

# Contests to Incentivize a Target Group\*

Edith Elkind, Abheek Ghosh and Paul W. Goldberg

Department of Computer Science, University of Oxford

{edith.elkind,abheek.ghosh,paul.goldberg}@cs.ox.ac.uk

## Abstract

We study how to incentivize agents in a target sub-population to produce a higher output by means of rank-order allocation contests, in the context of incomplete information. We describe a symmetric Bayes–Nash equilibrium for contests that have two types of rank-based prizes: (1) prizes that are accessible only to the agents in the target group; (2) prizes that are accessible to everyone. We also specialize this equilibrium characterization to two important sub-cases: (i) contests that do not discriminate while awarding the prizes, i.e., only have prizes that are accessible to everyone; (ii) contests that have prize quotas for the groups, and each group can compete only for prizes in their share. For these models, we also study the properties of the contest that maximizes the expected total output by the agents in the target group.

## 1 Introduction

Contests—situations where multiple agents compete for valuable prizes—are prevalent in many real-life situations, including sports, R&D races, crowdsourcing, college admissions and job recruitment.<sup>1</sup>

The dominant paradigm in contest theory is to assume that the principal designs the prize structure so as to incentivize the agents to produce a higher total output. In this model, the designer values the contribution from each agent equally, i.e., the marginal output produced by any agent contributes the same marginal value to the designer’s objective. In this paper, we break away from this assumption: instead, we assume that the agents belong to two different groups and the designer may want to prioritize one group—the target group—over the other, non-target group.

Contests that give special importance to agents from a particular group are not uncommon in practice. Conferences give best student paper awards, in addition to best paper awards.

\*Full version of the paper is available at <http://arxiv.org/abs/2204.14051>.

<sup>1</sup>See the book by Vojnovic [2016] for an introduction to contest theory and relevant resources.

Many competitions and hackathons are organized to elicit engagement from underrepresented groups: for example, recently Microsoft organized a hackathon as an intervention to get women interested in AI.<sup>2</sup> Contests are widely used for crowdsourcing, which is also an important source of training data for machine learning algorithms; in this context, eliciting input from disadvantaged groups is particularly important, as it helps the algorithms to learn decision-making rules that reflect opinions and preferences of such groups. Our work provides a better understanding of how to encourage contributors from such groups.

In this paper, we study an incomplete information (Bayesian) model. We assume that each agent is associated with an *ability*, which captures the amount of output they can produce per unit effort. In an incomplete information model for contests, it is generally assumed that the abilities of the agents are selected i.i.d. from a given distribution  $F$ . In our model, based on the assumption that the agents belong to one of the two groups, the abilities of the agents in the target and the non-target group are sampled from two distinct distributions  $F$  and  $G$ , respectively. Each agent belongs to the target group with a fixed probability  $\mu$ .<sup>3</sup> The agents know the prize allocation scheme, their own ability, and the prior distributions of other agents’ abilities. They act strategically to maximize their expected utility, where the utility of an agent is the prize they receive minus the effort they exert, and reach a Bayes–Nash equilibrium. The contest designer knows the prior distributions of the agents’ abilities, and therefore can reason about the equilibrium behavior of the agents. She designs the prize allocation scheme so as to elicit equilibrium behavior that optimizes her own objective. We assume that the contest designer has a budget of 1 to use for the prizes.

The contest ranks the agents based on their outputs and awards the prizes based on the ranks. In our general model, we assume that contests have two sequences of prizes, one available to all agents, and the other available to the target group only. A conference sponsoring a best student paper

<sup>2</sup><https://news.microsoft.com/europe/features/bridging-the-stem-divide-ai-hackathon-helps-young-women-excel-in-computer-science/>

<sup>3</sup>We assume that  $\mu$  is known based on historical data on the number of participants from the target group compared to the total number of participants. Similarly, the distributions  $F$  and  $G$  are also known from historical data.

award in addition to a best paper award is an example of such a contest. We also study two specific variants of this model. In the first variant, all prizes in the contest are equally accessible to all agents, and there are no group-specific prizes. One could come across such a design preference in situations where the contest designer does not want to discriminate among the agents, but still wants to incentivize an under-represented agent group. In this case, the contest does not discriminate while awarding the prizes, but the design of the prize structure (i.e., the value of the first prize, second prize, and so on) can incorporate the information about the ability distribution of the agents. This model captures an equal opportunity employer that treats all applicants equally, but has a preference to increase the diversity in the workplace. In the second variant, the contest has group-specific prizes only, with no prizes available to the entire population; examples include a hackathon for women in CS, or a hiring/scholarship contest with a strict group-specific quota.

We assume that the contest designer wants to maximize the expected output of the agents in the target group only. With minor changes, the techniques can be extended to an objective that maximizes the weighted sum of the group outputs and other similar objectives.

