function of α taking the equilibrium that provides the highest expected profit. For $\alpha < 0.03$ this is the pure equilibrium where only the strong agent and one of the other agents participate. For $\alpha < 0.59$ it is the mixed equilibrium. Then, for $0.59 < \alpha < 0.61$ it is the pure equilibrium where only the strong agent and three weak agents participate. For $0.61 < \alpha < 0.8$ it is the equilibrium where the four weak agents participate and finally for $0.8 < \alpha$ the one where none of the agents participate. As can be observed from the figure, the expected profit of the manager is indeed maximized at a ratio $\alpha > 0.5$. Note that even if taking the other equilibria in cases where the pure strategy equilibria are used in the graph (i.e., always preferring the symmetric equilibria according to which all weak agents use the same strategy) this latter result does not change and the expected-profit-maximizing α is still greater than 0.5.

Figure 6: The expected profit of the manager for different prize allocations. See main text for setting parameters.

Taking a close view at the dynamics that lead to a preference for prize-splitting, in this case it is the ability to leave the strong player in the game (using a pure-participation strategy, while increasing the participation probability of the other agents through allocating a greater amount to second prize at the expense of first prize). Interestingly, here with $\alpha = 0$ (i.e., without prize-splitting) the mixed equilibrium is dominated by the pure equilibrium according to which only the strong agent and one of the weak agents participate. It is the prize splitting that accounts for the change in the dominance relationship between the two, providing a greater profit through the mixed equilibrium.

The Homogeneous Case

In many real-life contest settings agents are a priori homogeneous in the sense that they are all characterized by a similar participation cost ($c_i = c$, $\forall i$) and the quality of their submissions derives from a similar probability distribution function ($f_i(x) = f(x)$, $\forall i$). This is very common whenever the differences in competence are minor thus their performance at the time of contest is mostly influenced by some probabilistic factors (e.g., luck, refereeing and weather conditions) or whenever contributions are evaluated subjectively by a referee whose taste cannot be a priori predicted. In fact, the majority of prior work on crowdsourcing contests, and in particular work that used its auction-mapping (either all-pay or winner-pay), as well as general contest-theory literature, has considered contestants to be homogeneous in the sense that they are either ex post (i.e., have exactly the same type) or ex ante (i.e., their types follow the same probabilistic distribution) identical (Koutsopoulos 2013; Luo, Tan, and Xia 2014; Chawla, Hartline, and Sivan 2012; DiPalantino and Vojnovic 2009; Archak and Sundararajan 2009; Moldovanu and Sela 2001; Cohen, Kaplan, and Sela 2008; Kaplan et al. 2002; Levy, Sarne, and Rochlin 2017). Interestingly, we find that with the homogeneous variant of the model studied in this paper, the manager always maximizes its expected profit by allocating its entire prize budget to the first prize (i.e., winner-takes-all).

When the agents are homogeneous a natural mixed equilibrium that holds (in the absence of equilibria in which all agents participate or do not participate) is the symmetric equilibrium by which all agents use the same participation probability p. This equilibrium can coexist alongside pure-strategy equilibria in which some of the agents participate, however since the agents are a priori homogeneous it is the most appealing and fair one, and therefore the one we relate to.

Assume we have k agents and we divide M into n prizes. The expected profit of each agent if participating is given by:

$$B(P) = \sum_{i=n-1}^{k-1} \frac{M}{i+1} \binom{k-1}{i} p^i (1-p)^{k-i-1} + \sum_{i=0}^{n-2} \frac{\sum_{w=1}^{i+1} M_w}{i+1} \binom{k-1}{i} p^i (1-p)^{k-i-1} \quad (5)$$

The first term relates to the case where there are $i \geq n-1$ other agents competing, therefore the entire prize budget M is awarded, with an equal chance for any of the $i+1$ competing agents to receive any prize M_w ($w \leq n$). The second term relates to the case where we have $i < n-1$ other agents competing, therefore a prize budget of $\sum_{w=1}^{i+1} M_w$ is awarded. Based on $B(P) = B(\neg P) = 0$ we can extract the equilibrium p.

We first prove that in the homogeneous case, with a single prize (winner-takes-all) the agents will choose to participate in the contest at least to the extent of participation achieved through splitting the prize budget.

