Cardinal Contests

Arpita Ghosh
Cornell University
Ithaca, NY, USA
arpitaghosh@cornell.edu

Patrick Hummel
Google Inc.
Mountain View, CA, USA
phummel@google.com

ABSTRACT

Contests are widely used as a means for effort elicitation in settings ranging from government R&D contests to on-line crowdsourcing contests on platforms such as Kaggle, Innocentive, or TopCoder. Such rank-order mechanisms—where agents’ rewards depend only on the relative ranking of their submissions’ qualities—are natural mechanisms for incentivizing effort when it is easier to obtain ordinal, rather than cardinal, information about agents’ outputs, or where absolute measures of quality are unverifiable. An increasing number of online contests, however, rank entries according to some numerical evaluation of their absolute quality—for instance, the performance of an algorithm on a test dataset, or the performance of an intervention in a randomized trial. Can the contest designer incentivize higher effort by making the rewards in an ordinal rank-order mechanism contingent on such *cardinal* information?

We model and analyze cardinal contests, where a principal running a rank-order tournament has access to an absolute measure of the qualities of agents’ submissions in addition to their relative rankings, and ask how modifying the rank-order tournament to incorporate cardinal information can improve incentives for effort. Our main result is that a simple threshold mechanism—a mechanism that awards the prize for a rank if and only if the absolute quality of the agent at that rank exceeds a certain threshold—is *optimal* amongst all mixed cardinal-ordinal mechanisms where the fraction of the j^{th} prize awarded to the j^{th} -ranked agent is any arbitrary non-decreasing function of her submission’s quality. Further, the optimal threshold mechanism uses exactly the same threshold for each rank. We study what contest parameters determine the extent of the benefit from incorporating such cardinal information into an ordinal rank-order contest, and investigate the extent of improvement in equilibrium effort via numerical simulations.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social & Behavioral Sciences.

General Terms

Economics.

Keywords

Contests; Optimal contest design; Crowdsourcing; Game theory.

1. INTRODUCTION

Contests have a long history as a means for procuring innovations, with government-sponsored contests for research and development dating back to at least 1714.¹ Contests provide an effective incentive structure for eliciting effort in settings where the quality of an agent’s output as well as her effort are unverifiable or difficult to measure, making conventional contracts based on input or output-dependent rewards infeasible—either by virtue of being too costly for the sponsor to implement, or because they cannot be credibly enforced due to unverifiability of output ([8], [41]). In several such situations an *ordinal* comparison—identifying a relative ranking of agents’ submissions—might nonetheless be feasible, allowing the principal to commit to an enforceable contract that awards rank-based prizes to some subset of entrants in a contest. This has led to a large literature on the optimal design of rank-order mechanisms for effort elicitation; see §1.1.

In contrast with more traditional settings, however, an increasing number of online contests procure innovations whose quality is evaluated and ranked via verifiable *cardinal* measurements. For instance, the well-known Netflix contest, designed to procure improved algorithms for movie recommendations, evaluated entries according to how well they predicted user preferences on a test subset of its user database, with the algorithm that obtained the highest score being declared the winner. In a different context, contests for designing mobile apps for health or education might compute scores for submissions based on their performance on metrics of efficacy in randomized trials², and rank entries

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author’s site if the Material is used in electronic media.
WWW 2015, May 18–22, 2015, Florence, Italy.
ACM 978-1-4503-3469-3/15/05.
<http://dx.doi.org/10.1145/2736277.2741652>.

¹The British parliament ran a contest with a 20000 pound prize for a method for determining longitude-at-sea to within half a degree.

²See, for one example, <http://www.robinhood.org/prize>

based on these scores. A third example is the contest platform Kaggle, which hosts contests where the innovation being procured is a data-mining algorithm: again, submissions are typically evaluated on a test dataset provided by the requesters and ranked according to the score they obtain. A number of other such contests abound. In summary, there is an increasingly large family of contests where submissions can be assigned a meaningful numerical quality score that reflects their value to the principal.

Suppose a principal running a contest has access to such *cardinal* measurements of the qualities of agents' submissions, in addition to the ordinal ranking of their outputs. What effect does making prizes contingent on absolute performance, in addition to relative ranking, have on contestants' incentives for effort, and what is the optimal way to incorporate such cardinal measurements of output into an ordinal rank-order contest?

Our contributions. We model and analyze *cardinal contests*, where a principal, who wants to maximize some increasing function of the quality of received submissions, can evaluate the quality of each submission via a real-valued score in addition to observing the rank-ordering of contestants' outputs. We ask whether and how such a principal can improve her utility by incorporating such cardinal information to determine agents' rewards.

Specifically, consider a rank-order tournament with prizes $(A_1, A_2, \dots, A_n)$ for ranks $1, \dots, n$, in a model where contestants are strategic and have a cost to effort. Can the principal modify the rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$, using participants' quality scores q_j in addition to their ranks j to decide how much of the prize A_j to actually award, to improve incentives for effort? And if so, how should the function $g_j(q_j)$ that determines the fraction of the prize A_j to be awarded to the j^{th} -ranked contestant based on her output quality q_j (such as fractions that grow linearly or as a convex function of quality) be chosen to create the optimal incentives for the agents?

Our main result is that a very simple threshold mechanism—a mechanism that awards the prize for a rank if and only if the output of the agent at that rank exceeds a certain threshold—is *optimal* amongst all mechanisms induced by any increasing functions $g_j(\cdot)$: We show that for any given rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$, the functions $g_j(q_j)$ that incentivize the highest effort are precisely step functions that increase from 0 to 1 at a threshold score (Theorem 3.1). Additionally, if the noise density that stochastically maps an agent's effort into her output quality is single-peaked at 0, we find that the optimal threshold mechanism also uses exactly the *same threshold* for each rank. Interestingly, real-world contests where prizes are made contingent on cardinal information measuring absolute performance seem to commonly use such threshold-based prize structures.³

After deriving comparative statics for the optimal threshold and equilibrium effort in such mechanisms in §3.2, we ask to what extent does incorporating cardinal information (via an optimal threshold mechanism) improve equilibrium effort relative to using only ordinal information. We address this question via numerical simulations as well as theory, and find that the extent to which cardinal information incen-

tivizes higher effort depends on the contest parameters: the most substantial improvements are achieved in contests with a smaller number of participants, contests with participants whose cost functions are 'not too convex' (in a sense made precise in §4), and contests that award a larger rather than smaller number of prizes.

We note that we ask the question of how to optimally *modify* a given rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$, rather than the question of choosing the overall optimal reward structure $(A_1(q_1, \dots, q_n), \dots, A_n(q_1, \dots, q_n))$ that incorporates all available cardinal information in determining the reward for each rank. There are a number of reasons we address this question rather than the more general mechanism design problem, the foremost of which is practical: a principal announcing a contest might choose, or be committed to, a certain rank-based prize structure for reasons such as sponsorship constraints, publicity, simplicity, cost of precise evaluations of a full rank-order due to scale, and so on.⁴ However, the principal might still want, and more easily be able to, incorporate an entry's absolute quality in determining whether, and how much of the announced rank-based prize to actually award (for instance, she may wish to award no prize if the highest-ranked submission performs worse than the current state-of-the-art innovation). See §2 and §5 for further discussion.

1.1 Related Work

There is a large body of work on contest design in the economics literature. In addition to foundational work on the theory of contests ([18], [24], [26], [27], [28], [29], [33], [34], [35], [36], [37], [38], [39], [40]), there has also been a variety of work motivated by specific applications. [17], [21], [23], and [31] address the design of rank-order tournaments for the purpose of incentivizing employees to work hard. [8] and [41] study contest design in the context of research tournaments. And there is a growing literature motivated by online crowdsourcing contests ([1], [6], [7], [11], [16]). There is also an extensive empirical literature analyzing observed strategic behavior by real subjects in contests in a variety of settings ([2], [3], [4], [5], [9], [10], [12], [13], [22], [25], [30]).