## 1.1 Our Contribution

We initiate a theoretical study of contest design to elicit higher participation by agents from under-represented groups, and propose tractable models for doing so. We analyze the equilibrium behavior of the agents for three cases: *general prizes only*, where there is a sequence of rank-ordered prizes accessible to all agents, whether in the target or the non-target group; *group-specific prizes only*, where there are two different sequences of prizes, each accessible to one of the groups; *both general and target group-specific prizes*, where there are prizes accessible to both groups as well as prizes accessible to the target group only. The techniques used to analyze the equilibrium in these contests can be extended to related prize structures.

Based on these equilibrium characterizations, we study the properties of the contest that maximizes the total expected output of the agents from the target group. For the setting where all prizes are general, we prove that the optimal contest awards a prize of  $1/\ell$  to the first  $\ell$  agents and 0 to the remaining agents, where  $\ell$  is a function of the relative frequency of target group agents,  $\mu$ , and the ability distributions of the target group,  $F$ , and the non-target group,  $G$ . We give a closed-form formula for  $\ell$ . We also study the effect of first-order and second-order stochastic dominance of  $F$  on  $G$ , and vice versa, on the optimal contest. For the setting where all prizes are group-specific, we show that it is optimal to award the full prize budget to the top-ranked agent, irrespective of the distributions  $F$  and  $G$ . This result matches similar results on maximizing total output [Moldovanu and Sela, 2001]. We also compare general prizes and group-specific prizes, and show that either of these choices may be preferred depending upon the distributions  $F$  and  $G$ , and also upon the frequency of the target group agents in the population,  $\mu$ , even if  $F = G$ .

All proofs and a few additional examples are given in the full version of the paper.

## 1.2 Related Work

We discuss literature in contest theory directly relevant to our work; we point the readers to the book by Vojnovic [2016] for a broader survey on contest theory.

Moldovanu and Sela [2001] characterize the Bayes–Nash equilibrium for contests with rank-order allocation of prizes assuming incomplete information with i.i.d. types; Chawla and Hartline [2013] prove the uniqueness of this equilibrium. Moldovanu and Sela [2001] also show that awarding the entire prize to the top-ranked agent is optimal when the agents have (weakly) concave cost functions, but the optimal mechanism may have multiple prizes for convex cost functions. In contrast, we observe that even for linear cost functions, it may be beneficial to have multiple prizes to maximize the expected output from the target group.

The standard assumption in mechanism design and its sub-area of contest design is that the agents’ types are sampled i.i.d. from some distribution. Our paper extends the equilibrium analysis of Moldovanu and Sela [2001] to an asymmetric model where the agents are from two groups with different distributions. Amann and Leininger [1996] study an asymmetric model with two agents with types sampled independently from two different distributions; our analysis uses the idea of the function ‘ $k$ ’ introduced by them, see Theorem 1. Characterizing equilibrium in contests with many agents with types sampled independently from different distributions has remained technically challenging; recently Olszewski and Siegel [2016; 2020] made progress by assuming very large numbers of agents and prizes, where an individual agent has infinitesimal effect on the equilibrium.

Bodoh-Creed and Hickman [2018] model affirmative action in college admissions using contests. Their model has general non-linear utility and cost functions, but, like Olszewski and Siegel [2016; 2020], they assume that the numbers of agents and prizes are very large. They give first-order equilibrium conditions and show that two types of affirmative actions—(i) admissions preference schemes, where the outputs of agents from the target group are (artificially) amplified before ranking and prize allocation, and (ii) quotas, where there are separate pools of seats for the different groups (similar to our group-specific prizes only model, Section 3.3)—have the same sets of equilibria with identical actions by the agents and identical outcomes. Our model makes a stronger linearity assumption regarding the utility and the cost functions, which allows us to better characterize the equilibrium and study the optimal contest (prize) design problem.

There is a long line of literature on the effect of different types of affirmative action on contests. This literature generally assumes that there is only one prize, which either gets allocated to the top-ranked agent or gets allocated proportionally (Tullock [1980] contests and their generalizations), and the contest designer introduces different kinds of interventions in the allocation of this prize. One such intervention is *favoritism*, where the objective is to maximize the total output by giving head-starts or handicaps to certain agents (see, e.g., [Kirkegaard, 2012; Fu and Wu, 2020; Deng *et al.*, 2021]). A head-start adds a bonus to an agent’s output, while a handicap decreases an agent’s score by a fixed percentage. See [Chowdhury *et al.*, 2020] for a survey on

contests and affirmative action.

Our work deviates from the widely studied objective of maximizing total output, by focusing on the output of a target group. Several papers have studied objectives other than the total output, such as maximum individual output [Chawla *et al.*, 2019], cumulative output from the top  $k$  agents [Archak and Sundararajan, 2009; Gavious and Minchuk, 2014], total output of agents producing output in a given range [Elkind *et al.*, 2021]. Elkind *et al.* [2021] target agents based on their output level, and therefore, effectively, based on their abilities; our work targets agents based on their group association, with the assumption that exogenous factors may have caused agents in some groups to have acquired less ability.

There are also several papers that perform equilibrium analysis and optimal contest design in the complete information setting (e.g., [Baye *et al.*, 1996; Barut and Kovenock, 1998; Siegel, 2009]).

