Simple vs Optimal Contests with Convex Costs

Amy Greenwald
Brown University
Computer Science
Providence, RI, USA

Takehiro Oyakawa
Brown University
Computer Science
Providence, RI, USA

Vasilis Syrgkanis
Microsoft Research
New England
Cambridge, MA, USA

ABSTRACT

We study an optimal contest design problem where contributors abilities are private, their costs are convex as a function of their effort, and the designer seeks to maximize their total effort. We address the design of approximately optimal mechanisms that are robust in that they are independent of the ability distribution and the precise form of the cost function. We show that a very simple all-pay contest where the prize is distributed equally among the top quartile of contributors is always a constant factor approximation to the optimal for a large class of convex cost functions, when the number of contributors is larger than some constant. This result stands in contrast to contests with linear costs, where awarding a prize to a single top contributor is approximately-optimal; when costs are convex, this latter allocation is far from optimal. Our result is enabled by novel results in the space of optimal mechanism design with convex costs, which could be of independent interest. Finally, we validate the performance of our approximately-optimal contests via simulation experiments, and portray much better empirical performance than the worst-case guarantees.

ACM Reference format:
Amy Greenwald, Takehiro Oyakawa, and Vasilis Syrgkanis. 2016. Simple vs Optimal Contests with Convex Costs. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference'17),* 10 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The Netflix challenge was a contest in which Netflix solicited prediction algorithms from the general public, and promised a $1 million prize to any team whose accuracy exceeded their own by at least 10%. Such crowdsourcing contests have become prevalent in the web economy: Kaggle competitions, user-generated content, and Topcoder are all examples of crowdsourced contests. As the host's objective usually depends on the quality of the contributions, an obvious question arises: how to design contests that spur the greatest innovation among potential contributors. Although a problem of this nature was first posed by Francis Galton in 1902, the prevalence of contests on the web has given rise to a long line of recent work at the intersection of computer science and economics addressing the optimal design of contests [3, 9, 11, 12, 14–17, 20].

Starting from a seminal work in the economics literature [22], typical models of contest design formulate a game in which a prize is to be distributed to participants based on the quality of their contributions. This level of quality is a strategic decision, and the participants incur costs (in the form of requisite effort) that depend on this choice. Participants also have private abilities, which offset their costs: i.e., more talented participants can produce the same quality contribution with less effort. The goal of each participant, then, is to maximize their share of the prize by making a worthy contribution, while simultaneously minimizing their efforts/costs; their utility is the difference between these two quantities.

Much of the existing work on optimal contest design assumes either that participants have essentially the same ability[1] [14–17], or that cost is a linear function of the quality of a contribution [3, 9, 11, 20].[2] However, the effort required to construct a high quality contribution naturally increases in a nonlinear fashion. It is easy to design a cookie-cutter solution; it is much harder to produce an original design. The famous 80/20 rule, which states that "20% of the work gets you 80% of the way there," captures this nonlinearity. In other words, effort is often characterized by (significantly) diminishing returns. The difficulty in addressing contest design with both nonlinear costs and participants of heterogeneous ability hinges on the fact that the combination of these two elements equates the contest design problem to that of mechanism design with non quasi-linear utility functions. Hence, unlike the case of linear costs, standard results from Myerson's theory of optimal mechanism design [23] are not fully available.

In this work, we address optimal contest design with convex costs and heterogeneous participant abilities drawn from some prior distribution. Moreover, following the recent literature on prior free mechanism design [19, 25], we strive for contests that do not depend on the prior distribution over abilities, or the details of the cost function. Our goal is to design simple approximately-optimal contests that are independent these *details* of the setting.

We work with the general model presented in the classic work of Moldovanu and Sela [22]. In particular, the goal of the designer is to maximize the sum of the participants' contribution qualities, each of which has a cost for contributing p of the form $c(p)/v$, where v is an ability and c is a convex function. If a participant wins a prize of x, then their overall utility is $u(x, p; v) = x - c(p)/v$. Despite defining this general model, [22] obtain optimality or near optimality results only in the case of linear costs. For convex costs, they analyze mechanisms that can distribute the prize to at most two contributors, in which case they state conditions under which

[1] Equivalently, they are unaware of their ability at the time of the contribution quality decision.
[2] Other work analyzes the equilibrium and performance properties of various fixed contest mechanisms [12].

Conference'17, Washington, DC, USA

DOI: 10.1145/nnnnnnn.nnnnnnn

allocating only to top contributors is optimal. However, for convex costs, allocating to only a few contributors can be very suboptimal.[3]

Contributions. Our main result is to show that a simple contest which distributes the prize equally among the top quartile of contributors achieves a constant factor approximation for a large class of distributions of abilities (bounded Monotone-Hazard-Rate distributions) and for the class of convex cost functions of the form $c(p) = p^d$, where $d \geq 2$ is a measure of convexity, when the number of contributors is larger than a constant.

Our results are enabled by a simulation theorem (which holds for *all* convex cost functions) whereby the outcome of a direct mechanism—where the designer elicits the private abilities of the players and then decides (on their behalf) how much each one should contribute and how much each one will be allocated—can be implemented as the Bayes-Nash equilibrium (BNE) of a contest where the participants choose their own contribution qualities. This simulation theorem, together with the relevation principle, establishes an equivalence between optimal mechanism design and optimal contest design. There are two main implications of this equivalence:

(1) It points to the difficulty of optimal contest design with convex costs, as it is an instance of optimal mechanism design with non quasi-linear utilities. Consequently, although our main results hinge on a very specific form of monomial cost functions, the difficulty of taking even this step should be apparent.
(2) It shows that restricting to contests does not require that a designer accept any sacrifice in performance (assuming they are willing to accept BNE implementation). Even if the designer were to try to devise some more complex mechanism, any such mechanism is implementable by a direct mechanism by the revelation principle, and any direct mechanism is implementable as a BNE of a contest by our simulation theorem. Hence, the total expected contributions of any indirect mechanism is achievable at a BNE of some contest.

Given this simulation theorem, we restrict our attention to direct mechanisms. As already noted, Myerson's theory of optimal mechanism design is not fully applicable in our convex-cost model. Hence, we begin by deriving an upper bound on the total expected contributions achievable by a direct mechanism. Subsequently, we analyze a simple class of threshold mechanisms, where the prize is equally divided among all contributors whose ability surpasses some threshold. We show that very simple thresholds, including the median of the ability distribution as well as a randomly sampled ability, achieve a constant factor approximation to our upper bound on the optimal, and hence the optimal.