This work addresses a variety of questions related to the economics of contests such as comparing tournaments to incentive schemes that are based solely on an individual's personal performance ([18], [21], [31], [34]), relationships between contests and all-pay auctions ([7], [8], [11]), taxing entry to improve the quality of contributions ([16], [41]), dynamic contests in which agents dynamically decide how much effort to exert when they continuously get information about how they are doing in the contest compared to their competitors ([14], [20]), and incentives for agents to work hard in teams [19]. The most relevant subset of this literature to our paper is that relating to optimal contest design ([1], [17], [24], [26], [27], [28], [29], [40]), which asks how to best choose the rank-based rewards for each rank under various models of effort and constraints on the rewards.

The key difference between this literature and our work is that this literature almost exclusively studies contests that are structured as *rank-order* mechanisms, where the

³For example, one of the largest contests hosted on Kaggle awarded prizes to the top three entries, provided their scores were above a minimum baseline.

⁴Similar contests on the same platform (e.g. Kaggle) with similar total prize money do sometimes differ in how they split the total prize across ranks, suggesting that the choice of the rank-order mechanism might perhaps not be driven only by optimality of effort incentives.

announced prizes depend only on the *rank* of an agent's output relative to that of her competitors, whereas we consider mixed cardinal-ordinal mechanisms of the form described in §2. Specifically, rather than ask which ordinal rewards $(A_1, A_2, \dots, A_n)$ incentivize optimal outcomes, we ask how cardinal information about an agent's output can be optimally incorporated into a *given* ordinal reward mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$ to incentivize the highest effort, a question that has not been addressed previously in this literature. The only relevant exception, to the best of our knowledge, is the work in [7] which does address the question of optimal contest design using both cardinal and ordinal information, albeit in a completely different model of output and for risk-neutral agents only. Interestingly, although the models are completely different, the optimal mechanism in that model also turns out to use cardinal information via a threshold.

In addition to this key difference, we also note that we do not make any of a variety of simplifying assumptions used in several papers in this literature, such as risk-neutrality of the agents, linearity of agents' cost functions, the assumption that there is no stochastic component that affects the final quality of an agent's contribution, or assuming a particular functional form for the principal's objective function.

2. MODEL

We consider a setting where a pool of agents competing in a contest strategically choose their effort levels which stochastically determines their submission qualities, and each agent receives a prize based on the relative rank of her output as well as possibly its absolute quality. The model is a natural extension of [23].

Agents. There are n agents who compete in a contest. Each agent simultaneously chooses a level of effort $e_i \geq 0$ to exert. The quality q_i of agent i 's output is determined both by her effort, and a random noise term ϵ_i as $q_i = e_i + \epsilon_i$, where each ϵ_i is an independent and identically distributed draw from a cumulative distribution function $F(\epsilon)$. The noise ϵ_i has a number of possible interpretations. Most simply, ϵ_i could model the fact that a given amount of effort (for example, the time spent working on a research problem) does not deterministically guarantee a certain level of output, but rather only influences its expected value. As another interpretation, the noise ϵ could also model randomness in the measurement, or perception, of an agent's output quality by the principal. Most interestingly, ϵ_i could be thought of as modeling heterogeneity amongst agents' abilities to solve the specific problem or execute the specific task required for the contest; we discuss this in detail at the end of this section. We will assume throughout that the probability density function $f(\epsilon)$ corresponding to F is a bounded continuously differentiable function of ϵ with a bounded first derivative, and the distribution F and the number of agents n are known to all agents.

Utilities. An agent's utility is the difference between her benefit from any prize she wins and her cost of effort. An agent who receives a prize A_i derives a benefit of $v(A_i)$, where $v(\cdot)$ is a strictly increasing and concave function satisfying $v(A) \geq 0$ for all $A \geq 0$ and $v(0) = 0$. Exerting effort e_i incurs a cost $c(e_i)$, where $c(\cdot)$ is a strictly increasing and convex function satisfying $c(0) = 0$ and $c'(0) = 0$.

The utility to an agent who exerts effort e_i and receives prize A_i is the difference between this benefit and cost, $u_i = v(A_i) - c(e_i)$. We assume that each agent chooses

e_i to maximize her expected utility $E[v(A_i)] - c(e_i)$, where the expectation is over the n random draws of $\epsilon_{i'}$ that determine each agent's output quality (and therefore their prizes). Our model allows for both risk-neutral and risk-averse agents, as $v(A)$ may either be a linear function of A (for risk-neutral agents) or a strictly concave function of A (risk-averse agents).

Mechanisms. We suppose that the principal running the contest can observe the quality q_i of each agent's output. (Our results extend immediately to a model where the principal's observations of output qualities are noisy, as long as the noise in each observation is an IID draw for each agent). We use $\mathcal{M}(A_1, A_2, \dots, A_n)$ to denote a rank-order mechanism which assigns a reward A_j to the agent with the j^{th} -highest output regardless of its absolute quality, and assume throughout that $A_1 \geq \dots \geq A_n \geq 0$.⁵

Mixed cardinal-ordinal mechanisms. Let q_j denote the quality of the j^{th} -ranked submission. We consider mixed cardinal-ordinal modifications of a rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$ of the form $\mathcal{M}(g_1(q_1)A_1, \dots, g_n(q_n)A_n)$, which awards the agent with the j^{th} -ranked submission of quality q_j a prize $P_j = g_j(q_j)A_j$, where $g_j(q)$ is a non-decreasing function of q satisfying $0 \leq g_j(q) \leq 1$. That is, $g_j(q)$ represents the fraction of the 'maximum' prize A_j for achieving rank j that an agent obtains if she produces a contribution with absolute quality q . Note that the rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$ corresponds to setting g_j to be the constant function $g_j(q) = 1$ for all j .

We note that a principal with access to cardinal measurements of the qualities of each submission could conceivably use more general mechanisms by allowing the function g_j to depend not only on the quality of the corresponding j^{th} -ranked submission, but rather on the entire vector of qualities $(q_1, \dots, q_n)$. We restrict ourselves to mechanisms that use functions $g_j(q_j)$ for simplicity, both of analysis and implementation. In addition to leading to simpler mechanisms, a principal might, practically speaking, prefer to announce a contest where the prize awarded to a winner is contingent on the absolute quality *only* of her own submission, and not on the absolute qualities of the submissions produced by her competitors; see also the discussion in §1.

Principal's objective. We assume that the principal's objective is to maximize some utility function $W(q_1, \dots, q_n)$ of agents' output qualities in equilibrium, where W is non-decreasing in its arguments, *i.e.*, $W(q_1, \dots, q_n)$ is such that if $q'_i \geq q_i$ for all i , then $W(q'_1, \dots, q'_n) \geq W(q_1, \dots, q_n)$. We will be interested throughout in symmetric pure-strategy Nash equilibria. If agents all use the same level of effort in equilibrium, then the principal's expected utility is non-decreasing in the effort choice of the agents: therefore, all such increasing objectives are simultaneously improved by mechanisms which elicit higher equilibrium effort from agents, assuming a symmetric equilibrium.

Heterogeneity. Our model, with noise terms ϵ_i all drawn from the same distribution $F(\cdot)$ and costs to effort given by the same function $c_i(\cdot) = c(\cdot)$, suggests contributors who are *a priori* homogeneous, corresponding to a pool of contestants who all have similar skills or ability for the subject of the contest (for example, programmers with similar levels of expertise or graphic designers with similar skill levels).

⁵It is worth noting, however, that most of the proofs and results in this paper will continue to hold even if this assumption on the monotonicity of rewards is relaxed.

Such a model captures scenarios where it is predominantly differences in effort, rather than differences in ability, that dominate differences in the quality produced. It also captures scenarios where the set of potential contestants may be self-selected to have rather similar abilities or expertise levels, and therefore similar costs to producing a particular quality.