Proposition 1. *When switching from a multi-prize to a single prize in the homogeneous case, the equilibrium value of p (i.e., the participation probability) cannot decrease.*

Proof. If the equilibrium with the optimal multi-prize case is having all agents opt not to participate then necessarily $c > M_1$. Therefore using a single prize $M > M_1$ will either result in a similar equilibrium (if $c > M$) or with one where $p > 0$ (otherwise). Similarly, if the equilibrium with the optimal multi-prize case is to have all agents participate then necessarily $c < \sum M_i/k = M/k$ and therefore the same equilibrium holds using a single prize M. Therefore all that is left to prove is that when the equilibrium with the optimal multi-prize case is to use $0 < p < 1$ then switching to a single prize M results in an equilibrium characterized by at least the same p value.
From (5) we obtain:

$$B(P) = \sum_{i=0}^{k-1} \frac{M}{i+1} \binom{k-1}{i} p^i (1-p)^{k-i-1} + \sum_{i=0}^{n-2} \frac{(\sum_{w=1}^{i+1} M_w - M)}{i+1} \binom{k-1}{i} p^i (1-p)^{k-i-1} \quad (6)$$

Note that $(\sum_{w=1}^{i+1} M_w - M)$ is negative. Therefore since we require $B(P) = B(\neg P) = 0$ in equilibrium, we need the term $\sum_{i=0}^{k-1} \frac{M}{i+1} \binom{k-1}{i} p^i (1-p)^{k-i-1}$ to be greater than zero. With a single prize M we obtain an equilibrium $p = p^*$ value such that the latter term equals zero as it captures $B(P)$ and we require $B(P) = B(\neg P) = 0$. Therefore, we need to prove that $\sum_{i=0}^{k-1} \frac{M}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$ can be positive only by using $p < p^*$, meaning that the term increases as p decreases. Using some standard algebraic manipulations obtains:[3]

$$\sum_{i=0}^{k-1}\frac{1}{i+1}\binom{k-1}{i}p^i(1-p)^{k-1-i} = \frac{1}{k}\sum_{i=0}^{k-1}\binom{k}{i+1}p^i(1-p)^{k-1-i}$$
$$= \frac{1}{k}\sum_{i=1}^{k}\binom{k}{i}p^{i-1}(1-p)^{k-i} = \frac{1}{kp}\sum_{i=1}^{k}\binom{k}{i}p^i(1-p)^{k-i}$$
$$= \frac{1}{kp}(1-(1-p)^k)$$

Now taking the derivative with respect to p we get:

$$-\frac{1}{kp^2} + \frac{kp(1-p)^{k-1} + (1-p)^k}{kp^2} \tag{7}$$

We want to show that it is less than or equal to 0. With further algebraic manipulations we get:

$$(1-p)^{k-1}((k-1)p+1) \leq 1$$

But: $1 + (k-1)p \leq (1+p)^{k-1}$, so:

$$(1-p)^{k-1}((k-1)p+1) \leq (1-p)^{k-1}(1+p)^{k-1}$$
$$= (1-p^2)^{k-1} \leq 1$$

□

The intuition behind Proposition 1 derives directly from Equation 5 - in the multi-prize case there is a chance that a portion of the prize budget is not awarded (whenever there are less participating agents than prizes), therefore there is a lesser incentive to participate overall.

Proposition 1 by itself does not guarantee that the manager's expected profit cannot improve by using more than a single prize. Indeed with the decreased p the expected contribution of each agent decreases, however at the same time the decrease in p results in awarding a smaller expected prize. Theorem 1 suggests that even when weighing in the tradeoff between the decrease in p and the decrease in the expected prize awarded, the expected profit of the manager always decreases due to splitting the prize.

Theorem 1. *In the homogeneous case the expected profit of the contest manager cannot be improved by splitting the prize budget into two or more prizes, i.e., winner-takes-all is the dominating prize scheme.*

Proof. Consider the single prize m' in the winner-takes-all case that results in the same p as when using the optimal n-prize scheme. Since in both cases $B(P) = B(\neg P) = 0$, we obtain (using (6)):

$$\sum_{i=0}^{k-1}\frac{M}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1} + \tag{8}$$
$$\sum_{i=0}^{n-2}\frac{(\sum_{w=1}^{i+1}M_w - M)}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$$
$$= \sum_{i=0}^{k-1}\frac{m'}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$$

Resulting in:

$$\sum_{i=0}^{k-1}\frac{M-m'}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1} + \tag{9}$$
$$\sum_{i=0}^{n-2}\frac{(\sum_{w=1}^{i+1}M_w - M)}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1} = 0$$