We then show that when the number of contributors is larger than a constant,[4] an approximately-optimal threshold mechanism can be well approximated (at the loss of an extra constant factor) by a contest that awards the prize to a carefully chosen top fraction of contributors. For instance, awarding equally the top half of contributors is very close to the threshold mechanism that awards to all contributors whose ability is above the median; likewise, for other quantiles of the ability distribution. This observation allows us to establish our main result: that allocating to the top quartile of contributors is a constant factor approximation to the optimal contest. We note, once again, that the flavor of this result is quite different from that of the linear-cost model, where a "winner-take-all" design (i.e., allocating only to top contributors) is optimal. When costs are convex, it is preferable to incentivize additional effort from multiple participants, rather than only from those of the greatest ability.

We conclude by evaluating the performance of the proposed contests via simulation experiments. We show that the top quartile contest and other more refined contests whose thresholds depend on the degree of convexity of the cost function perform significantly better on average than the worse case bound of our theorems.

Our results also have implications for revenue maximization with non quasilinear utilities [1, 7, 13, 18, 21, 24]. One implication is a prior-free approximately-optimal mechanism for revenue maximization, assuming our form of utility functions, which could be of independent interest in the revenue maximization literature.

2 ALL-PAY CONTEST MODEL

We consider a contest model, which we dub an "all-pay" contest model, because it can be understood as an all-pay simultaneous reverse auction. In this model, a contest designer has one unit of prize money to award to a set of n contributors/players. Each player $i \in N = \{1, \ldots, n\}$ has a private ability value v_i, drawn independently from an atomless distribution F, with continuous probability density f that is strictly positive on the support, which is the closed interval $T = [0, \bar{v}]$. We write $\mathbf{v} = (v_1, \ldots, v_n) \in T^n$ to denote a sample ability vector, drawn from distribution F^n.

Conditional on their ability, each player i chooses a quality/level of contribution $b_i \in \mathbb{R}^n$, which hereafter we refer to solely as contribution. To contribute b_i, player i incurs a cost $c_i(b_i)/v_i$, where $c_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is an increasing continuous convex cost function with $c(0) = 0$. Given a vector of contribution qualities $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{R}^n$ the designer then chooses an allocation of the prize, $\mathbf{x}(\mathbf{b}) \in [0, 1]^n$, such that $\sum_{i \in N} \mathbf{x}(\mathbf{b}) = 1$. Player i's utility is the prize awarded less the cost incurred:

$$u_i(b_i, \mathbf{b}_{-i}; v_i) = x_i(b_i, \mathbf{b}_{-i}) - c_i(b_i)/v_i. \tag{1}$$

For vectors such as $\mathbf{b}$, we use the notation $\mathbf{b} = (b_i, b_{-i})$ to distinguish between player i's role in the contest and all others $N \setminus \{i\}$.

Solution concept. We assume each player chooses a contribution in such a way as to maximize their expected utility conditional on their own ability and in expectation over their opponents' abilities. Formally, a vector of functions $\mathbf{b}(\mathbf{v}) = (b_i(v_i), \ldots, b_n(v_n))$ is a Bayes-Nash equilibrium (BNE) if: $\forall i \in N$, $\forall v_i \in [0, \bar{v}]$, $\forall b_i' \in \mathbb{R}_{\geq 0}$

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i}}\left[u_i(b_i(v_i), \mathbf{b}_{-i}(\mathbf{v}_{-i}); v_i)\right] \geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i}}\left[u_i(b_i', \mathbf{b}_{-i}(\mathbf{v}_{-i}); v_i)\right] \tag{2}$$

Designer objective. The goal of the contest designer is to design a prize allocation rule so as to maximize **total expected contributions** at the BNE of the resulting all-pay contest, i.e.: $\mathbb{E}_{\mathbf{v}}\left[\sum_{i \in N} b_i(v_i)\right]$.

Utility transformation. For technical reasons, it is convenient to translate the above contest model to a mathematically equivalent

[3] In the full version, we present a concrete example where the inefficiency of a winner-takes-all contest grows with the number of players.
[4] Which depends logarithmically on the ratio between the upper bound and the median of the ability distribution.

one, where players have utility of the form:

$$u_i(b_i, \mathbf{b}_{-i}; v_i) = v_i x_i(b_i, \mathbf{b}_{-i}) - c_i(b_i). \quad (3)$$

Since players maximize utility conditional on their ability, maximizing the utility in Equation (1) is equivalent to maximizing the utility in Equation (4), which is simply the former multiplied by v_i. This latter form of utilities is more convenient because it allows us to interpret the setting in the forward direction, where player i has value v_i per-unit of good (prize) and incurs a cost that is a function of his payment (contribution) b_i. We can then make use the extensive toolbox developed in the revenue maximization literature. This utility transformation approach was also used in prior work on contests with linear costs [9, 11, 20].

Direct Mechanisms. By the classic revelation principle [23], an all-pay contest is outcome equivalent (i.e., induces the same contribution from each player, conditional on his value, and the same allocation of prizes) to what is known as a *direct mechanism*. In a direct mechanism, the designer elicits the private information of the participants directly, which in our setting is an ability report $w_i \in T$ from each player i.

Given a vector of reports $\mathbf{w} = (w_1, \dots, w_n) \in \mathbb{R}^n$, for all $i \in N$, a mechanism is defined by an *allocation rule* $\mathbf{x}(\mathbf{w}) \in [0,1]^n$, such that $\sum_{i \in N} x_i(\mathbf{w}) = 1$, together with a *contribution rule* $\mathbf{p}(\mathbf{w}) \in \mathbb{R}^n_{\geq 0}$, that specifies the contribution $p_i(\mathbf{w})$ that is required of each player i. The utility of each player i is then:

$$u_i(w_i, \mathbf{w}_{-i}; v_i) = v_i x_i(w_i, \mathbf{w}_{-i}) - c_i(p_i(w_i, \mathbf{w}_{-i})), \quad (4)$$

For readability, we often write $c_i(w_i, \mathbf{w}_{-i})$ instead of $c_i(p_i(w_i, \mathbf{w}_{-i}))$. This is player i's *cost*, which depends on their contribution $p_i(\mathbf{w})$. Similar to the contribution rule, we refer to $\mathbf{c}(\mathbf{w}) \in \mathbb{R}^n_{\geq 0}$, comprised of $c_i(w_i, \mathbf{w}_{-i})$'s, as the *cost rule* of the direct mechanism.