Our model nonetheless allows capturing agent heterogeneity in two different ways. First, two contestants with the same effort choice will still come up with solutions of different qualities depending on their draws of ϵ_i , corresponding to situations where agents with similar skills nonetheless produce different outputs for the *specific* task posed by a particular contest. Second, the incentives of agents in contests where there is indeed heterogeneity in abilities that affects agents' final output qualities, but where agents do not know these abilities (beyond their distribution) prior to making their strategic effort choices, are identical to those in our model with 'abilities' drawn from $F(\cdot)$ after the agents choose their effort levels. Therefore, this model of heterogeneity includes many situations with non-homogeneous contributors, *as long as* agents do not learn their abilities prior to choosing their effort levels.

In a different model of heterogeneity amongst agents that appears in the contest design literature, an agent's output q_i is her ability-scaled effort $a_i e_i$, where agents' abilities are all randomly drawn from a single distribution F : a logarithmic transformation of variables from $q_i = a_i e_i$ in those models yields exactly our model where $q_i = e_i + \epsilon_i$. The key difference between these two models of heterogeneity is *timing of information*—agents observe their random draws of a_i *before* making their strategic effort choices in those models, whereas agents do *not* observe their draws of ϵ_i prior to making their effort choices in our model (for instance, as with a graphic designer who does not know exactly how good a design she will produce until she attempts it).

3. OPTIMAL CARDINAL MODIFICATIONS OF ORDINAL MECHANISMS

A contestant's cardinal quality score can be incorporated into a given rank-order mechanism in many different ways. A mechanism might choose to increase the reward for each rank linearly with an agent's output; more generally, the reward might increase as some convex function of her cardinal quality score up to some maximum reward. The reward scheme could also vary discontinuously with the quality of the agent's output, for instance so that the reward for a particular rank is determined by which of several intervals the corresponding quality score lies in. *A priori*, each of these cardinal modifications to a rank-order mechanism might create stronger incentives for effort than a purely ordinal mechanism by making reward more strongly dependent on the absolute quality. What choice of functions $g_j(\cdot)$ creates the strongest incentives for effort amongst all possible non-decreasing functions g_j ?

While analyzing this question, we make two mild simplifying assumptions throughout. First, we assume that there is a symmetric pure strategy equilibrium in which all agents exert the same effort level e in the game; we show in the full version of the paper [15] that such symmetric pure strategy equilibria will exist under the fairly standard assumption that a player's cost function is sufficiently convex. Second,

we consider functions $g_j(q)$ for which there is some small $\delta_j > 0$ such that $g_j(q)$ may only assume values that are integral multiples of δ_j . While this assumption is purely for technical simplicity and our results continue to hold without this assumption (albeit with a more complex proof), we note that this assumption is realistic because in any practical application there will be some minimum unit of a currency that represents the smallest possible amount by which one can change the value of an agent's prize. (For example, if prizes were paid in US dollars, any prize would necessarily have to be some integral multiple of some small fraction of a penny, as a principal would not be able to divide an agent's prize further than this).

Our main result, Theorem 3.1, shows that the question of how to choose the functions $g_j(q)$ to optimally modify any given rank-order mechanism using cardinal information has a strikingly simple answer: For any given rank-order mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$, no other functions $g_j(q)$ can incentivize higher equilibrium effort than the optimal step functions that increase from 0 to 1 at some threshold score. We refer to such mechanisms as *threshold mechanisms*.

THEOREM 3.1. *Consider a rank-order tournament in which the agent who finishes in j^{th} place is awarded a prize $g_j(q)A_j$, where $g_j(q)$ is a non-decreasing function satisfying $0 \leq g_j(q) \leq 1$ for all q . There exist functions $g_j(q)$ of the form $g_j(q) = 0$ for $q < q_j^*$ and $g_j(q) = 1$ for $q \geq q_j^*$ for some constants q_j^* that incentivize the highest equilibrium effort amongst all possible mechanisms characterized by some functions $g_j(q)$.*

PROOF. Consider a set of *ladder functions* $l_j(q)$ that are characterized by a set of m_j cutoffs $q_{1,j}^* < q_{2,j}^* < \dots < q_{m_j,j}^*$ for each of the j prizes such that $g_j(q) = r_{0,j}$ for $q < q_{1,j}^*$, $g_j(q) = r_{m,j}$ for all $q \in [q_{m,j}^*, q_{m+1,j}^*)$, and $g_j(q) = r_{m_j,j}$ for $q \geq q_{m_j,j}^*$, where $0 \leq r_{0,j} < r_{1,j} < \dots < r_{m_j,j} \leq 1$. Note that any non-decreasing function $g_j(q)$ such that $g_j(q)$ is an integral multiple of δ_j and $0 \leq g_j(q) \leq 1$ for all q can be written in this form. Thus in order to prove that threshold mechanisms are optimal amongst the set of all mechanisms, it suffices to show that threshold mechanisms are also optimal amongst the set of all mechanisms induced by ladder functions.

Let $y_j(e, e_i, \epsilon_i)$ denote the probability that agent i finishes in j^{th} place for a given realization of ϵ_i given that agent i exerts effort e_i and all other agents exert effort e . If the prize for being ranked in the j^{th} place with a contribution of quality q is $l_j(q)A_j$ for some ladder function $l_j(q)$, then an agent i 's expected utility from exerting effort e_i when all other agents are exerting effort e is

$$\begin{aligned} E[u_i] &= \sum_{j=1}^n \int_{-\infty}^{\infty} v(g_j(e_i + \epsilon_i)A_j) y_j(e, e_i, \epsilon_i) f(\epsilon_i) d\epsilon_i - c(e_i) \\ &= \sum_{j=1}^n \sum_{k=0}^{m_j} \int_{q_{k,j}^* - e_i}^{q_{k+1,j}^* - e_i} v(r_{k,j}A_j) y_j(e, e_i, \epsilon_i) f(\epsilon_i) d\epsilon_i - c(e_i) \end{aligned}$$

where we abuse notation by letting $q_{0,j}^* \equiv -\infty$ and $q_{m_j+1,j}^* \equiv \infty$. In a symmetric pure-strategy equilibrium, all agents, and specifically agent i , choose effort e , and this is a best response, *i.e.*, the derivative of u_i with respect to e_i must be zero at $e_i = e$. That is, the equilibrium effort e must satisfy

the first order conditions given below:

$$\begin{aligned} c'(e) = & \sum_{j=1}^n \sum_{k=0}^{m_j} v(r_{k,j}A) [y_j(e, e, q_{k,j}^* - e) f(q_{k,j}^* - e) \\ & - y_j(e, e, q_{k+1,j}^* - e) f(q_{k+1,j}^* - e) \\ & + \int_{q_{k,j}^* - e}^{q_{k+1,j}^* - e} \frac{\partial y_j(e, e_i, \epsilon_i)}{\partial e_i} \Big|_{e_i=e} f(\epsilon_i) d\epsilon_i]. \quad (1) \end{aligned}$$

But note that the right-hand side of this equation is a linear function of $v(r_{k,j}A)$ for all k and j . From this it follows that for all $k \leq m_j$, the right-hand side of this equation is either non-decreasing in $r_{k,j}$ or non-increasing in $r_{k,j}$. Thus if $m_j \geq 2$, then one can instead set the value of $r_{1,j}$ to either be equal to $r_{0,j}$ or $r_{2,j}$ without decreasing the right-hand side of equation (1), meaning that one can make this change without decreasing equilibrium level of effort.

But setting the value of $r_{1,j}$ to equal to $r_{0,j}$ or $r_{2,j}$ would be equivalent to replacing the ladder function $l_j(q)$ with a ladder function that has $m_j - 1$ points of discontinuity rather than m_j points of discontinuity. From this it follows that if one is using a mechanism based on ladder functions $l_j(q)$ such that some $l_j(q)$ has $m_j \geq 2$ points of discontinuity, then the mechanism designer can induce at least as large a level of effort by instead using some mechanism based on ladder functions such that $l_j(q)$ has $m_j - 1$ points of discontinuity. By induction, it then follows that the mechanism designer can also induce at least as large a level of effort by instead using some mechanism based on ladder functions such that each $l_j(q)$ has no more than one point of discontinuity.