And since $\sum_{i=0}^{k-1}\frac{M-m'}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1} = \frac{(1-(1-p)^k)(M-m')}{kp}$:

$$\frac{(1-(1-p)^k)(M-m')}{kp} \tag{10}$$
$$= \sum_{i=0}^{n-2}\frac{(M - \sum_{w=1}^{i+1}M_w)}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$$

Now from the contest manager's point of view the expected best contribution in the contest is the same (because the agents are using the same p both when using the prizes $M_1, ..., M_n$ and when using the single prize m'). However with the single-prize contest the manager needs to award an expected prize of $m'(1-(1-p)^k)$ and with the n-rewards it awards $M - \sum_{i=0}^{n-1}\sum_{w=i+1}^{n}M_w\binom{k}{i}p^i(1-p)^{k-i}$. Subtracting the amount the manager pays in the two cases we obtain:

$$M - \sum_{i=0}^{n-1}\sum_{w=i+1}^{n}M_w\binom{k}{i}p^i(1-p)^{k-i} - m'(1-(1-p)^k) = \tag{11}$$
$$M(1-(1-p)^k) - \sum_{i=1}^{n-1}\sum_{w=i+1}^{n}M_w\binom{k}{i}p^i(1-p)^{k-i} - m'(1-(1-p)^k)$$
$$= (M-m')(1-(1-p)^k) - \sum_{i=1}^{n-1}\sum_{w=i+1}^{n}M_w\binom{k}{i}p^i(1-p)^{k-i}$$

Notice that $\sum_{i=1}^{n-1}\sum_{w=i+1}^{n}M_w\binom{k}{i}p^i(1-p)^{k-i} = pk\sum_{i=0}^{n-2}\frac{(M-\sum_{w=1}^{i+1}M_w)}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$. And also $(1-(1-p)^k)(M-m') = kp\sum_{i=0}^{n-2}\frac{(M-\sum_{w=1}^{i+1}M_w)}{i+1}\binom{k-1}{i}p^i(1-p)^{k-i-1}$ (according to (10)). Therefore the difference captured by (11) above is zero. Meaning that using m' such that it provides the same p as with the optimal n-prize contest will get us the exact same expected profit (as with using the optimal n-prize contest). And since m' is not the optimal prize to use with the single prize contest, the optimal single prize will result in an even greater expected profit. Therefore, winner-takes-all dominates any multi-prize scheme when the agents are a priori homogeneous. □

We emphasize that the winner-takes-all scheme was found to be the dominating one in several other contest models with homogeneous contestants studied in prior literature (Taylor 1995; Ghosh and McAfee 2012; Liu et al. 2014), whereas in others it has been shown that such dominance does not necessarily hold (Moldovanu and Sela 2001; Archak and Sundararajan 2009; Kaplan et al. 2002).

[3] The first step uses the identity $\binom{k}{i+1} = \frac{k}{i+1}\binom{k-1}{i}$, and the last uses the binomial theorem.

Interestingly, in Ghosh and Kleinberg (2016) that use a similar model as ours, except that a contestant realizes the value of her contribution before making her participation decision, the optimal prize structure is awarding equal prizes to the top j contestants. This is very different from the winner-takes-all dominance we prove for our model, which assumes the value of contribution is realized only after making the participation decision.

Optimal Prize Structure

In order to determine the optimal prize budget M to be set and its division into prizes, the manager needs to solve a complex optimization problem which essentially attempts to maximize $B^{manager}$ according to (4) by controlling the parameters $\{M_1, ..., M_k\}$ (where in case it is optimal to use $n < k$ prizes we get $M_j = 0\ \forall j > n$). The complexity of the optimization derives from the fact that, as demonstrated in the analysis section, even the slightest change in the prize structure may result in substantial changes in the resulting equilibria and consequently in the participation scheme of the different agents, the expected performance and the expected prize awarded, all affecting the manager's expected profit.