Given a direct mechanism, we define the **interim allocation**, **interim contribution** and **interim cost** rules, respectively as: $\hat{x}_i(w_i) = \mathbb{E}_{\mathbf{v}_{-i}}[x_i(w_i, \mathbf{v}_{-i})]$, $\hat{p}_i(w_i) = \mathbb{E}_{\mathbf{v}_{-i}}[p_i(w_i, \mathbf{v}_{-i})]$ and $\hat{c}_i(w_i) = \mathbb{E}_{\mathbf{v}_{-i}}[c_i(w_i, \mathbf{v}_{-i})]$. These variables comprise each player's expected allocation, contribution, and cost as a function of their report, assuming all the other players report their abilities truthfully.

We call a mechanism **Bayesian incentive compatible** (BIC) if utility is maximized by truthful reports in expectation: $\forall i \in N$ and $\forall v_i, w_i \in T$, $v_i \hat{x}_i(v_i) - \hat{c}_i(v_i) \geq v_i \hat{x}_i(w_i) - \hat{c}_i(w_i)$. **Interim individual rationality** (IIR) insists on non-negative utilities in expectation: $\forall i \in N$ and $\forall v_i \in T$, $v_i \hat{x}_i(v_i) - \hat{c}_i(v_i) \geq 0$.

The goal of the designer is to find a BIC and IIR mechanism that maximizes **total expected contributions**, which is defined as $\mathbb{E}_{\mathbf{v}}[\sum_{i \in N} p_i(\mathbf{v})] = \sum_{i \in N} \mathbb{E}_{v_i}[\hat{p}_i(v_i)]$.

Lemma 2.1 (Application of the Revelation Principle). *The total expected contributions achieveable by an all-pay contest at some BNE can also be achieved by a direct BIC/IIR mechanism.*

This lemma follows from the well-known revelation principle, so its proof is omitted. Intuitively, the allocation and contribution rule of the direct mechanism will simulate the allocation and payment rules specified by the BNE of the all-pay contest, and it will be BIC and IIR since it is simulating a BNE. More interestingly, in the next section, we show that the other direction also holds, so that the total expected contributions of any direct mechanism that is BIC and IIR can be achieved as a BNE of an all-pay contest. In this way, we establish an equivalence between optimal direct mechanisms and optimal all-pay contests.

Characterization of Direct Mechanisms. Myerson [23] showed that for a mechanism to satisfy BIC and IIR, several conditions need to hold. Among them, a specific cost formula must be used. We restate his result below, adapted to our setting. For simplicity, we assume the utility of reporting an ability of zero is zero: i.e., $u_i(0, \mathbf{b}_{-i}) = 0$ for all $i \in N$, as is the case for all our mechanisms.

Lemma 2.2 ([23]). *A mechanism is BIC and IIR if and only if the following conditions hold:*

- *The allocation rule is monotone:*

$$\hat{x}_i(v_i) \geq \hat{x}_i(w_i), \quad \forall i \in N, \forall v_i \geq w_i \in T, \quad (5)$$

- *Interim costs satisfy the following condition:*

$$v_i \hat{x}_i(v_i) - \hat{c}_i(v_i) = \int_0^{v_i} \hat{x}_i(z_i)\, \mathrm{d}z_i, \quad \forall i \in N, \forall v_i \in T, \quad (6)$$

Moreover, the total expected cost of such a mechanism can be described using a virtual value function φ_i, defined as follows:

$$\varphi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}. \quad (7)$$

Theorem 2.3 ([23]). *The total expected costs incurred by participants in a BIC and IIR mechanism satisfies:*

$$\sum_{i \in N} \mathop{\mathbb{E}}_{v_i \sim F_i}[\hat{c}_i(v_i)] = \sum_{i \in N} \mathop{\mathbb{E}}_{v_i \sim F_i}[\varphi_i(v_i) \hat{x}_i(v_i)] \quad (8)$$

In the traditional quasi-linear setting, where the cost function is the identity function, Myerson's theorem tells us that expected virtual surplus is equivalent to total expected contributions. From this result and the forthcoming simulation theorem, we can deduce that, to maximize contributions, the prize should be allocated to the players with the highest non-negative virtual values. However, in the convex cost setting, Myerson's lemma does not pin down the interim contributions of the players; only their interim costs. Hence the virtual value characterization theorem does not shed much light on the optimal mechanism.

3 BNE OF ALL-PAY CONTESTS ARE OPTIMAL

In this section, we present our simulation theorem, a converse to the application of the revelation principle presented above: i.e., the total expected contributions achievable by a BIC/IIR direct mechanism are also achievable as a BNE of some all-pay contest.

Before we can prove this claim, we first prove a lemma that allows us to make a connection between the interim cost and the interim contribution of a player in any optimal BIC/IIR direct mechanism. Concretely, we show that for all contribution rules $p_i(\mathbf{v})$ that might be dependent on $\mathbf{v}_{-i}$, there exists a corresponding contribution rule $h_i(v_i)$ that is independent of $\mathbf{v}_{-i}$, and that achieves at least as high (quality) total expected contributions.

Consider a randomized direct mechanism, Mechanism A, where $p_i^A(v_i, \mathbf{v}_{-i}, r)$ denotes player i's contribution in Mechanism A and r is the outcome of some randomization device. We define another direct mechanism, Mechanism B, with contribution rule $p_i^B(v_i, \mathbf{v}_{-i}, r) = c_i^{-1}\left(\mathbb{E}_{\mathbf{v}_{-i}, r}\left[c_i\left(p_i^A(v_i, \mathbf{v}_{-i}, r)\right)\right]\right)$. We abbreviate the latter by $h_i(v_i)$, and note that it depends only on v_i.