To complete the proof, it suffices to show that if the mechanism designer is using a mechanism based on ladder functions $l_j(q)$ that each have no more than one point of discontinuity, then these single-step ladder functions must correspond to threshold mechanisms, *i.e.*, that $r_{0,j} \in \{0, 1\}$ and $r_{1,j} \in \{0, 1\}$. Since equilibrium effort is again given by the solution to equation (1), and the right-hand side of this equation is a linear function of $v(r_{k,j}A)$ for all k and j , we again have that the right-hand side of this equation is either non-decreasing in $r_{0,j}$ or non-increasing in $r_{0,j}$. Thus if $l_j(q)$ has exactly one point of discontinuity, then one can set the value of $r_{0,j}$ to either be equal to 0 or $r_{1,j}$ without decreasing the right-hand side of equation (1), meaning that one can make this change without decreasing equilibrium effort.

Now if $r_{0,j} = 0$, then the same argument illustrates that one can set the value of $r_{1,j}$ to either be equal to 0 or 1 without decreasing the right-hand side of equation (1), meaning that one can make this change without decreasing equilibrium effort. And if $r_{0,j} = r_{1,j}$, then $l_j(q)$ has no points of discontinuity, and the same argument again illustrates that one can set the value of $r_{0,j} = r_{1,j}$ to be either 0 or 1 without decreasing the right-hand side of equation (1), meaning that one can make this change without decreasing equilibrium effort. Thus the mechanism designer can induce the agents to exert at least as much effort in equilibrium by using a threshold mechanism in which $r_{0,j} \in \{0, 1\}$ and $r_{1,j} \in \{0, 1\}$. $\square$

While there are potentially a variety of far more nuanced ways to incorporate cardinal information in determining agents' prizes in a mixed cardinal-ordinal mechanism, Theorem 3.1 says that there is always an optimal mechanism with an exceedingly simple form—it awards the entire value of the j^{th} prize to the agent who finishes in j^{th} place if this agent's output quality meets some threshold, and awards her no

prize at all otherwise. This result may be surprising, since it would seem far more natural to use a mechanism where an agent's prize varies smoothly with the quality of her output than one where it exhibits a sharp discontinuity with respect to quality, especially if agents are risk-averse—risk-averse agents are likely to prefer a prize structure in which they have a good chance of receiving a moderate prize over one in which they have a high chance of receiving nothing and a high chance of receiving a large prize. Nonetheless, Theorem 3.1 shows that such threshold mechanisms can always incentivize agents to choose optimal effort levels.⁶

3.1 Optimal Thresholds

Our main result shows that the optimal mechanism is a threshold mechanism. We now ask what the optimal thresholds are. To do this, we will need an equilibrium analysis of threshold mechanisms.

Since we are interested in symmetric equilibria in which all players exert effort e , we begin by computing a player's expected utility from a given threshold mechanism when the player exerts effort e_i and all other players exert effort e . Consider a given threshold mechanism $\mathcal{M}(A_1, t_1, A_2, t_2, \dots, A_n, t_n)$, where the j^{th} -ranked agent receives the prize A_j if and only if her output quality exceeds the threshold t_j . Agent i 's expected utility in $\mathcal{M}(A_1, t_1, A_2, t_2, \dots, A_n, t_n)$ from choosing effort e_i , when all other agents are exerting effort e , is the difference between i 's expected prize minus the cost of her effort $c(e_i)$, where i 's prize depends both on the rank of the quality of her output $q_i = e_i + \epsilon_i$ relative to the qualities of the other agents' outputs $q_{i'} = e_{i'} + \epsilon_{i'}$, and whether her output q_i exceeds the threshold t_j specified for her rank j . We thus begin by considering the probability that agent i 's output has rank j when she chooses effort e_i and all other agents choose effort e .

If $\epsilon_{[j]}$ denotes the j^{th} -largest of the noise terms ϵ drawn by the $n - 1$ remaining agents, then agent i will have the j^{th} -highest quality output if and only if $e + \epsilon_{[j]} \leq e_i + \epsilon_i \leq e + \epsilon_{[j-1]}$. The probability of this event, for a given draw of ϵ_i , is the probability that the number $e_i + \epsilon_i - e$ lies between the random values of $\epsilon_{[j]}$ and $\epsilon_{[j-1]}$. Thus the probability of this event is equal to the probability that exactly $j - 1$ of $n - 1$ randomly drawn values from the cumulative distribution function $F(\cdot)$ exceed $e_i + \epsilon_i - e$. Using standard expressions for binomial probabilities, we know that this event takes place with probability

$$\binom{n-1}{j-1} (1 - F(e_i - e + \epsilon_i))^{j-1} F(e_i - e + \epsilon_i)^{n-j}. \quad (2)$$

Now suppose agent i produces output $q_i = e_i + \epsilon_i$. To actually receive the prize A_j in the threshold mechanism $\mathcal{M}(A_1, t_1, A_2, t_2, \dots, A_n, t_n)$, i 's output must also exceed the threshold t_j , *i.e.*, satisfy $e_i + \epsilon_i \geq t_j$. Therefore, the probability that agent i finally receives the j^{th} prize, unconditional on the realization of ϵ_i , is the integral of the probability in equation (2) over $\epsilon_i \geq t_j - e_i$, *i.e.*,

⁶It is worth noting that this result will also hold even if there is a common shock to the agents' output in the sense that $q_i = e_i + \epsilon_i + \eta$ for some randomly drawn value of η that is common to all agents. A substantively identical proof can be used to prove Theorem 3.1 in this slightly more general model.

$$\int_{t_j - e_i}^{\infty} \binom{n-1}{j-1} (1 - F(e_i - e + \epsilon_i))^{j-1} F(e_i - e + \epsilon_i)^{n-j} f(\epsilon_i) d\epsilon_i.$$

Thus an agent's total expected utility from exerting effort e_i , when other agents choose e , is the sum of her expected benefit over all ranks j minus her cost $c(e_i)$ or

$$E[u_i] = \sum_{j=1}^k v(A_j) \int_{t_j - e_i}^{\infty} \binom{n-1}{j-1} (1 - F(e_i - e + \epsilon_i))^{j-1} \times F(e_i - e + \epsilon_i)^{n-j} f(\epsilon_i) d\epsilon_i - c(e_i).$$

We now address the question of what the optimal thresholds are. The proof of optimality of threshold mechanisms does not say anything about how the thresholds corresponding to each rank j vary with j , and thus allows for the possibility that the thresholds that are optimal for each rank might be substantially different for each rank. However, while it might seem intuitive that the optimal threshold could either consistently increase or consistently decrease across ranks, this turns out not to be the case, as we show below:

THEOREM 3.2. *Suppose the noise density $f(\cdot)$ is single-peaked at 0. Then the optimal threshold mechanism $\mathcal{M}(A_1, t_1^*, \dots, t_n^*)$ applies the same threshold to each rank j for any monotone mechanism $\mathcal{M}(A_1, A_2, \dots, A_n)$, i.e., $t_j^* = t^*$ for $j = 1, \dots, n$.*

PROOF. If the mechanism designer uses a threshold mechanism in which the agent who finishes in j^{th} place receives a prize if and only if this agent's observed quality exceeds t_j , then an agent i 's expected utility from exerting effort e_i when all other agents are exerting effort e is given by

$$E[u_i] = \sum_{j=1}^n \int_{t_j - e_i}^{\infty} v(A_j) y_j(e, e_i, \epsilon_i) f(\epsilon_i) d\epsilon_i - c(e_i)$$

where we let $y_j(e, e_i, \epsilon_i) \equiv \binom{n-1}{j-1} (1 - F(e_i - e + \epsilon_i))^{j-1} F(e_i - e + \epsilon_i)^{n-j}$ denote the probability that agent i finishes in j^{th} place given that all other agents exert effort e , agent i exerts effort e_i , and the value of agent i 's noise term is ϵ_i . From this it follows that the derivative of the agent's utility with respect to e_i is given by the following expression:

$$\sum_{j=1}^n v(A_j) [y_j(e, e_i, t_j - e_i) f(t_j - e_i) + \int_{t_j - e_i}^{\infty} \frac{\partial y_j(e, e_i, \epsilon_i)}{\partial e_i} f(\epsilon_i) d\epsilon_i] - c'(e_i)$$