## 2 Model and Preliminaries

There are  $n$  agents. Any given agent belongs to the target (resp., non-target) group with probability  $\mu$  (resp.,  $(1 - \mu)$ ), independently of the other agents. Let  $\mathbf{v} = (v_1, v_2, \dots, v_n)$  be the ability profile of the agents, where  $v_i$  values are independent and identically distributed (i.i.d.) random variables from continuous and differentiable distributions  $F$  or  $G$  with support  $[0, 1]$ , depending upon whether agent  $i$  belongs to the target or the non-target group, respectively. Let  $f$  and  $g$  be the probability density functions (PDFs) of  $F$  and  $G$ . The  $n$  agents simultaneously produce the output profile  $\mathbf{b} = (b_i)_{i \in [n]}$ , so that the cost of agent  $i$  is  $b_i/v_i$ .

The contest awards two sequences of prizes  $w_1 \geq w_2 \geq \dots \geq w_n$  and  $\omega_1 \geq \omega_2 \geq \dots \geq \omega_n$ , where the prize  $w_j$  is given to the  $j$ -th ranked agent overall (we call these prizes *general prizes*), and the prize  $\omega_j$  is given to the  $j$ -th ranked agent among the agents in the target group (we call these *group-specific prizes*), with ties broken uniformly.<sup>4</sup> We assume that  $\sum_j w_j + \sum_j \omega_j \leq 1$ , i.e., the contest designer has a unit budget. Given a vector of outputs  $\mathbf{b}$ , let the allocation of general prizes be given by  $\mathbf{x}(\mathbf{b}) = (x_{i,j}(\mathbf{b}))_{i,j \in [n]}$ , where  $x_{i,j} = 1$  if agent  $i$  is awarded the  $j$ -th prize and  $x_{i,j} = 0$  otherwise. Similarly, let  $\mathbf{y}(\mathbf{b}, \mathbf{t}) = (y_{i,j}(\mathbf{b}, \mathbf{t}))_{i,j \in [n]}$  be the allocation of target group prizes, where  $\mathbf{t} = (t_i)_{i \in [n]}$  is the group label vector;  $t_i \in \{T, N\}$ ,  $t_i = T$  if the agent is in the target group,  $t_i = N$  if not. We shall suppress the notation for  $\mathbf{t}$  and write  $\mathbf{y}(\mathbf{b})$  instead of  $\mathbf{y}(\mathbf{b}, \mathbf{t})$ .

If agent  $i$  is in the target group, her utility is given by

$$u_T(v_i, \mathbf{b}, \mathbf{t}) = \sum_{j \in [n]} w_j x_{i,j}(\mathbf{b}) + \sum_{j \in [n]} \omega_j y_{i,j}(\mathbf{b}) - b_i/v_i$$

<sup>4</sup>In our model, it does not matter how we break ties, so w.l.o.g., we can assume uniform tie breaking. In more detail, for a tie to happen, two players must produce exactly identical output. In our setting, this happens with zero probability, as there are no point masses in the probability distributions  $F$  and  $G$ , and the distributions of the output generated by the agents (that we derive as a result of our equilibrium analysis) also do not have any point masses.

$$\equiv v_i \left( \sum_{j \in [n]} w_j x_{i,j}(\mathbf{b}) + \sum_{j \in [n]} \omega_j y_{i,j}(\mathbf{b}) \right) - b_i, \quad (1)$$

where the equivalence is true because scaling the utility function by a constant does not change the strategy of an agent. Similarly, the utility of an agent  $i$  in the non-target group is given by

$$u_N(v_i, \mathbf{b}, \mathbf{t}) = v_i \sum_{j \in [n]} w_j x_{i,j}(\mathbf{b}) - b_i. \quad (2)$$

### 2.1 Mathematical Preliminaries

Let  $p_j^H(v)$  denote the probability that a value  $v \in [0, 1]$  is the  $j$ -th highest among  $n$  i.i.d. samples from a distribution  $H$ , given by the expression

$$p_j^H(v) = \binom{n-1}{j-1} H(v)^{n-j} (1-H(v))^{j-1}.$$

A key role in the Bayes–Nash equilibrium of rank-order allocation contests is played by the order statistics. Let  $f_{n,j}^H$  be the PDF of the  $j$ -th highest order statistic out of  $n$  i.i.d. samples from  $H$  (with PDF  $h$ ), given by the expression

$$f_{n,j}^H(v) = \frac{n!}{(j-1)!(n-j)!} H(v)^{n-j} (1-H(v))^{j-1} h(v). \quad (3)$$

We shall also frequently use the following two identities for order statistics:

$$\begin{aligned} H(v) f_{n-1,j}^H(v) &= \frac{n-j}{n} f_{n,j}^H(v); \\ (1-H(v)) f_{n-1,j}^H(v) &= \frac{j}{n} f_{n,j+1}^H(v). \end{aligned}$$

**Definition 1** ((Weak) First-Order Stochastic (FOS) Dominance). *A distribution  $F$  FOS dominates another distribution  $G$ , if for every  $x$ ,  $F$  gives at least as high a probability of receiving at least  $x$  as does  $G$ :  $1 - F(x) \geq 1 - G(x) \iff F(x) \leq G(x)$ .*

**Definition 2** ((Weak) Second-Order Stochastic (SOS) Dominance). *A distribution  $F$  SOS dominates another distribution  $G$ , if for every  $x$ :  $\int_{-\infty}^x [F(t) - G(t)] dt \leq 0$ .*

FOS dominance implies SOS dominance. Another sufficient condition for SOS dominance:  $F$  SOS dominates  $G$  if  $G$  is a mean-preserving spread of  $F$ .