Lemma 3.1. *An arbitrary allocation rule* $\mathbf{x}$*, together with the corresponding contribution rule* $\hat{\mathbf{p}}^A$ *or* $\hat{\mathbf{p}}^B$*, satisfies BIC and IIR for Mechanism A if and only if it satisfies BIC and IIR for Mechanism B. Moreover, Mechanism B's total expected contributions are at least that of Mechanism A's.*

Proof. We begin with the first part. Observe that the proposed contribution rule preserves interim player costs, and hence, utilities. More concretely:

$$\begin{aligned}\hat{c}_i^B(v_i) &= \mathop{\mathbb{E}}_{\mathbf{v}_{-i},r}\left[c_i\left(p_i^B(v_i,\mathbf{v}_{-i},r)\right)\right] = c_i\left(h_i(v_i)\right)\\ &= c_i\left(c_i^{-1}\left(\mathop{\mathbb{E}}_{\mathbf{v}_{-i},r}\left[c_i\left(p_i^A(v_i,\mathbf{v}_{-i},r)\right)\right]\right)\right) = \hat{c}_i^A(v_i).\end{aligned}$$

Hence, for all players $i \in N$ and abilities $v_i, w_i \in T$, $v_i\hat{x}_i(v_i) - \hat{c}_i^B(v_i) \geq v_i\hat{x}_i(w_i) - \hat{c}_i^B(w_i)$ iff $v_i\hat{x}_i(v_i) - \hat{c}_i^A(v_i) \geq v_i\hat{x}_i(w_i) - \hat{c}_i^A(w_i)$ and $v_i\hat{x}_i(v_i) - \hat{c}_i^B(v_i) \geq 0$ iff $v_i\hat{x}_i(v_i) - \hat{c}_i^A(v_i) \geq 0$.

Now we prove the second part. Since $c_i(\cdot)$ is convex, by Jensen's inequality,

$$\begin{aligned}h_i(v_i) &= c_i^{-1}\left(\mathop{\mathbb{E}}_{\mathbf{v}_{-i},r}\left[c_i\left(p_i^A(v_i,\mathbf{v}_{-i},r)\right)\right]\right)\\ &\geq \mathop{\mathbb{E}}_{\mathbf{v}_{-i},r}\left[c_i^{-1}\left(c_i\left(p_i^A(v_i,\mathbf{v}_{-i},r)\right)\right)\right] = \mathop{\mathbb{E}}_{\mathbf{v}_{-i},r}\left[p_i^A(v_i,\mathbf{v}_{-i},r)\right].\end{aligned}$$

In other words, the contributions in Mechanism B can only exceed those of Mechanism A. Therefore, the total expected contributions of Mechanism B is at least that of Mechanism A. □

This lemma immediately implies, that in our search for an optimal direct mechanism, it suffices to restrict our attention to mechanisms like Mechanism B, in which the contribution of each player i is a (deterministic) function his ability alone:

Corollary 3.2. *For any optimal direct mechanism that is BIC and IIR, there exists a corresponding mechanism that achieves at least as high (quality) total expected contributions, where each player's contribution is a deterministic function of only his value, given by:* $\forall i \in N, \forall v_i \in T$,

$$h_i(v_i) = c_i^{-1}\left(\hat{c}_i(v_i)\right) = c_i^{-1}\left(v_i\hat{x}_i(v_i) - \int_0^{v_i}\hat{x}_i(z_i)\,\mathrm{d}z_i\right). \quad (9)$$

This corollary also allows us to prove our promised simulation theorem, namely that there always exists an all-pay contest and an associated BNE that achieves the total expected contributions of the optimal (BIC and IIR) direct mechanism.

Theorem 3.3 (All-Pay Contest Optimality). *For any optimal direct mechanism that is BIC and IIR, there exists an all-pay contest with a corresponding BNE that achieves the total expected contributions of the optimal direct mechanism.*

We will see later (cf. Lemma 7.1) that Equation (9) leads to an easy calculation of the contributions of the optimal direct mechanism and thereby also the optimal all-pay contest, by Theorem 3.3. However, this connection still does not lead to a clean characterization of the optimal mechanism, which is the outcome of a convex optimization problem, and can be arbitrarily complex. Working towards our main theorem, we propose simple mechanisms that are approximately optimal in Section 5. To establish their approximation ratios, we first prove an upper bound on the total expected contributions of the optimal contest.

4 UPPER BOUND ON OPTIMAL

We first prove an upper bound on the total expected contributions of the optimal mechanism solely based on the IIR property. We first invoke a simple argument from [18] that in the case of a symmetric setting, i.e. cost functions $c_i(\cdot)$ and distributions of abilities are the same for all players, the optimal mechanism must be symmetric.

Lemma 4.1 ([18]). *If all cost functions* $c_i(\cdot)$ *are the same and player ability distributions are identical, then there exists an optimal BIC and IIR mechanism that is symmetric across players.*

Theorem 4.2. *If,* $\forall i \in N$, $c_i(b_i) = b_i^d$*, then* $\mathrm{OPT} \leq n\left(\frac{\mu}{n}\right)^{1/d}$*, for* $d \geq 2$.

Proof. By Lemma 4.1 it suffices to look at symmetric mechanisms across players. Moreover, by Corollary 3.2, it suffices to upper bound the expected contributions of a mechanism whose contribution rule is of the form given in Equation (9) (observe that crucially Lemma 3.1 preserves symmetry of the original mechanism when costs and distributions of abilities are identical across players). Hence, it suffices to look at symmetric mechanisms whose contribution rule is only a function of the player's ability. For any such mechanism, we have that by IIR:

$$h_i(v_i) \leq c_i^{-1}(v_i\hat{x}_i(v_i)) = v_i^{1/d}\hat{x}_i(v_i)^{1/d} \quad (10)$$

Thus the expected contribution of a player is upper bounded by:

$$\mathop{\mathbb{E}}_{v_i}\left[h_i(v_i)\right] \leq \mathop{\mathbb{E}}_{v_i}\left[v_i^{1/d}\hat{x}_i(v_i)^{1/d}\right] \quad (11)$$

By Cauchy-Schwarz inequality:

$$\mathop{\mathbb{E}}_{v_i}\left[v_i^{1/d}\hat{x}_i(v_i))^{1/d}\right] \leq \sqrt{\mathop{\mathbb{E}}_{v_i}\left[v_i^{2/d}\right]}\sqrt{\mathop{\mathbb{E}}_{v_i}\left[\hat{x}_i(v_i)^{2/d}\right]} \quad (12)$$

When $d \geq 2$, the function $x^{2/d}$ is concave for $x \geq 0$. Hence, by Jensen's inequality:

$$\mathop{\mathbb{E}}_{v_i}\left[v_i^{2/d}\right] \leq \mathop{\mathbb{E}}_{v_i}\left[v_i\right]^{2/d} = \mu^{2/d} \quad (13)$$

$$\mathop{\mathbb{E}}_{v_i}\left[\hat{x}_i(v_i)^{2/d}\right] \leq \mathop{\mathbb{E}}_{v_i}\left[\hat{x}_i(v_i)\right]^{2/d} \quad (14)$$