By setting this derivative equal to zero when $e_i = e$, it then follows that it is an equilibrium for all agents to exert effort e if and only if $c'(e)$ is equal to

$$\sum_{j=1}^n v(A_j) [y_j(e, e, t_j - e) f(t_j - e) + \int_{t_j - e}^{\infty} \frac{\partial y_j(e, e_i, \epsilon_i)}{\partial e_i} \Big|_{e_i=e} f(\epsilon_i) d\epsilon_i]. \quad (3)$$

In the optimal threshold mechanism, the thresholds t_j must be chosen in such a way as to make the expression in equation (3) as large as possible for a given e . A necessary condition for this is that the derivative of this equation

with respect to t_j must be zero for all j . When we differentiate this equation with respect to t_j , the only term that remains is

$$v(A_j) y_j(e, e, t_j - e) f'(t_j - e). \quad (4)$$

Now note that the principal would never have an incentive to choose a value of t_j that is equal to $-\infty$ or ∞ for all j . To see this, note that for values of t_j that are arbitrarily negative, it must be the case that $t_j - e < 0$ in equilibrium, meaning $f'(t_j - e) > 0$ and the derivative of equation (3) with respect to t_j is positive. Thus for any sufficiently negative values of t_j , the principal can always increase the expression in equation (3) by increasing t_j , and thereby increase the equilibrium level of effort. From this it follows that a threshold of $t_j = -\infty$ can never be optimal for any j . A similar argument shows that a threshold of $t_j = \infty$ can also never be optimal. Thus there is some large finite value of $T > 0$ such that all thresholds $t_j \notin [-T, T]$ are dominated by some threshold $t_j \in [-T, T]$ for all j . Since this set of feasible thresholds is compact, and the equilibrium level of effort varies continuously with the thresholds, it then follows that some set of optimal thresholds exists.

Now if the optimal threshold t_j is not equal to $-\infty$ or ∞ , then it must be the case that the derivative in equation (4) is zero at the optimal t_j . Now we know that $f(\cdot)$ is single-peaked at 0, so when t_j is not equal to $-\infty$ or ∞ , the above derivative is only zero when $t_j = e$. Thus the optimal thresholds satisfy $t_j = e$ for all j and the optimal thresholds are always equal for all prizes. $\square$

Theorem 3.2 says that optimally incorporating cardinal scores into a rank-order contest to maximize effort is, in fact, even simpler than the method suggested by Theorem 3.1—a mechanism designer only needs to compare all agents' outputs against the *same* baseline, reducing the problem of finding the optimal modification of $\mathcal{M}(A_1, A_2, \dots, A_n)$ to one of choosing a single optimal threshold t^* for $\mathcal{M}(A_1, A_2, \dots, A_n)$. Our proof of Theorem 3.2 also contains a proof of the following result, which we state here since it is used repeatedly in the remainder of our analysis:

COROLLARY 3.1. *Suppose that $f(\cdot)$ is single-peaked at 0. Then the optimal threshold that induces the highest equilibrium effort in the threshold mechanism is equal to the equilibrium level of effort at that threshold.*

3.2 Comparative Statics

We now address the question of how the optimal threshold t^* varies with changes in the problem parameters, specifically with the number of contestants and changes in the prize structure. First we consider comparative statics with respect to the number of agents. Our mechanisms $\mathcal{M}(A_1, A_2, \dots, A_n)$ so far have been specified in terms of the prizes for each of the n ranks, where n is the number of players. Since we want to now vary n , we assume that there is a fixed number k of prizes, $A_1, A_2, \dots, A_k$, and the number of players n is greater than or equal to k . In this scenario, we prove the following result:

THEOREM 3.3. *Consider any given rank-order mechanism $\mathcal{M}$ with prizes $A_1, \dots, A_k$, and let $t^*(n)$ denote the optimal threshold for $\mathcal{M}$ when there are n agents in the contest. If the noise density $f(\cdot)$ is single-peaked at 0, then the optimal threshold $t^*(n)$ is decreasing in the number of players n .*

for all $n \geq k$. Further, the equilibrium effort in the optimal threshold mechanism also decreases with n .

PROOF. Note that a player i 's expected value for her prize from exerting effort e_i when all other players are exerting effort e is just equal to the sum, over all $j = 1, \dots, k$, of the difference between the value the player obtains from receiving the j^{th} prize A_j and her value from the $j+1^{\text{th}}$ prize A_{j+1} , multiplied by the probability that she finishes in at least j^{th} place and meets the threshold t . Now let $G_j(q)$ denote the probability that no more than $j-1$ of the $n-1$ values of $\epsilon_{i'}$ are greater than q . Then the probability that agent i finishes in at least j^{th} place for any given realization of ϵ_i when she exerts effort e_i and all other players are exerting effort e is $G_j(e_i - e + \epsilon_i)$. From this it follows that the probability agent i finishes in at least j^{th} place and meets the threshold t unconditional on the realization of ϵ_i is

$$\int_{t-e_i}^{\infty} G_j(e_i - e + \epsilon_i) f(\epsilon_i) d\epsilon_i.$$

By using the insights in the previous paragraph, it follows that player i 's expected utility from exerting effort e_i when all other players are exerting effort e is

$$\sum_{j=1}^k (v(A_j) - v(A_{j+1})) \int_{t-e_i}^{\infty} G_j(e_i - e + \epsilon_i) f(\epsilon_i) d\epsilon_i - c(e_i)$$

By differentiating this expression with respect to e_i , setting the derivative equal to zero, and using the fact that $e_i = e$ must hold in any symmetric pure strategy equilibrium, it follows that at the equilibrium level of effort e , $c'(e)$ must equal

$$\sum_{j=1}^k (v(A_j) - v(A_{j+1})) [G_j(t-e) f(t-e) + \int_{t-e}^{\infty} \frac{dG_j(\epsilon)}{d\epsilon} f(\epsilon) d\epsilon].$$

From Corollary 3.1, the equilibrium effort e at the optimal threshold t will satisfy $t - e = 0$. Thus equilibrium effort in the optimal threshold mechanism is the solution to

$$\sum_{j=1}^k (v(A_j) - v(A_{j+1})) [G_j(0) f(0) + \int_0^{\infty} \frac{dG_j(\epsilon)}{d\epsilon} f(\epsilon) d\epsilon] = c'(e). \quad (5)$$

Here $\frac{dG_j(\epsilon)}{d\epsilon}$ represents the density corresponding to this distribution. Write $G_j(\epsilon; n)$ to denote the dependence of G_j on n . Then, the distribution $G_j(\epsilon; n')$ first order stochastically dominates $G_j(\epsilon; n)$ for all $n' > n$. Further, $f(\epsilon)$ is non-increasing in ϵ for all $\epsilon \geq 0$, since we assumed that $f(\cdot)$ is single-peaked at 0. From this it follows that increasing n decreases the value of the expression

$$G_j(0) f(0) + \int_0^{\infty} \frac{dG_j(\epsilon)}{d\epsilon} f(\epsilon) d\epsilon = E_{\epsilon \sim G_j} [f(\max\{0, \epsilon\})]$$

for all j (by stochastic dominance). Therefore, for equality to hold in equation (5), the equilibrium effort $e(n)$ must be such that $c'(e(n))$ also decreases with n , i.e., $e(n)$ decreases with n . Since the equilibrium effort in the optimal threshold mechanism decreases with the number of players, and the optimal threshold equals equilibrium effort in the optimal threshold mechanism, by Corollary 3.1, the optimal threshold $t^*(n)$ is also decreasing in n . $\square$

Next we ask how the optimal threshold varies with the number of prizes awarded. To formulate this question meaningfully, we consider contests where the top k participants who meet the threshold all receive the same prize, for some k , and consider two ways that the total number of prizes might increase—first, where the total prize pool stays the same, but the prizes are split amongst a larger number of players, and second, where the value of each prize stays the same, but more prizes of this value are awarded (contingent on meeting the threshold). The optimal threshold varies predictably with these changes in the prize structure, as the following theorem illustrates:

THEOREM 3.4. Suppose that $f(\cdot)$ is single-peaked at zero and the number of prizes is less than the number of players.