Let us denote the inner product of two functions  $\psi, \phi : [0, 1] \rightarrow \mathbb{R}_+$  by

$$\langle \psi, \phi \rangle = \int_0^1 \psi(x) \phi(x) dx.$$

When appropriately normalized,  $\langle \psi, \phi \rangle$  measures the similarity between the functions  $\psi$  and  $\phi$ .

### 2.2 Equilibrium and Objective Function

We shall study the Bayes–Nash equilibrium of the agents, where the agents in the target group use a symmetric strategy  $\alpha(v)$  (symmetric across the agents in the group), where  $\alpha(v)$

is the output of an agent with ability  $v$ ; similarly, the agents in the non-target group use a symmetric strategy  $\beta(v)$ . The objective of the contest designer is to maximize the expected total output generated by the agents in the target group:

$$\sum_{i \in [n]} \mathbb{E}_{v_i \sim F}[\alpha(v_i)] \mathbb{P}[t_i = T] = n \cdot \mu \cdot \mathbb{E}_{v_i \sim F}[\alpha(v_i)],$$

which is equivalent to maximizing  $\mathbb{E}_{v \sim F}[\alpha(v)]$  because  $n$  and  $\mu$  are (fixed) parameters given to the model.

### 3 Equilibrium Analysis

In this section, we characterize a symmetric Bayes–Nash equilibrium (symmetric across the agents in a given group). We first solve for the equilibrium of the contest with both general and target group-specific prizes, then we specialize it for the two special cases: (1) only general prizes; (2) only group-specific prizes.

#### 3.1 Both Target Group and General Prizes

In the next theorem, we focus on contests that have a sequence of general prizes  $(w_j)_{j \in [n]}$  and a sequence of target group-specific prizes  $(\omega_j)_{j \in [n]}$ . We will characterize a Bayes–Nash equilibrium where a player in the target group plays an action  $\alpha(v)$  as a function of her ability  $v \sim F$  and a player in the non-target group plays an action  $\beta(v)$  as a function of her ability  $v \sim G$ . Note that all the players in the target group use the same strategy  $\alpha$  and the players in the non-target group use the same strategy  $\beta$ .

**Theorem 1.** *A contest with general prizes  $(w_j)_{j \in [n]}$  and target group-specific prizes  $(\omega_j)_{j \in [n]}$  has a Bayes–Nash equilibrium where an agent with ability  $v \in [0, 1]$  uses strategy  $\alpha(v)$  if she is in the target group and strategy  $\beta(v)$  if she is in the non-target group, where  $\alpha(v)$  and  $\beta(v)$  are defined as:*

$$\beta(v) = \int_0^v (\mu f(k(y))k'(y) + (1 - \mu)g(y))A(y)dy,$$

$$\alpha(v) = \begin{cases} \beta(k^{-1}(v)), & 0 \leq v \leq k(1) \\ \beta(1) + \int_{k(1)}^v \mu f(y)C(y)dy, & k(1) \leq v \leq 1 \end{cases},$$

where:

- $k$  is a function defined over  $[0, 1]$  and is the solution to the following ordinary differential equation, with boundary condition  $k(0) = 0$ ,

$$k'(v) = \frac{(1 - \mu)g(v)(v - k(v))A(v)}{\mu f(k(v))(k(v)A(v) + k(v)B(v) - vA(v))}.$$

- $A(v) = \sum_{j \in [n]} w_j \psi_j(\mu F(k(v)) + (1 - \mu)G(v))$ .
- $B(v) = \sum_{j \in [n]} \omega_j \psi_j(\mu F(k(v)) + (1 - \mu))$ .
- $C(v) = \sum_{j \in [n]} (w_j + \omega_j) \psi_j(\mu F(v) + (1 - \mu))$ .
- $\psi_j(x) = \binom{n-1}{j-1} x^{n-j-1} (1-x)^{j-2} ((n-j) - (n-1)x)$ .

Notice that Theorem 1 does not provide a closed-form solution for  $\alpha(v)$  and  $\beta(v)$ , it rather provides a characterization using the function  $k(v)$  that relates  $\alpha(v)$  and  $\beta(v)$ . Standard equilibrium analysis techniques using first-order equilibrium

conditions will give a system of differential equations in  $\alpha(v)$  and  $\beta(v)$ . Using the function  $k(v) = \alpha^{-1}(\beta(v))$ , we convert this system of differential equations in  $\alpha(v)$  and  $\beta(v)$  to a comparatively simpler single first-order explicit ordinary differential equation in only one dependent variable  $k(v)$ . After solving the ordinary differential equation for  $k(v)$ , we can use  $k(v)$  to compute  $\alpha(v)$  and  $\beta(v)$  as described in the theorem statement. This technique of using the function  $k(v)$  is motivated by the work of Amann and Leininger [1996], where they use it to analyze two-agent asymmetric contests.