By symmetry of the mechanism we have that $\forall i : \mathbb{E}_{v_i}\left[\hat{x}_i(v_i)\right] = \kappa$, for some constant κ. By feasibility, then we have $\kappa = 1/n$, since $\sum_i \mathbb{E}_{v_i}\left[\hat{x}_i(v_i)\right] = \mathbb{E}_{\mathbf{v}}\left[\sum_i x(\mathbf{v})\right] = 1$. Combining all the above we have:

$$\mathop{\mathbb{E}}_{v_i}\left[h_i(v_i)\right] \leq \mu^{1/d}\frac{1}{n^{1/d}} \quad (15)$$

The theorem follows since $\mathrm{OPT} = \sum_i \mathbb{E}_{v_i}\left[h_i(v_i)\right]$. □

Further Refinement for MHR Distributions. We conclude this section by providing a more useful upper bound on the optimal contributions for the case of Monotone Hazard Rate (MHR) distributions. To do so, we introduce some useful notation and terminology with respect to properties of the distribution F. For any distribution F, let $q(v) = 1 - F(v)$ be the quantile function, and let $v(q) = q^{-1}(\cdot)$ be the inverse quantile function. In words, the quantile of an ability v is the probability that a random draw from F exceeds v. Observe that quantiles are distributed uniformly on $[0, 1]$. We will also denote by $\kappa = v(1/2)$ the median of the distribution.

In the usual quasi-linear setting, the **revenue function** $R(q)$ associated with F is defined as $R(q) = v(q)q = v(q)(1 - F(v(q)))$.

Intuitively, this is the expected revenue of a seller who posts a reserve price of $v(q)$ (i.e., one that is surpassed with probability q). Since, F is atomless with support $[0, \bar{v}]$, $R(0) = R(1) = 0$.

We will be looking at two standard classes of distributions. The smaller class is that of monotone hazard rate (MHR) distributions, which require that $h(v) = f(v)/(1 - F(v))$ be monotone non-decreasing. The larger class is that of regular distributions, which require that $R(q)$ be a concave function, or equivalently $\phi(v(q)) = R'(q) = v(q) - \frac{1-F(v(q))}{f(v(q))}$, be monotone decreasing. Since $\phi(v(q)) = v(q) - \frac{1}{h(v(q))}$, it is easy to see that an MHR distribution is also regular. MHR distributions contain several well-known families, such as the uniform and the exponential (see e.g. [4]).

We now state two lemmas describing bounds on revenue curves, depending on the assumptions made on the ability distributions.

Lemma 4.3 ([10, 26]). *For any regular distribution, $R(q^*) \leq \kappa$. For any MHR distribution, $R(q^*) \geq \frac{\mu}{e}$.*

This lemma implies that for MHR distributions: $\mu \leq eR(q^*) \leq e\kappa$. Combining this result with Theorem 4.2, we conclude:

Corollary 4.4. *If the ability distribution F is an MHR distribution then,* $\text{OPT} \leq n\left(\frac{e\kappa}{n}\right)^{1/d}$.

5 QUANTILE THRESHOLD MECHANISMS

We now turn to the design of simple detail-free mechanisms and contests. We begin our analysis by looking at *quantile threshold mechanisms*, i.e., mechanisms that distribute the prize uniformly among all players whose ability v_i is above some ability threshold r, or equivalently, to players whose quantile q_i is below some quantile threshold $\hat{q}$. In these mechanisms, the player is asked to contribute based on the characterization in Equation (9), which renders the mechanism BIC and IIR. In the case of quantile threshold mechanisms, this contribution takes the simple form:

$$h_i(v_i) = c_i^{-1}\left(v(\hat{q})\mathbb{1}_{v_i \geq v(\hat{q})} \mathop{\mathbb{E}}_{\mathbf{v}_{-i}}\left[\frac{1}{1 + \sum_{j\neq i} \mathbb{1}_{v_j \geq v(\hat{q})}}\right]\right) \tag{16}$$

All-pay implementation. By the same argument as in Theorem 3.3, a quantile threshold mechanism with quantile $\hat{q}$ can be implemented as an all-pay contest. In fact, for quantile threshold mechanisms, this contest has a simple form. We omit the proof of the following lemma, since it is similar to that of Theorem 3.3.

Lemma 5.1 (Contest Implementation of Quantile Threshold Mechanism). *Consider an all-pay contest that distributes the prize uniformly to all players i who contribute b_i above:*

$$b_i^* = c_i^{-1}\left(v(\hat{q}) \mathop{\mathbb{E}}_{\mathbf{v}_{-i}}\left[\frac{1}{1 + \sum_{j\neq i} \mathbb{1}_{v_j \geq v(\hat{q})}}\right]\right). \tag{17}$$

There exists a BNE of this all-pay contest that implements the same contributions as a quantile threshold mechanism with threshold $\hat{q}$.

When all players have the same cost function, and abilities are drawn from indepdendent and identical distributions, then this all-pay contest sets the same contribution threshold for all players, and takes an intuitive form that is prevalent in practice (e.g., badge mechanisms [2]). We also note that this threshold mechanism might afford multiple equilibria, but only one of them implements the same outcome as the original mechanism. In the next section, we will construct approximately optimal contests with unique equilibria.

We begin with the analysis of the expected contributions of a *quantile threshold mechanism*, with some quantile threshold $\hat{q}$, in the case where players have cost functions of the form b^d.