(1) The optimal threshold t^* in a contest with k equal prizes of value A increases with k .

(2) The optimal threshold t^* in a contest with k equal prizes of value A/k each also increases with k if players are sufficiently risk-averse in the sense that the coefficient of absolute risk aversion on $v(\cdot)$ is sufficiently large.

PROOF. We know from equation (5) in the proof of the previous theorem that if $G_j(q)$ denotes the probability that no more than $j-1$ of the $n-1$ other values of $\epsilon_{i'}$ are greater than q , then the following condition must be satisfied by the equilibrium level of effort e in the optimal threshold mechanism:

$$\sum_{j=1}^k (v(A_j) - v(A_{j+1})) [G_j(0) f(0) + \int_0^{\infty} \frac{dG_j(\epsilon)}{d\epsilon} f(\epsilon) d\epsilon] = c'(e)$$

Substituting in the fact that $A_j = A_k$ for $j \leq k$ and $A_j = 0$ for $j > k$, it then follows that the equilibrium level of effort e satisfies the following equation:

$$v(A_k) [G_k(0) f(0) + \int_0^{\infty} \frac{dG_k(\epsilon)}{d\epsilon} f(\epsilon) d\epsilon] = c'(e),$$

which can be rewritten as

$$v(A_k) E_{\epsilon \sim G_k} [f(\max\{0, \epsilon\})] = c'(e). \quad (6)$$

Now if one increases the number of prizes by awarding additional prizes that are the same as those originally awarded to agents who finished in the top k and met the threshold, then $v(A_k)$ is independent of k . And if one increases the number of prizes by splitting the same total prize pool amongst a larger number of players, then in the limit as the coefficient of absolute risk aversion on $v(\cdot)$ becomes arbitrarily large, $\frac{v(A_k)}{v(A)}$ becomes arbitrarily close to 1, and $v(A_k)$ approaches a function that is also independent of k .

But we also know that $E_{\epsilon \sim G_k} [f(\max\{0, \epsilon\})]$ is increasing in k since the fact that $G_k(\epsilon)$ is a distribution corresponding to the probability that no more than $k-1$ of the $n-1$ values of ϵ_j are greater than ϵ implies that G_j first order stochastically dominates G_k for all $j < k$. And we also know that $f(\max\{0, \epsilon\})$ is decreasing in ϵ for all $\epsilon > 0$. By combining these facts, it follows that $E_{\epsilon \sim G_k} [f(\max\{0, \epsilon\})]$ is increasing in k . Thus under either of the conditions of the theorem, it must be the case that $v(A_k) E_{\epsilon \sim G_k} [f(\max\{0, \epsilon\})]$ is increasing in k . By combining this fact with equation (6), it then follows that under either of these conditions, equilibrium effort in the optimal threshold mechanism is increasing in k . Since the optimal threshold is equal to equilibrium effort in

the optimal threshold mechanism, it then follows that the optimal threshold is also increasing in k . $\square$

The intuition for this result has to do with how an agent's incentives to try to meet the threshold vary with the number of prizes. As the number of prizes increases in either of the two manners considered in Theorem 3.4, the expected value of the prize that an agent obtains for meeting the threshold unambiguously increases. Thus agents have a stronger incentive to try to meet the threshold when the number of prizes increases, and agents will thus exert more effort in equilibrium when there are a larger number of prizes. Since the optimal threshold is equal to the equilibrium level of effort (from Corollary 3.1), it then follows that increasing the number of prizes also increases the optimal threshold.

4. HOW USEFUL IS CARDINAL INFORMATION?

Our results thus far show that making the award of the prizes in a rank-order contest $\mathcal{M}(A_1, A_2, \dots, A_n)$ contingent on submission qualities exceeding a suitable threshold creates the strongest incentives for effort, and in particular, outperforms the contest that ignores cardinal information. In this section, we briefly address the question of *how much* of an improvement is obtained from using this cardinal information relative to $\mathcal{M}(A_1, A_2, \dots, A_n)$, which awards prizes based only on relative rankings.

We first present simulation results to obtain a quantitative sense for the size of the improvement in equilibrium effort as a function of the number of contestants, the nature of participants' cost functions, the reward structure $(A_1, A_2, \dots, A_n)$, and the noise distribution governing the stochastic perturbations that influence an agent's output. We ask to what extent the size of these gains are affected by changes in these underlying parameters, and discuss these results in the context of typical parameter values in online contests. Finally, we conclude with a result (proven in the full version of the paper [15]) that formalizes the observations from these simulations.

Simulations. We simulate contests where n players each have cost function $c(\hat{e}) = \frac{\hat{e}^a}{a}$ for some constant $a > 1$ (a larger a corresponds to a more convex cost function). We assume that players are risk-neutral with values $v(A) = A$ for prize A . For the distribution of the noise terms ϵ_i that randomly influence a player's submission quality as $q_i = e_i + \epsilon_i$, we consider IID draws from (i) a standard normal distribution and (ii) a standard Laplace distribution. For the set of prize structures, we consider contests with prizes of A_1 and A_2 for the top two ranks and $A_j = 0$ for all lower ranks for simplicity and brevity. Each of n , a , and the split into A_1 and A_2 are parameters describing the contest that we will vary in our simulations.

Computing equilibrium effort. To quantify the extent of improvement in equilibrium effort from using cardinal information in a contest $\mathcal{M}(A_1, A_2, \dots, A_n)$, we first need to calculate the equilibrium level of effort both with no threshold and with the optimal threshold. When there is no threshold, the equilibrium level of effort is the same as it would be if the threshold were $t = -\infty$. Applying equation (3) to the special case in which $t = -\infty$, we see that the equilibrium level of effort e_ϕ^* with no thresholds must satisfy

$$\sum_{j=1}^n v(A_j) \left[\int_{-\infty}^{\infty} \left[\frac{(n-1)!}{(j-1)!(n-j-1)!} (1-F(\epsilon_i))^{j-1} F(\epsilon_i)^{n-j-1} - \frac{(n-1)!}{(j-2)!(n-j)!} (1-F(\epsilon_i))^{j-2} F(\epsilon_i)^{n-j} \right] f^2(\epsilon_i) d\epsilon_i \right] = e_\phi^{*a-1}$$

where we abuse notation by defining $\frac{(n-1)!}{(j-2)!(n-j)!}$ to be zero when $j = 1$.

When the mechanism designer uses the *optimal* threshold t^* , we know from Corollary 3.1 that the equilibrium level of effort e_O^* equals the optimal threshold. Therefore, setting $t = e$, we see that e_O^* must satisfy

$$e_O^{*a-1} = \sum_{j=1}^n v(A_j) \left[\left(\frac{n-1}{j-1} \right) \left(\frac{1}{2} \right)^{n-1} f(0) + \int_0^\infty \left[\frac{(n-1)!}{(j-1)!(n-j-1)!} (1-F(\epsilon_i))^{j-1} F(\epsilon_i)^{n-j-1} - \frac{(n-1)!}{(j-2)!(n-j)!} (1-F(\epsilon_i))^{j-2} F(\epsilon_i)^{n-j} \right] f^2(\epsilon_i) d\epsilon_i \right] \quad (7)$$

We can compute the percentage increase in equilibrium effort for a particular set of contest parameters by using the equations above to compute e_ϕ^* and e_O^* . We compute these improvements for a wide range of different parameters to observe how the various parameters affect the extent of the increase in equilibrium effort from using cardinal information via the optimal threshold mechanisms. While we cannot present the full results for want of space, we summarize the results for a few parameter values in Tables 1, 2, and 3.

n	Normal distribution	Laplace distribution
2	20.71%	50.00%
3	4.95%	16.67%
4	1.59%	7.14%
5	0.58%	3.33%
6	0.23%	1.61%
7	0.10%	0.80%
8	0.04%	0.40%
9	0.02%	0.19%

Table 1: Percentage increase in equilibrium effort from using the optimal threshold under a winner-take-all contest when the cost to effort $\hat{e}$ is $c(\hat{e}) = \frac{\hat{e}^2}{2}$ for different noise distributions and varying values of the number of players n .