If the distributions  $F$  or  $G$ , or the prizes  $(w_j)_{j \in [n]}$  or  $(\omega_j)_{j \in [n]}$ , are non-trivial, then it may be intractable to calculate the analytical solutions for  $\alpha(v)$  and  $\beta(v)$  derived in Theorem 1, but numerical solutions may be computed. An example in the full version of the paper illustrates the use of Theorem 1 to compute  $\alpha(v)$  and  $\beta(v)$  for a particular instantiation of the problem.

Theorem 1 characterizes an equilibrium assuming that all players in a given group use the same strategy ( $\alpha(v)$  or  $\beta(v)$ ), and that  $\alpha(v)$  and  $\beta(v)$  are strictly increasing and (almost everywhere) differentiable functions of  $v$ . We believe that these assumptions are reasonable and the equilibrium is the most natural one for this model. It remains open to prove the uniqueness of this equilibrium or to characterize the set of all possible equilibria.

#### 3.2 General Prizes Only

A contest designer may want to only award prizes that are equally accessible to all the agents. This model is particularly appealing for several real-life applications, because, although the design of the prize structure may incorporate distributional information about the ability of the agents in the two groups, the contest itself is completely unbiased.

In our model, this design choice corresponds to setting  $\omega_j = 0$  for all  $j \in [n]$ . For this case, the expected utility for a target group agent and a non-target group agent for a given pair of ability  $v$  and output  $b$  is the same. We can specialize the equilibrium characterization in Theorem 1 to this case as follows:

**Theorem 2.** *A contest with only general prizes  $(w_j)_{j \in [n]}$  has a Bayes–Nash equilibrium where an agent with ability  $v \in [0, 1]$  in the target or non-target group uses the strategy  $\alpha(v)$ , where  $\alpha(v)$  is defined as:*

$$\alpha(v) = \sum_{j \in [n-1]} (w_j - w_{j+1}) \int_0^v y f_{n-1,j}^{\mu F + (1-\mu)G}(y) dy.$$

This equilibrium is unique, as follows from a result by Chawla and Hartline [2013]. They prove that if a contest is anonymous (does not discriminate among the agents) and the abilities of the agents are sampled i.i.d. from a distribution, then the equilibrium is unique. With only general prizes, the contest is (i) anonymous because the general prizes are awarded without any bias towards players from any group, and (ii) effectively, the abilities of the agents are sampled i.i.d. from the distribution  $\mu F(v) + (1 - \mu)G(v)$ .

#### 3.3 Group-Specific Prizes Only

In this case, there are no general prizes accessible to agents from both groups: all the prizes are allocated based on strict

quotas for the groups. This model is similar to the model of seat quotas for college admissions studied by Bodoh-Creed and Hickman [2018]. Technically, in this case we assume that  $w_j = 0$  for all  $j \in [n]$ . Here we focus on the target group; a similar result also holds for the non-target group, if there are any prizes reserved for them.

**Theorem 3.** *A contest with target only group-specific prizes,  $(\omega_j)_{j \in [n]}$ , has a Bayes-Nash equilibrium where an agent with ability  $v \in [0, 1]$  in the target group uses strategy  $\alpha(v)$ , where  $\alpha(v)$  is defined as:*

$$\alpha(v) = \sum_{j \in [n-1]} (\omega_j - \omega_{j+1}) \int_0^v y f_{n-1,j}^{\mu F + (1-\mu)G}(y) dy.$$

Chawla and Hartline [2013]’s uniqueness result also applies to this equilibrium characterization, as it did for only general prizes (Section 3.2). We can transform an instance of our problem (without changing its set of equilibria) to satisfy these requirements: the abilities of the players are picked i.i.d. from the distribution  $\mu F(v) + (1 - \mu)$  and the contest awards prizes without discriminating among the players.

## 4 Designing Prizes to Incentivize the Target Group

Based on the equilibrium characterizations in the previous section, in this section we investigate the properties of the optimal contest that maximizes the total output of the target group. We characterize the optimal contests for only general prizes and for only group-specific prizes. Characterization of the optimal contest for both general and group-specific prizes is a non-trivial open problem because of the absence of a closed-form equilibrium characterization for the case (the characterization in Theorem 1 involves an ordinary differential equation).

In the full version of the paper, we compare the choice between only general and only group-specific prizes; we observe that either of them may be a better choice depending upon the situation. Even if the distributions of the target and the non-target group are the same, the relative frequency of target group agents, measured by  $\mu$ , determines which prize allocation scheme is better. If  $\mu$  is sufficiently high, then it is better to have only group-specific prizes, while the converse is true if  $\mu$  is very low.