Lemma 5.2. *Consider a convex cost of the form $c_i(b) = b^d$, $\forall i \in N$, where $d \geq 1$. Let $\text{APX}(\hat{q})$ be the expected contributions in a mechanism that allocates uniformly across all players with quantile $q_i \leq \hat{q}$ and requires contributions given by Equation (16). Then:*

$$\text{APX}(\hat{q}) \geq n\left(\frac{v(\hat{q})}{1 + (n-1)\hat{q}}\right)^{1/d} \hat{q} \geq n^{1-1/d} v(\hat{q})^{1/d} \hat{q}. \tag{18}$$

Proof. By the contribution identity (Equation 16), since $1/(1 + x)$ is a convex function, by Jensen's inequality, and because the probability that $v_i \geq v(\hat{q})$ is $1 - F(v(\hat{q})) = \hat{q}$:

$$h_i(v_i) \geq \left(\frac{v(\hat{q})\mathbb{1}_{v_i \geq v(\hat{q})}}{1 + \sum_{j\neq i} \mathbb{E}_{\mathbf{v}_j}[\mathbb{1}_{v_j \geq v(\hat{q})}]}\right)^{1/d} = \left(\frac{v(\hat{q})\mathbb{1}_{v_i \geq v(\hat{q})}}{1 + (n-1)\hat{q}}\right)^{1/d},$$

By the definition of expected contributions:

$$\text{APX}(\hat{q}) = \sum_{i=1}^{n} \mathop{\mathbb{E}}_{v_i \sim F}\left[\left(\frac{v(\hat{q})\mathbb{1}_{v_i \geq v(\hat{q})}}{1 + (n-1)\hat{q}}\right)^{1/d}\right] = \sum_{i=1}^{n}\left(\frac{v(\hat{q})}{1 + (n-1)\hat{q}}\right)^{1/d} \hat{q}.$$

We have thus established the first part of the Equation (18). The second part follows by noting that $\hat{q} \in [0, 1]$. ◻

Detail-free mechanisms. We now show that if the designer has some knowledge of the distribution of abilities, then mechanisms can generate more expected contributions than otherwise. Specifically, knowing the median of the distribution is sufficient for getting a much better approximation ratio. We conclude this section by providing two examples of detail-free mechanisms that optimize over the quantile threshold $\hat{q}$ to obtain better approximation ratios.

Theorem 5.3 (Median Threshold). *The approximation ratio of the mechanism with median ability threshold κ (i.e., $\hat{q} = \frac{1}{2}$) when, $\forall i \in N$, $c_i(b) = b^d$, where $d \geq 2$ with MHR distributions, is:*

$$\frac{\text{APX}}{\text{OPT}} \geq \frac{1}{2}\left(\frac{2n}{e(n+1)}\right)^{1/d}. \tag{19}$$

The median mechanism can be implemented at a BNE of the all-pay contest which allocates uniformly to all players who contribute above:

$$b^* = \left(\kappa \frac{2 - 2^{n-1}}{n}\right)^{1/d} \tag{20}$$

Proof. Lemma 5.2 tells us that when $\hat{q} = \frac{1}{2}$,

$$\text{APX} \geq n\left(\frac{\kappa}{1 + (n-1)/2}\right)^{1/d} \frac{1}{2} = \frac{n}{2}\left(\frac{2\kappa}{n+1}\right)^{1/d},$$

and by Theorem 4.4,

$$\frac{\text{APX}}{\text{OPT}} \geq \frac{n}{2}\left(\frac{2\kappa}{n+1}\right)^{1/d} \frac{1}{n^{(d-1)/d}(e\kappa)^{1/d}} = \frac{1}{2}\left(\frac{2n}{e(n+1)}\right)^{1/d}.$$

The last part of the theorem follows from Lemma 5.1, and by observing that for a median quantile:

$$\mathop{\mathbb{E}}_{\mathbf{v}_{-i}}\left[\frac{1}{1+\sum_{j\neq i}\mathbb{1}_{v_j\geq v(\hat{q})}}\right]=\frac{1}{2^{n-1}}\sum_{t=0}^{n-1}\binom{n-1}{t}\frac{1}{t+1}=\frac{2-2^{1-n}}{n}$$

□

Equation (19) tells us that for a large number of players, the approximation ratio is $\frac{1}{2}\left(\frac{2}{e}\right)^{1/d}$. As $d \geq 2$, this is at least 0.42. Further, as d tends towards infinity, the approximation ratio approaches $\frac{1}{2}$.

Next, we claim that if we use a quantile threshold that is also dependent on degree of complexity d of the cost function. then we can further improve the approximation ratio. The proof and analysis of this theorem is deferred to Appendix A.2.

Theorem 5.4 (Cost-Optimized Threshold). *Assume that the ability distributions satisfy the Monotone Hazard Rate (MHR) condition, and that the cost function for each player $i \in N$ is of the form $c_i(b) = b^d$, for $d \geq 2$. Then a quantile threshold mechanism with threshold $\hat{q} = \max\left\{\frac{1}{2}, 1 - \frac{1}{d-1}\right\}$ achieves contributions that are at least the following fractions of optimal: $\left(\frac{n}{n+1}\right)^{1/d}\frac{1}{2\sqrt{e}}$ for $d \in [2,3)$ and $\left(\frac{n}{n+1}\right)^{1/d}\frac{1}{(4e(d-2))^{1/d}}$ for $d \geq 3$.*

Observe that for $d \to \infty$, this bound converges to 1, since $(d-2)^{1/d} \to 1$. Thus, as the payment functions become more convex, the cost-optimized reserve mechanism converges to full optimality. Moreover, like the median threshold mechanism, the cost-optimized threshold mechanism can also be implemented as the BNE of an all-pay contest with an appropriately chosen contribution threshold.

We have seen that simple contests that allocate the prize uniformly to all players who contribute above a certain threshold achieve a constant factor approximation. Although they are detail-free, in the sense that they do not require full distributional knowledge, they require some knowledge of the distribution or of the cost function in order to be implemented. In the next section, we will completely remove the dependence on the distribution and the cost function at the expense of only assuming that the number of contributors is larger than some constant.

Remark 1 (Implications for Revenue Maximization). *Before we move on to our endeavor for robust optimal contests we remark that the results in this section have implications for the revenue maximization problem with non quasi-linear utilities. In particular, if we replace contributions with payments, then the optimal mechanism design problem for our setting is a revenue maximization problem where players incur costs that are a convex function of their payments. Under this interpretation, Theorem 5.3 implies that allocating uniformly at random to all players whose value is above the median of the value distribution is a constant approximation to the optimal revenue. Using similar techniques, we can also show that rather than choosing the median threshold, choosing as a threshold a value that is drawn randomly from the value distribution also leads a constant factor approximation. Notably, this mechanism can be implemented in a completely prior-free manner, where we use one player as a the threshold setter, and allocate uniformly at random to all remaining players whose value is above the value of this threshold setter. This result is an analogue of the classic result of Bulow and Klemperer [6] for our setting. Finally, all these mechanisms can be implemented in a manner that renders truthful reporting an ex-post dominant strategy, rather than simply ensuring that the mechanism is BIC. Due to space constraints, and since these results are tangential to our main goal, we defer them to the full version of the paper.*

6 MAIN RESULT: PRIOR AND COST INDEPENDENT CONTEST

In this section we present our main result: an all-pay contest which is approximately optimal and which does not depend on the distribution of abilities or the degree of convexity d of the cost function.