Fortunately, in some cases, the multi-prize model enables reaching an equilibrium solution which is identical to the fully cooperative solution, in which case the determination of the optimal prize scheme (both M and its division into prizes) can be substantially simplified. The fully cooperative solution is the one where the manager can fully control which of the agents will take part in the contest, attempting to maximize the expected net benefit when taking into consideration the value distributions and costs. Given a subset $A' \subseteq A$ of agents taking part in the contest, the expected net profit, denoted $B^{cooperative}$, is $B^{cooperative} = \int\limits_{y=-\infty}^{\infty} y \frac{d(F^M(y))}{dy} dy - \sum_{A_i \in A} c_i$, where $F^M(x) = \prod\limits_{A_j \in A'} F_j(x)$.[4] The optimal cooperative solution is to have all agents in $A^* \subseteq A$ participate, where A^* is the subset for which $B^{cooperative}$ is maximized. This requires considering all subsets of A, which is combinatorial however of a lesser concern given the relatively moderate number of contestants in most crowdsourcing contest platforms.[5] The maximum expected profit in the fully cooperative case is an upper bound for the manager's expected profit in a contest, as any non-cooperative equilibrium solution is dominated (from the manager's point of view) by a similar cooperative solution in which the agents' profit is zero (i.e., the expected award to each participating agent merely covers its participation cost). Therefore if a prize scheme can constructed such that it results in the optimal fully cooperative solution, then it is necessarily the optimal one. A prize scheme $\{M_1, ..., M_{|A^*|}\}$ will result in an equilibrium solution identical to the fully cooperative one if: (a) for each agent $A_i \in A^*$ (i.e., participating according to the fully cooperative solution) $B_i(P) = 0$; and (b) for each agent $A_i \notin A^*$ $B_i(P) \leq 0$. In both cases $B_i(P)$ is calculated according to (2), substituting $p_j = 1$ in (1) for each $A_j \in A^*$.

Therefore, one should first solve the set of $|A^*|$ linear equations of type $B_i(P) = 0$ (each being an instance of (2), representing a different agent $A_i \in A^*$, where the prize allocation is the set of variables). If the solution is a set $\{M_1, ..., M_{|A^*|}\}$ such that $M_j \geq 0\ \forall j \leq |A^*|$ and satisfies $B_i(P) \leq 0$ for each $A_i \notin A^*$ then this is the optimal prize scheme for the multi-prize case and the expected profit is equal to the theoretical-optimal fully-cooperative case. Note that having an equilibrium solution of the latter type does not preclude the existence of additional equilibrium solutions with the same prize scheme, though as discussed above the determination of which of these will hold in such cases is beyond the scope of the paper.

Consider for example the setting used for Figure 3. Here, the optimal fully cooperative solution is to have both agents participate. The solution to the set of equations of type (a) above for this case yields a prize scheme $(0.10175, 0.28325)$, which is optimal for the contest, as it is equal to the optimal fully cooperative solution. In this example the equilibrium of having both agents participate with this prize scheme is unique so it is necessarily the one to be used. We emphasize that this solution could not have been achieved with a single prize—it is the division into several prizes that enables the necessary additional flexibility in the design of the mechanism.

Related Work

Contest design in general (Riley and Samuelson 1981; Myerson 1981; Morgan, Orzen, and Sefton 2012; Liu et al. 2013; Siegel 2009) and particularly crowdsourcing contest design (Glazer and Hassin 1988; Ghosh and McAfee 2012; DiPalantino and Vojnovic 2009; Archak and Sundararajan 2009) has attracted much attention in research over the past decades, mostly due to the applicability of the mechanism (Taylor 1995; Vojnovi 2016; Liu et al. 2014) and its effectiveness in eliciting effort (Ghosh and Kleinberg 2016; Glazer and Hassin 1988; Green and Stokey 1983). Along the rich theoretical work there is also much empirical work aiming to understand the mechanics and dynamics of contests by analyzing contest data from web-sites such as 99designs and Taskcn (de Araújo 2013; Liu et al. 2014). Common to all contest models, that contestants need to spend costly efforts (e.g., time, resources) in order to become eligible to win or to increase the chance of winning one or more prizes (Dechenaux, Kovenock, and Sheremeta 2014). In that sense, many have used all-pay auction models, which have similar characteristics, as a framework for analyzing contests (Chawla, Hartline, and Sivan 2012; Kaplan et al. 2002; Luo et al. 2016).

The model analyzed in this paper differs from most models used in contest literature primarily in the sense that it considers a "simple" contest, i.e., one where participants'

[4]A solution to the fully cooperative problem can be extracted by mapping the process into a cooperative exploration problem, e.g., as in (Rochlin, Sarne, and Mash 2014).