### 4.1 General Prizes Only

The strategy used by the agents in the target group is  $\alpha(v)$ , as derived in Theorem 2. It can be rewritten as

$$\begin{aligned} \alpha(v) &= \sum_{j \in [n-1]} (\omega_j - \omega_{j+1}) \int_0^v y f_{n-1,j}^{\mu F + (1-\mu)G}(y) dy \\ &= \sum_{j \in [n-1]} \gamma_j \frac{1}{j} \int_0^v y f_{n-1,j}^{\mu F + (1-\mu)G}(y) dy, \end{aligned} \quad (4)$$

where  $\gamma_j = j(\omega_j - \omega_{j+1})$ . Instead of optimizing over  $(\omega_j)_{j \in [n]}$ , we can optimize over  $(\gamma_j)_{j \in [n-1]}$  with constraints  $\gamma_j \geq 0$  and  $\sum_j \gamma_j = 1$ , to find the optimal contest that maximizes the target group’s output.

**Theorem 4.** *The contest with only general prizes that maximizes the expected total output of the agents in the target group awards a prize of  $1/k^*$  to the  $k^*$  top-ranked agents and a prize of 0 to the remaining  $n - k^*$  agents, where  $k^*$  is defined as*

$$k^* \in \arg \max_{j \in [n-1]} \langle \psi, f_{n,j+1}^{\mu F + (1-\mu)G} \rangle$$

$$\text{for } \psi(x) = \frac{x(1-F(x))}{1-\mu F(x)-(1-\mu)G(x)}.$$

Let  $H(x)$  denote  $\mu F(x) + (1 - \mu)G(x)$ . The function  $\psi(x) = x \frac{1-F(x)}{1-H(x)}$  in Theorem 4 is independent of the prize structure: it depends only on  $\mu$ ,  $F$ , and  $G$ , which are parameters given to the contest designer. The function  $f_{n,j+1}^H(x)$  is the PDF of the  $(j + 1)$ -th highest order statistic of the distribution  $H$ . The optimal  $j = k^*$  maximizes the inner product between  $\psi$  and  $f_{n,j+1}^H$ , and therefore, the similarity between them. In other words, the optimal  $j = k^*$  selects the order statistic that is as similar to  $\psi$  as possible.

Note that when we are maximizing the total output, then instead of an inner product of the order statistic with  $\psi(x)$  in Theorem 4, we take an inner product with just  $x$ . As  $x$  is monotonically increasing, the order statistic that is most similar to  $x$  and that maximizes the inner product with  $x$  is the one with  $j = 1$ , as shown by Moldovanu and Sela [2001]. This provides an argument in favor of allocating the entire prize budget to the first prize. In contrast, when we focus on the target group only, it may be optimal to distribute the prize among several top-ranked agents. This also means that we may lose a portion of the total output, and this loss can be  $\Omega(n)$  as shown in Example 1. Also, by the *single-crossing* property, see Vojnovic [Vojnović, 2016] Chapter 3, if we flatten the prize structure then we increase the output of the agents with low ability and decrease the output of the agents with high ability.

We have partially omitted calculations from the examples in this section. These examples are presented in the full version of the paper with more details.

**Example 1.** *Let  $F(x) = 1 - (1 - x)^{n-1}$ ,  $G(x) = ((n - 1)x - F(x))/(n - 2)$ , and  $\mu = 1/(n - 1)$ . We get  $H(x) = \mu F(x) + (1 - \mu)G(x) = x$ , which is the uniform distribution over  $[0, 1]$ . Note that  $\psi(x) = \frac{x(1-F(x))}{1-H(x)} = x(1 - x)^{n-2}$ .*

*It can be checked that the order-statistic that maximizes the inner-product with  $\psi(x)$  is  $f_{n,n-1}^H(x) = n(n - 1)x(1 - x)^{n-1}$ . So, the optimal value is  $k^* = n - 2$ . The total output generated for  $k^* = n - 2$  is  $\frac{2}{n(n+1)}$ . On the other hand, we know that the expected total output is maximized by  $j = 1$  and is equal to  $\frac{(n-1)}{n(n+1)}$ . The ratio between the two quantities is  $(n - 1)/2 = \Omega(n)$ .*

In the above example, the distribution of the target group  $F$  was first-order stochastic (FOS) dominated by the distribution of the general population  $H$ . For this case, we observed that awarding prizes to more than one agent increases the expected output of the target group.  $F$  being FOS dominated by  $H$  means that the agents in the target group have lower ability than the general population. As flattening the prize structure increases the output of the lower ability agents and decreases

the output of the higher ability agents, it makes sense that having a flatter prize structure incentivizes them.