We will start by drawing intuition from our quantile threshold mechanisms and their all-pay contest implementations. We saw that these contests are approximately optimal but are faced with a difficulty: the appropriate contribution threshold to set in the contest so that it translates to a target quantile threshold in the quantile space of abilities is dependent on the exponent d. Therefore, it seems unlikely that there exists a universal contribution threshold that would work for all d. A further difficulty stems from the fact that contests in which the allocation depends on the absolute value of players' contributions tend to not have unique equilibria [8, 20].

Given these issues, we instead investigate allocation functions in which each player's allocation depends only on the rank of his contribution relative to those of the other players. The key idea is that we can try to approximate the median threshold mechanism with such a relative ranking mechanism by allocating uniformly to the top half of the contributors and nothing to the bottom half. Assuming that players contribute according to a symmetric Bayes-Nash equilibrium with a strictly monotone contribution function $b(\cdot)$ (which is the unique equilibrium based on the results of [8]), then at equilibrium the contributors whose ability is among the top 50% of abilities will be allocated a prize of $2/n$.[5]

If the number of contributors is at least some constant, then this interim allocation function will strongly resemble the interim allocation function of the median threshold mechanism: contributors with ability strictly above the median ability by some error margin ϵ will be allocated $2/n$ with probability approaching 1, while players whose ability is strictly below the median ability by some error margin ϵ will get an allocation that is near 0. These results follow from a concentration of measure inequality argument. Hence, this mechanism will achieve an approximation ratio APX/OPT that is of the form $\frac{1}{\alpha} - O\left(\sqrt{\frac{\log(n)}{n}}\right)$, where α is some small constant, since the median mechanism is approximately optimal.

In order to avoid dependence of our analysis on a lower bound on the density of the distribution of abilities, we actually propose a more competitive all-pay contest where we allocate to a smaller fraction of the contributors. In the theorem that follows we pick the top quarter, but we note that the constants could be further optimized by picking a more involved fraction. The proof of this theorem appears in Appendix A.3.[6]

[5] For simplicity, we assume that n is even.

[6] We also note that the constants in the analysis of this theorem could be optimized to obtain a better result. For simplicity of exposition we omit such optimization, as the main point of the theorem is that a constant approximation can be achieved by a mechanism that is independent of the cost function.

Finally, we note that if the designer has an idea about the degree of convexity d of the cost function, then, drawing intuition from the cost-optimized threshold mechanism, we might expect a less competitive contest to perform better. We revisit this intuition in the experimental section and indeed show improved performance of such convexity-tailored contests.

THEOREM 6.1 (MAIN THEOREM). *Consider an all-pay contest which allocates the prize uniformly to the top quarter of the contributors, i.e.: $x_i(b_i, \mathbf{b}_{-i}) = \frac{4}{n}$ if b_i is among the top $\frac{n}{4}$ highest contributors. Assuming that $c_i(\cdot)$ is a strictly monotone function, and that it is the same for all players, and further assuming that the distribution of abilities is atomless and has a continuous CDF with support $[0, \bar{v}]$, then at the unique Bayes-Nash equilibrium of this contest, each player will choose a contribution $b_i(v_i)$ such that*

$$b_i(v_i) = c_i^{-1}\left(v_i\hat{x}_i(v_i) - \int_0^{v_i} \hat{x}_i(z_i)\,\mathrm{d}z_i\right), \tag{21}$$

where $\hat{x}_i$ is the interim prize allocation that corresponds to awarding an allocation of $\frac{4}{n}$ to the $\frac{n}{4}$ players of highest ability, i.e.,

$$\hat{x}_i(v_i) = \frac{4}{n}\Pr\left(v_i \text{ is among } \tfrac{n}{4} \text{ highest abilities} \mid v_i\right). \tag{22}$$

Finally, assume that $c_i(b) = b^d$ for $d \geq 2$, that $n \geq 32\log(16\bar{v}/\kappa)$, and that the distribution of abilities is MHR. Then this contest achieves expected contributions APX *at the unique Bayes-Nash equilibrium, which satisfy: $\frac{\text{APX}}{\text{OPT}} \geq \frac{1}{16}$.*

7 EXPERIMENTS

In this section, we provide empirical evidence that our proposed contests yield near-optimal performance. We first show that in a symmetric setting, it is possible to compute the optimal allocation and contributions in polynomial time. The characterization makes use of Border's theorem [5] to reduce the number of interim feasibility constraints to a manageable size; we refer the reader to [18, Section 8.5.1] for an elaborate description of this approach. As the following lemma is an adaptation to our setting of the results in [18], we omit the proof. Intuitively the function $z(\cdot)$ in the program corresponds to the derivative of the interim allocation rule of the direct mechanism.

LEMMA 7.1. *In a symmetric setting with convex costs, where values are drawn from an atom-less regular distribution F with support in $T = [0, 1]$ and density function f, the contributions of the optimal contest can be described by the following program:*

$$\max_{z(\cdot)\geq 0} \int_0^1 f(t)\cdot c^{-1}\left(\int_0^t \tau z(\tau)\,\mathrm{d}\tau\right)\,\mathrm{d}t \tag{23}$$

$$\text{s.t.} \int_0^1 z(\tau)\left(1 - F\left(\max\{t,\tau\}\right)\right)\mathrm{d}\tau \leq \frac{1-F(t)^n}{n}, \quad \forall t \in T \tag{24}$$

Experiments and Results. We conclude this paper with an investigation of how well our proposed contests perform empirically, compared to our theoretical guarantees. In particular, we examine the performance of an all-pay contest that distributes the prize uniformly to the top quartile of contributors (Theorem 6.1) and the cost-optimized version of this contest which distributes uniformly to the top $\hat{q}$ fraction of contributors for $\hat{q} = \max\left\{\frac{1}{2}, 1 - \frac{1}{d-1}\right\}$ (the latter is the relative contribution analogue of the cost-optimized threshold mechanism of Theorem 5.4). Moreover, we examine the performance of the less robust but better performing contests that rely on a contribution threshold, namely the all-pay implementation of the median mechanism 5.3 and the all-pay implementation of the cost-optimized threshold mechanism 5.4.