Discussion of simulation results. The simulations lead to a number of interesting insights. First, they show that when the number of players is larger, the benefit to setting the optimal threshold is relatively smaller. We also see that when the parameter a in the exponent of the cost function $c(\hat{e}) = \frac{\hat{e}^a}{a}$ increases, so that players' cost functions are more convex, less is gained by using the optimal threshold. Finally, when more of the prizes are given to the lower-ranked players, there is again a larger benefit from using the optimal threshold.

These observations all have intuitive explanations. Since equilibrium effort equals the optimal threshold (Corollary 3.1), when there are a large number of players, the final realized quality of the winning player is likely to far exceed

a	Normal distribution	Laplace distribution
1.25	6.51%	31.78%
1.5	3.21%	14.80%
1.75	2.12%	9.64%
2	1.59%	7.14%
3	0.79%	3.51%
4	0.53%	2.33%
5	0.40%	1.74%
6	0.32%	1.39%

Table 2: Percentage increase in equilibrium effort from using the optimal threshold for different noise distributions and varying values of the term a in the cost function $c(\hat{e}) = \frac{\hat{e}^a}{a}$, under a winner-take-all contest with $n = 4$.

(A_1, A_2)	Normal distribution	Laplace distribution
(1, 0)	0.58%	3.33%
(0.9, 0.1)	1.06%	5.17%
(0.8, 0.2)	1.59%	7.14%
(0.7, 0.3)	2.20%	9.26%
(0.6, 0.4)	2.90%	11.54%
(0.5, 0.5)	3.72%	14.00%

Table 3: Percentage increase in equilibrium effort from using the optimal threshold for different noise distributions and varying distributions of the top prizes, with five players when $c(\hat{e}) = \frac{\hat{e}^2}{2}$.

the threshold, so the threshold provides little incentive to exert additional effort. By contrast, when there are only a small number of players, there is a significant risk that even the winning player’s final realized quality will be in the same range as the threshold, and the threshold may therefore provide significant incentives to exert effort. Thus using the optimal threshold should have a relatively smaller effect when there are a larger number of players.

We note that both regimes with large numbers of players and small numbers of players are relevant empirically. For example, in Kaggle contests, there are typically at least dozens and usually hundreds of participants, so in such a contest setting the optimal threshold is likely to have little effect on incentives to exert effort. By contrast, many tasks on TaskCN and Topcoder⁷ only attract a few participants⁸, so in such contests, using the optimal threshold is likely to significantly improve effort. There are also online contests with elimination rounds leading up to a final round with a few finalists (such as the Robin Hood Foundation Prize⁹). In such settings, our results would suggest that using an optimal threshold to award prizes is likely to significantly improve effort in the final round of the contest.

⁷<http://community.topcoder.com/tc?module=CompList&pt=35>

⁸For example, [25] notes that most programming tasks on TaskCN attracted fewer than six participants and [42] notes that it is far more common to receive a small number of solutions for a task on TaskCN than some larger number of solutions.

⁹<https://www.robinhood.org/prize>

To understand why convexity affects equilibrium effort, note that when the players’ cost functions are more convex, a change in a player’s effort does relatively more to increase the player’s marginal cost of exerting effort, which in turn implies that a player can only increase his effort by a smaller amount in response to improved incentives before the player’s marginal cost from exerting more effort will equal his marginal benefit. Thus using the optimal threshold will have less effect when the players’ cost functions are more convex.

Finally we address the effect of the distribution of prizes on effort. The lower-ranked players are more likely to have final realized qualities that are close to the threshold, so the threshold provides a greater incentive for a lower-ranked player to exert effort than a higher-ranked player. Thus when a relatively greater amount of the prize pool is given to the lower-ranked players, this has the effect of shifting prizes from players who will be less strongly incentivized by the thresholds to players who will be more strongly incentivized by the thresholds. Thus using the optimal threshold has a relatively greater effect when a larger amount of the prize pool is distributed amongst lower ranks.

The insights from these simulations are, in fact, more general than this. In the full version of the paper [15], we prove the following result under some technical assumptions:

THEOREM 4.1 (INFORMAL). *The absolute and relative differences between the equilibrium effort elicited in the optimal threshold mechanism and an ordinal rank-order mechanism with no threshold is (i) decreasing in the number of players, (ii) increasing in the number of prizes if all the prizes are equal, and (iii) decreasing in the convexity of the cost function, where the degree of convexity of the cost function is measured by the ratio $\frac{c''(e)}{c'(e)}$.*

5. DISCUSSION AND FURTHER WORK

In this paper, we addressed the problem of how a principal running a contest might optimally incorporate cardinal information regarding the absolute qualities of contestants’ entries into an existing rank-order tournament, motivated by the observation that an increasing number of contests today evaluate entries according to some numerical metric. We found that threshold mechanisms, which compare a submission’s score against an absolute threshold—in fact, the *same* threshold for each prize—are optimal amongst the class of all mixed cardinal-ordinal mechanisms which award the agent with the j^{th} -ranked submission with quality q_j an arbitrary quality-dependent fraction $g_j(q_j)$ of the prize A_j . Therefore, using cardinal information as coarsely as by comparing against a single threshold provides the optimal modification of a rank-order mechanism in terms of incentives for effort. Finally, we saw that gains from incorporating cardinal information relative to using only ordinal information are highest for small contests with more than one prize and with participants whose costs to effort are not too convex.

Further results. There are a number of interesting questions regarding incentives in cardinal contests that we address in the full version of this paper [15].

Endogenous entry. A first natural question is how the results would be affected by the possibility of endogenous entry. While we have assumed throughout this paper that the number of players in the contest is known and fixed,

there are many real-world situations in which participation, even with very low effort, is costly and an agent must strategically decide whether to participate at all. We address how the possibility of endogenous entry affects the equilibrium in [15], and show that the choice of threshold in a threshold mechanism results in an interesting participation-effort tradeoff: the equilibrium level of participation may increase as a result of decreasing the threshold, but such equilibria will result in the agents exerting lower levels of effort. How the principal resolves this tradeoff will then depend on how much the principal values participation versus effort.

Optimal rank-based prizes. Another interesting question relates to whether the optimal division of prizes would change as a result of optimally incorporating cardinal information into a contest. In particular, suppose that the optimal division of the total prize pool A amongst the n places for a rank-order mechanism is $(A_1, \dots, A_n)$. Is it still the case that this is the optimal division of the prizes amongst the n places for the optimal mixed cardinal-ordinal mechanism that makes use of the optimal threshold? We address this question in [15] as well, and show that the answer to this question depends crucially on whether the agents are risk-averse: if agents are risk-neutral in the sense that $v(A)$ varies linearly with A , then the optimal division of prizes will be the same for both a pure rank-order mechanism and the optimal threshold mechanism. However, if agents are risk-averse in the sense that $v(A)$ is strictly concave in A , then this need not be the case. Instead, the principal will typically want to reward a larger percentage of the prize pool to the lower ranks under the optimal threshold mechanism than in a purely ordinal tournament.

Learning optimal thresholds. A number of questions also arise from practical considerations regarding uncertainty. For instance, the optimal threshold t^* for a given rank-order contest depends on the parameters of the population, which the designer might typically not have access to. As such, it might be difficult for a mechanism designer to set the optimal threshold in a contest. How could a mechanism designer who intends to run multiple iterations of a contest learn what the optimal threshold is from the earlier iterations of the contest? We show in [15] how a mechanism designer without complete knowledge of the contest population could make probabilistic inferences about whether the threshold selected was too high or too low from the results of previous contests and use this information to better set the threshold in future iterations.