The above result raises the question: What if the target group is equally able or stronger than the general population, i.e., if the distribution of the ability of an agent in the target group,  $F$ , FOS dominates the distribution of the ability of an agent in the non-target group,  $G$ ; is it then optimal to award the prize to the top-ranked agent only? When  $F = G = H$ , we know that it is optimal to give prize to the top-ranked agent only, so when  $F$  FOS dominates  $H$ , one might expect this to be the case as well. However, the example below illustrates that this is not true.<sup>5</sup>

**Example 2.** Let  $\mu = 1/8$ ,  $F(x) = 1 - S(x)$ ,  $G(x) = (x - \mu F(x))/(1 - \mu)$ , where  $S(x)$  is given by:

$$S(x) = \begin{cases} 1, & \text{if } x < 3/4 \\ 16(1-x)^2, & \text{if } 3/4 \leq x < 15/16 \\ 1-x, & \text{if } x \geq 15/16 \end{cases}.$$

It can be verified that  $F(x)$  and  $G(x)$  are continuous cumulative distribution functions. The distribution of the general population is  $H(x) = \mu F(x) + (1 - \mu)G(x) = x$ . Also,  $F(x) \leq H(x)$  for every  $x$  (and strict for some values of  $x$ ), and therefore,  $F$  FOS dominates  $H$ . The optimal number of prizes  $k^*$  need not be 1; for example, for  $n = 50$ ,  $k^* = 11$ .

Now we consider the case where  $F$  and  $H$  have the same mean but different variance, which implies second-order stochastic (SOS) dominance. The distribution with lower variance SOS dominates the one with higher variance, assuming both distributions have the same mean. The results are similar to FOS dominance: it may be optimal to have multiple prizes irrespective of whether  $F$  or  $H$  has a higher variance, as shown in the following examples. Note that FOS dominance implies SOS dominance, but not vice-versa.

**Example 3.** Let  $\mu = 2/3$ ,  $F(x) = 3x^2 - 2x^3$ , and  $G(x) = 3x - 6x^2 + 4x^3$ . Observe that  $H(x) = \mu F(x) + (1 - \mu)G(x) = x$ , and the variance of  $H$  is  $1/12$  while the variance of  $F$  is  $1/20$ , and both have mean  $1/2$ .

Solving for  $k^*$ , we get  $k^* = \frac{5n}{6} - \frac{\sqrt{7n^2+30n+39}}{6} + \frac{1}{2}$ . For example, for  $n = 50$ ,  $k^* = 19$ .

**Example 4.** Let  $\mu = 1/4$ ,  $F(x) = 1 - S(x)$ , and  $G(x) = (x - \mu F(x))/(1 - \mu)$ , where  $S(x)$  is given by:

$$S(x) = \begin{cases} 1 - 48x/31, & \text{if } x < 31/96 \\ 1/2, & \text{if } 31/96 \leq x < 3/4 \\ 8(1-x)^2, & \text{if } 3/4 \leq x < 7/8 \\ 1-x, & \text{if } x \geq 7/8 \end{cases}$$

It can be checked that  $F(x)$  and  $G(x)$  are valid and for the distribution of the general population we have  $H(x) =$

<sup>5</sup>In the preliminaries we assumed  $F$  to be continuous, strictly increasing, and differentiable in  $[0, 1]$ . In the following and the subsequent examples, the distribution  $F$  is continuous and weakly increasing but may not be strictly increasing and differentiable everywhere. If  $F$  is not strictly increasing, we can add a slight gradient to resolve the issue; on the other hand, we can smooth out the finite number of points where it is not differentiable. These changes will have a minimal effect on the objective value and our analysis holds. For ease of presentation, we shall not discuss these issues.

$\mu F(x) + (1 - \mu)G(x) = x$ . It can also be checked that the mean of all the distributions is  $1/2$ , and the variance of  $H$  is  $1/12 \approx 0.083$  while the variance of  $F$  is  $6703/55296 \approx 0.121 > 1/12$ . The optimal number of prizes  $k^*$  need not be 1; for example, for  $n = 50$ ,  $k^* = 11$ .

## 4.2 Group-Specific Prizes Only

In this section, we study the optimal prize structure when the prizes are accessible to the target group only. This models situations when there are separate contests for the target and the non-target group (there can be a separate contest for only non-target group agents, this will not have any affect on the strategy of the target group agents). We see that the optimal contest allocates the entire prize budget to the first prize, i.e., gives a prize of 1 to the top-ranked agent in the group (Theorem 5). The proof of Theorem 5 extends the techniques used to provide a similar result for total output [Moldovanu and Sela, 2001].

**Theorem 5.** The contest with only group-specific prizes that maximizes the output of the target group awards a prize of 1 to the top-ranked agent and 0 to others.

## 5 Conclusion

Our paper studies rank-order allocation of prizes that aim to maximize the output of a target group. We provide equilibrium characterization for these contests and study the properties of the optimal contest. For unbiased contests (i.e., with general prizes only), although it is preferable to award the prize to the top-performing player if the target group players are *similar* to the overall population, if the target group players are very *different*, then we may want to flatten out the prize structure. This does not only mean that we flatten out the prizes if the target group players are *weaker* than others; we also demonstrate that if the target group players are *stronger* but quite different from others, it may still be a good idea to flatten out the prizes. However, for contests with fixed prize quotas (i.e., with group-specific prizes only), it may not be useful to flatten the prizes. Comparing unbiased contests and fixed quota contests, the size of the target group in proportion to the size of the population plays an important role.