[5]The median number of users competing over a task is 2 in TopCoder and 4-15 (depending on the task type) in Taskcn (see Table 1.1 in Vojnovic 2016).

sole decision is whether or not to participate in the contest. This, as opposed to enabling participants full control over the amount of effort they exert which in turn influences their chance of winning in the contest (i.e., an "effort-based" contest) (Nti 1999; Moldovanu and Sela 2006; Cavallo and Jain 2013; Cohen, Kaplan, and Sela 2008; Liu et al. 2014). Studies that consider a model similar to ours include those of Ghosh et al (2016) and Levy et al (2017). Ghosh et al assume contestants learn about their performance measure in the contest prior to having to decide on participation, whereas in our model there is no certainty concerning performance at the time the participation decision is made. For that model, they found that whenever the agents are homogeneous and the manager's profit is taken to be the sum of performance obtained throughout the contest it is always optimal to award equal prizes to a subset of "best" contestants. For the heterogeneous case where the goal is to maximize the performance of the best-performer they find that the winner-takes-all scheme is not necessarily the optimal one. Levy et al (2017) use a model identical to ours, except that they only consider the winner-takes-all prize scheme. The locus of their work is a comparison of a parallel contest to a sequential one, pointing to the transition in preference between the two as a function of the setting parameter.

Prize splitting has been extensively studied for models of effort-based contests, yielding a plethora of results. For example Archak et al. (2009) found that when contestants are risk-neutral, the manager should optimally allocate all of its budget to the top prize even if it values multiple submissions. In contrast, if contestants are sufficiently risk-averse, the manager may optimally offer more prizes than the number of submissions it desires. Chawla et al (2012) found that winner-takes-all is the optimal choice, whenever prize division is determined prior to the contest itself, for a specific model where the manager only benefits from the highest submission and contestants' probability of winning the prize depends on the effort they exert. They also compare crowdsourcing contests with more conventional means of procurement such as simple highest-bid auctions. Luo et al (2015) studied a contribution-dependent prize function in Tullock contests (in which the winning probability is the ratio of the contestants effort and the total effort exerted by all contestants), deriving the optimal prize function that induces the maximum profit for the contest manager. Moldovanu et al (2001) studied the effect of contestants' cost function over the managers' profit, when the latter is taken to be the sum of efforts. Their main result is that in case of linear or convex cost functions a winner-takes-all contest is optimal. In case of convex functions they provide a necessary and sufficient condition for the optimality of a multi-prize scheme. DiPalantino et al (2009) studied the essential features of a crowdsourcing system, demonstrating the relationship between incentives and participation in such systems. They found that rewards yield logarithmically diminishing returns with respect to participation levels.

All the above, however, do not apply to our model due to the inherent difference in the modeling choice made concerning the ability of contestants to control the effort they exert and consequently their chances of winning the contest. Furthermore, none of the above work has demonstrated the preference of a multi-prize scheme where the first prize is smaller than other prizes offered.

Discussion, Conclusions and Future Research

Like with some of the contest models studied in prior work, we find that in the contest model analyzed in this paper switching from the winner-takes-all to a multi-prize scheme can be highly beneficial for the contest manager. In contrast to all prior work in this area, however, in our model it is possible that the optimal prize allocation results in awarding the contributor associated with the highest quality contribution (which accounts entirely to the manager's profit) a prize that is smaller than the prizes awarded to those with contributions of lesser quality (of which the manager has no benefit whatsoever), leaving the winner envy in those ranked below her.

For the homogeneous case we prove that the winner-takes-all scheme dominates the multi-prize one. The importance of this finding is twofold: First, it provides the manager with the preferred scheme for this important class of settings (which other than being quite common in contest literature, as discussed in the relevant section, is also quite common in real life). Second, it directly points to the sole source of improvement the multi-prize model achieves—it is the contestants' heterogeneity that accounts to the improvement in the expected profit, as with the additional flexibility enabled by the division into multiple prizes the manager can offer finer participation incentives. This is illustrated in the numerical examples provided along the analysis, pointing to several different sets of dynamics, of different natures, enabling the improvement achieved.

As demonstrated, in some cases the multi-prize contest scheme can even lead to the performance of the fully cooperative model, which is the upper bound for the manager's expected profit when running a contest. Both the identification of those settings and the extraction of the optimal prize allocation for them is simple and only requires solving a set of linear equations.

The analysis provided in the paper, much like most of the work on crowdsourcing contests (and on contests in general), is applicable to settings where both the agent initiating the contest (the manager) and the prospective participants are fully rational and aim to maximize their expected profit. Motivations for such theoretical work and arguments for its importance are abundant in prior work cited throughout the paper, hence we do not see a need to repeat those. Having said that, we do acknowledge the importance of empirical research in this application domain, and plan to study human behavior in the model analyzed under this framework. In that sense, the results reported in the paper provide the basis for rational decision making and a benchmark for comparison in any future empirical work aiming to study simple crowdsourcing multi-prize contests carried out with people. An additional direction for future research on the theoretical side is the analysis of social welfare and price of anarchy under a multi-prize mechanism compared to winner-takes-all.