Open questions. A number of open questions remain for further research. First, the specific question we ask about modifying given rank-order mechanisms to incorporate cardinal information is motivated by the fact that practical considerations beyond incentives for effort—such as simplicity, sponsorships of various prize levels, media or publicity considerations—might cause a principal to choose a particular prize structure for his contest. However, it is also very interesting to study the more general optimal contest design problem in such contests with access to absolute measurements of quality. What mechanism $\mathcal{M}(q_1, \dots, q_n)$ incentivizes the highest effort over all mechanisms with access to cardinal, and not just ordinal, information about outputs?¹⁰ And how does the answer to this question depend on the

specifics of the model such as agents' risk preferences and the objective function of the mechanism designer?

A second intriguing theoretical question regards a connection to optimal auctions and reserve pricing. There is a formal connection between auctions and contests ([7], [8], [11]) in an *alternative* model for contests where effort deterministically translates into output (see §2 for a discussion regarding the difference with our model). At first glance, this might suggest a mapping between the model and analysis in our paper and that in the literature on sponsored search auctions [32]. However, this simplified mapping is a fallacious analogy in contests where output is a stochastic perturbation of effort, because the fact that an agent's output is determined by a noise-perturbed version of effort completely changes both the structure of the problem and all the underlying analysis relative to auctions. Nevertheless, the parallels between the results on optimality of threshold-like structures in these two settings—contests and auctions—raises the intriguing question of whether there is indeed a way to formally relate our results on optimalities of thresholds in contests to optimality of reserve prices in auctions. A deeper understanding of the connection between threshold mechanisms and optimal auctions is an interesting open direction for further research.

6. REFERENCES

- [1] N. Archak. Money, glory and cheap talk: analyzing strategic behavior of contestants in simultaneous crowdsourcing contests on TopCoder.com. In *Proceedings of the 19th International Conference on the World Wide Web (WWW)*, pages 21–30, 2009.
- [2] N. Archak, A. Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the 30th International Conference on Information Systems (ICIS)*, pages 1–16, 2009.
- [3] K.J. Boudreau, N. Lacetera, K.R. Lakhani. Incentives and problem uncertainty in innovation contests: an empirical analysis. *Management Science*, 57(5): 843–863, 2011.
- [4] C.B. Cadsby, C. Bram, F. Song, F. Tapon. Sorting and incentive effects of pay-for-performance: an experimental investigation. *Academy of Management Journal*, 50(2): 387–405, 2007.
- [5] J. Carpenter, P.H. Matthews, J. Schirm. Tournaments and office politics: evidence from a real effort experiment. *American Economic Review*, 100(1): 504–517, 2010.
- [6] R. Cavallo, S. Jain. Winner-take-all crowdsourcing contests with stochastic production. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, pages 34–41, 2013.
- [7] S. Chawla, J.D. Hartline, B. Sivan. Optimal crowdsourcing contests. In *Proceedings of the 23rd ACM-SIAM Symposium of Discrete Algorithms (SODA)*, pages 856–868, 2012.
- [8] Y.-K. Che, I. Gale. Optimal design of research tournaments. *American Economic Review*, 93(3): 646–671, 2003.
- [9] Y. Chen, T.-H. Ho, Y.-M. Kim. Knowledge market design: a field experiment at Google Answers. *Journal of Public Economic Theory*, 12(4): 641–664, 2010.

¹⁰[7] addresses this question in a specific (but different) model; see §1.1.

- [10] S.M. Chowdhury, R.M. Sheremeta, T.L. Turocyd. Overbidding and overspreading in rent-seeking experiments. *Games and Economic Behavior*, 87: 224–238, 2014.
- [11] D. DiPalantino, M. Vojnović. Crowdsourcing and all-pay contests. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 119–128, 2009.
- [12] T. Dohmen, A. Falk. Performance pay and multidimensional sorting: productivity, preferences, and gender. *American Economic Review*, 101(2): 556–590, 2011.
- [13] T. Eriksson, M.C. Villeval. Performance pay, sorting, and social motivation. *Journal of Economic Behavior and Organization*, 68(2): 412–421, 2008.
- [14] Q. Fu, J. Lu. The optimal multi-stage contest. *Economic Theory*, 51(2): 351–382, 2012.
- [15] A. Ghosh, P. Hummel. Cardinal contests. Mimeo, 2015.
- [16] A. Ghosh, R.P. McAfee. Crowdsourcing with endogenous entry. In *Proceedings of the 20th International World Wide Web Conference (WWW)*, pages 137–146, 2012.
- [17] A. Glazer, R. Hassin. Optimal contests. *Economic Inquiry*, 26(1): 133–143, 1988.
- [18] J.R. Green, N.L. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3): 349–364, 1983.
- [19] B. Holmstrom. Moral hazard in teams. *Bell Journal of Economics*, 13(2): 324–340, 1982.
- [20] K.A. Konrad. Dynamic contests. Max Planck Institute for Intellectual Property Typescript, 2010.
- [21] E.P. Lazear. Pay equality and industrial politics. *Journal of Political Economy*, 97(3): 561–580, 1989.
- [22] E.P. Lazear. Performance pay and productivity. *American Economic Review*, 90(5): 1346–1361, 2000.
- [23] E.P. Lazear, S. Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89(5): 841–864, 1981.
- [24] B. Liu, J. Liu, R. Wang, J. Zhang. Prize and punishment: optimal contest design with incomplete information. National University of Singapore Typescript, 2013.
- [25] T.X. Liu, J. Yang, L.A. Adamic, Y. Chen. Crowdsourcing with all-pay auctions: a field experiment on Taskcn. *Management Science*, 60(8): 2020–2037, 2014.
- [26] D.B. Minor. Increasing revenue through rewarding the best (or not at all). Northwestern University Typescript, 2013.
- [27] B. Moldovanu, A. Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3): 542–558, 2001.
- [28] B. Moldovanu, A. Sela. Contest architecture. *Journal of Economic Theory*, 126(1): 70–96, 2006.
- [29] B. Moldovanu, A. Sela, X. Shi. Contests for status. *Journal of Political Economy*, 115(2): 338–363, 2007.
- [30] J. Morgan, H. Orzen, M. Sefton. Endogenous entry in contests. *Economic Theory*, 51(2): 435–463, 2012.
- [31] B.J. Nalebuff, J.E. Stiglitz. Prizes and incentives: towards a general theory of compensation and competition. *Bell Journal of Economics*, 14(1): 21–43, 1983.
- [32] M. Ostrovsky, M. Schwarz. Reserve prices in Internet advertising auctions: a field experiment. Stanford Graduate School of Business Typescript, 2009.
- [33] S. Parreiras, A. Rubinchik. Contests with three or more heterogeneous agents. *Games and Economic Behavior*, 68(2): 703–715, 2010.
- [34] C. Riis. Efficient contests. *Journal of Economics and Management Strategy*, 19(3): 643–665, 2010.
- [35] R. Siegel. All-pay contests. *Econometrica*, 77(1): 71–92, 2009.
- [36] R. Siegel. Asymmetric contests with conditional investments. *American Economic Review*, 100(5): 2230–2260, 2010.
- [37] R. Siegel. Participation in deterministic contests. *Economics Letters*, 116(3): 588–592, 2012.
- [38] R. Siegel. Asymmetric contests with productive effort. *American Economic Journal: Microeconomics*, 6(3): 59–105, 2014.
- [39] R. Siegel. Contests with productive effort. *International Journal of Game Theory*, 43(3): 515–523, 2014.
- [40] S. Szymanski, T.M. Valletti. Incentive effects of second prizes. *European Journal of Political Economy*, 21(2): 467–481, 2005.
- [41] C.R. Taylor. Digging for golden carrots: an analysis of research tournaments. *American Economic Review*, 85(4): 872–890, 1995.
- [42] J. Yang, L.A. Adamic, M.S. Ackerman. Crowdsourcing and knowledge sharing: strategic user behavior on TaskCN. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pages 246–255, 2008.