An important open problem is to understand the properties of the optimal contest that has both group-specific and general prizes. We characterize the equilibrium for this problem and analyze the optimal contest for specific sub-cases, but the general problem is open. One way to make this tractable may be to assume large numbers of agents and prizes; such assumptions allow for a stronger and more tractable equilibrium characterization by making the influence of an individual agent infinitesimal [Olszewski and Siegel, 2016]. Our work focused on contests with a rank-order allocation of prizes with incomplete information and linear cost functions. Instead of rank-order allocation of prizes, we can study a proportional allocation of prizes or its generalizations (see, e.g., [Tullock, 1980]); one motivation for proportional allocation is the randomness and unpredictability of outcome in the real world. Similarly, we may relax the assumption of linear cost functions, or we can study complete information models instead of incomplete information.

## Ethical Statement

There may be both short-term and long-term implications for designing contests specifically to incentivize a target group. For example, even if the contest itself is unbiased (Section 4.1), optimizing the expected total output of the players in the target group may mean that we decrease the expected total output of the entire population, and in particular decrease the expected total output of the non-target group. Moreover, such effects can be stronger if the target group forms a smaller fraction of the participants. Our analysis and examples try to capture such effects. For any given contest, this means that we are optimizing for the target group with a cost to society. But for a longer time frame, the answer is less obvious, and we believe that it is out of the scope of our paper. This question concerns whether practices such as affirmative action have an overall positive or negative impact on the society in the long run. Our paper addresses the following question: if we do want to use affirmative action in contests to incentivize the target group (with particular design choices, such as the prize allocation scheme being unbiased), how do we do it most effectively, and how does it affect the output of all the players in the given contest.

## Acknowledgements

We would like to thank the reviewers for their valuable feedback. The second author is supported by Clarendon Fund and SKP Scholarship.

## References

- [Amann and Leininger, 1996] Erwin Amann and Wolfgang Leininger. Asymmetric all-pay auctions with incomplete information: the two-player case. *Games and Economic Behavior*, 14(1):1–18, 1996.
- [Archak and Sundararajan, 2009] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the 13th International Conference on Information Systems*, 2009.
- [Barut and Kovenock, 1998] Yasar Barut and Dan Kovenock. The symmetric multiple prize all-pay auction with complete information. *European Journal of Political Economy*, 14(4):627–644, 1998.
- [Baye *et al.*, 1996] Michael R Baye, Dan Kovenock, and Casper G De Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, 1996.
- [Bodoh-Creed and Hickman, 2018] Aaron L Bodoh-Creed and Brent R Hickman. College assignment as a large contest. *Journal of Economic Theory*, 175:88–126, 2018.
- [Chawla and Hartline, 2013] Shuchi Chawla and Jason D Hartline. Auctions with unique equilibria. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, pages 181–196, 2013.
- [Chawla *et al.*, 2019] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. *Games and Economic Behavior*, 113:80–96, 2019.
- [Chowdhury *et al.*, 2020] Subhasish M Chowdhury, Patricia Esteve-González, and Anwesha Mukherjee. Heterogeneity, leveling the playing field, and affirmative action in contests. SSRN:3655727, 2020.
- [Deng *et al.*, 2021] Shanglyu Deng, Qiang Fu, and Zenan Wu. Optimally biased Tullock contests. *Journal of Mathematical Economics*, 92:10–21, 2021.
- [Elkind *et al.*, 2021] Edith Elkind, Abheek Ghosh, and Paul Goldberg. Contest design with threshold objectives. In *Proceedings of the 17th Conference on Web and Internet Economics*, page 554, 2021. ArXiv preprint 2109.03179.
- [Fu and Wu, 2020] Qiang Fu and Zenan Wu. On the optimal design of biased contests. *Theoretical Economics*, 15(4):1435–1470, 2020.
- [Gavious and Minchuk, 2014] Arie Gavious and Yizhaq Minchuk. Revenue in contests with many participants. *Operations Research Letters*, 42(2):119–122, 2014.
- [Kirkegaard, 2012] René Kirkegaard. Favoritism in asymmetric contests: Head starts and handicaps. *Games and Economic Behavior*, 76(1):226–248, 2012.
- [Moldovanu and Sela, 2001] Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001.
- [Olszewski and Siegel, 2016] Wojciech Olszewski and Ron Siegel. Large contests. *Econometrica*, 84(2):835–854, 2016.
- [Olszewski and Siegel, 2020] Wojciech Olszewski and Ron Siegel. Performance-maximizing large contests. *Theoretical Economics*, 15(1):57–88, 2020.
- [Siegel, 2009] Ron Siegel. All-pay contests. *Econometrica*, 77(1):71–92, 2009.
- [Tullock, 1980] Gordon Tullock. Efficient rent-seeking. In James Buchanan, Robert Tollison, and Gordon Tullock, editors, *Toward a Theory of the Rent Seeking Society*, pages 131–146. Texas A & M University, College Station, Texas, USA, 1980.
- [Vojnović, 2016] Milan Vojnović. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press, Cambridge, UK, 2016.