Acknowledgments

This research was partially supported by the ISRAEL SCIENCE FOUNDATION (grants No. 1162/17) and the ISF-NSFC joint research program (grant No. 2240/15).

References

Archak, N., and Sundararajan, A. 2009. Optimal design of crowdsourcing contests. In *Proceedings of the International Conference on Information Systems*, 200.

Cavallo, R., and Jain, S. 2013. Winner-take-all crowdsourcing contests with stochastic production. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing*.

Chawla, S.; Hartline, J. D.; and Sivan, B. 2012. Optimal crowdsourcing contests. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, 856–868.

Cohen, C.; Kaplan, T.; and Sela, A. 2008. Optimal rewards in contests. *RAND Journal of Economics* 39(2):434–451.

de Araújo, R. M. 2013. 99designs: An analysis of creative competition in crowdsourced design. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing*.

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

Difallah, D. E.; Catasta, M.; Demartini, G.; and Cudré-Mauroux, P. 2014. Scaling-up the crowd: Micro-task pricing schemes for worker retention and latency improvement. In *Proceedings of the Second AAAI Conference on Human Computation and Crowdsourcing*.

DiPalantino, D., and Vojnovic, M. 2009. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, 119–128.

Elmalech, A.; Sarne, D.; David, E.; and Hajaj, C. 2016. Extending workers attention span through dummy events. In *Proceedings of the Fourth AAAI Conference on Human Computation and Crowdsourcing*.

Faradani, S.; Hartmann, B.; and Ipeirotis, P. G. 2011. What's the right price? pricing tasks for finishing on time. In *Papers from the 2011 AAAI Human Computation Workshop*.

Ghosh, A., and Hummel, P. 2015. Cardinal contests. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, 377–387.

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

Ghosh, A., and McAfee, P. 2012. Crowdsourcing with endogenous entry. In *Proceedings of the 21st International Conference on World Wide Web*, WWW '12, 999–1008.

Glazer, A., and Hassin, R. 1988. Optimal contests. *Economic Inquiry* 26(1):133–143.

Green, J., and Stokey, N. 1983. A comparison of tournaments and contracts. *Journal of Political Economy* 91(3):349–64.

Kaplan, T.; Luski, I.; Sela, A.; and Wettstein, D. 2002. All-pay auctions with variable rewards. *The Journal of Industrial Economics* 50(4):417–430.

Koutsopoulos, I. 2013. Optimal incentive-driven design of participatory sensing systems. In *Proceedings of the IEEE INFOCOM*, 1402–1410.

Levy, P.; Sarne, D.; and Rochlin, I. 2017. Contest design with uncertain performance and costly participation. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. (to appear).

Liu, B.; Lu, J.; Wang, R.; and Zhang, J. 2013. Prize and punishment: Optimal contest design with incomplete information. *National University of Signapore Typescript*.

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.

Luo, T.; Das, S. K.; Tan, H. P.; and Xia, L. 2016. Incentive mechanism design for crowdsourcing: An all-pay auction approach. *ACM Trans. Intell. Syst. Technol.* 7(3):35:1–35:26.

Luo, T.; Tan, H.; and Xia, L. 2014. Profit-maximizing incentive for participatory sensing. In *2014 IEEE Conference on Computer Communications (INFOCOM)*, 127–135.

Moldovanu, B., and Sela, A. 2001. The optimal allocation of prizes in contests. *American Economic Review* 91(3):542–558.

Moldovanu, B., and Sela, A. 2006. Contest architecture. *Journal of Economic Theory* 126(1):70–96.

Morgan, J.; Orzen, H.; and Sefton, M. 2012. Endogenous entry in contests. *Economic Theory* 51(2):435.

Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.

Nti, K. O. 1999. Rent-seeking with asymmetric valuations. *Public Choice* 98(3):415–430.

Riley, J. G., and Samuelson, W. F. 1981. Optimal auctions. *The American Economic Review* 71(3):381–392.

Rochlin, I.; Sarne, D.; and Mash, M. 2014. Joint search with self-interested agents and the failure of cooperation enhancers. *Artificial Intelligence* 214:45–65.

Siegel, R. 2009. All-pay contests. *Econometrica* 77(1):71–92.

Taylor, C. 1995. Digging for golden carrots: An analysis of research tournaments. *American Economic Review* 85(4):872–90.

Vojnovi, M. 2016. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press.