To compare these mechanisms, we generated 50 random MHR distributions (for computational reasons we considered discretized versions of the distributions). For each distribution, we simulated 10,000 contests, and then computed the average contributions of each, for each of $n = \{1, \ldots, 300\}$ symmetric players. Each player had cost $c_i(b) = b^d$, for varying d. We computed the mean contributions across distributions as well as the 5% and 95% percentiles of their performance (i.e., the ratio between the contribution achieved and the optimal). We summarize the results of the experiments in Figure 1. In all cases, the experimental results suggest that our proposed mechanisms may perform much better in practice than our theoretical worst-case guarantees.

8 CONCLUSION

In this work, we studied the problem of optimal contest design, where the designer's objective is to maximize the total expected quality of contributions assuming participants incur convex costs. This convex-cost model is a natural model of effort, as it takes more and more effort to improve the quality of a contribution as that quality increases. We began by establishing an equivalence between all-pay contests and optimal direct mechanisms, in which players report their values, and are possibly allocated a prize and asked to contribute some amount. On the other hand, we also noted that Myerson's theory of optimal mechanism design is not fully applicable in our setting. Indeed, a closed-form characterizaton of the optimal BIC/IIR mechanism with convex costs remains open.

Nonetheless, we were able to establish an upper bound on the optimal, which we used to derive approximation ratios for two simple and detail-free all-pay contest mechanisms, one of which allocates the prize uniformly to all players whose ability is above the median ability (in the direct setting). We then noted that the spirit of the median mechanism can be achieved in the contest setting by allocating the prize uniformly to the top half of the contributors and nothing to the others. Indeed, our main theorem, for which we achieve a constant approximation, allocates uniformly at random to the top quarter of the contributors and nothing to the others.

Our main result stands in stark contrast to the optimal mechanism in the usual quasi-linear setting, where no one but the top contributors are ever allocated any portion of the prize. When costs are convex, and when the goal is to maximize the total quality of the contributions (i.e., total productivity), it can be very suboptimal to allocate only to the top contributors. If we accept the premise that the behavior of many Americans' is governed by the 80/20 rule, our results suggest a possible reason for the recent stagnant/slow growth in the United States economy, namely that the current system does not offer enough incentives to enough participants.

A OMITTED PROOFS

A.1 Proof of Lemma 3.3

Consider a direct mechanism with an ex-post allocation rule $x_i : T^n \to [0, 1]$ and a deterministic contribution rule $h_i : T \to \mathbb{R}_{\geq 0}$ for each player i. We will implement the direct mechanism as an

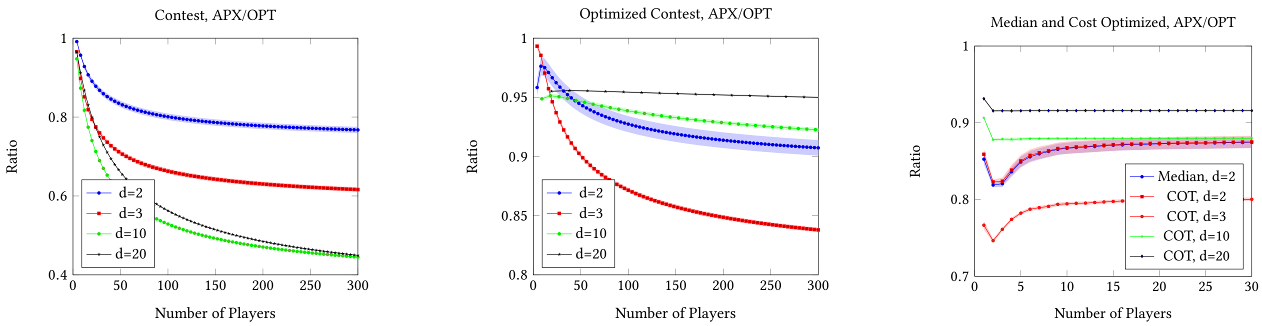

Figure 1: Empirical performance of proposed contests on random MHR distributions. The figure on the left shows the performance of a contest that allocates to the top $1/4$ players; the middle figure shows the contest that allocates to top $\max\left\{\frac{1}{2}, 1-\frac{1}{(d-1)}\right\}$ players; and the right figure shows the threshold contests (Theorems 5.3 and 5.4).

all-pay contest as follows: given the contribution of a player b_i, use the interim contribution rule of the direct mechanism to map it back to the set of abilities that contribute b_i in the direct mechanism, i.e. $S_i(b_i) = \{w_i \in T : h_i(w_i) = b_i\}$. Since the interim contribution rule is non-decreasing, the latter set is an interval $[L(b_i), U(b_i)]$ and moreover these intervals are disjoint for different b_i and are consequtive. If the player submits a contribution that is not in the range of $h_i(\cdot)$, i.e. if $S_i(b_i)$ is empty, then we map it to an ability of 0 and allocate 0 prize. Otherwise, we draw an ability from distribution F, conditional on the fact that the ability is in $S_i(b_i)$, i.e. if $S_i(b_i)$ is a singleton then we return this singleton otherwise we draw an ability from the distribution with support $S_i(b_i)$ and CDF $G(v; b_i) = \frac{F(v)-F(L(b_i))}{F(U(b_i))-F(L(b_i))}$. Let z_i be the resulting ability drawn from all this process for player i. Then allocate to player i a prize of $x_i(\mathbf{z})$ based on the ex-post allocation rule of the direct mechanism for the constructed ability vector $\mathbf{z}$.

We will show that all players contributing $b_i(v_i) = h_i(v_i)$ is a BNE of the resulting all-pay contest. The latter would imply that the all-pay contest achieves the same expected contributions as the direct mechanism.

We first note that if all other players bid according to this BNE, then if a player i is assigned a value z_i from the aforementioned process, then his interim allocation in the all-pay contest $\mathbb{E}_{\mathbf{z}_{-i}}[x_i(z_i, \mathbf{z}_{-i})]$ (i.e. his expected allocation in expectation over the draws of his opponents abilities and the randomness of the contest) is equal to $\hat{x}(\mathbf{z}_i)$, i.e. it is equal to the interim allocation of a player with ability z_i in the original direct mechanism. To argue this we only need to argue that the distribution of z_j for any $j \neq i$ follows exactly distribution F. Then the claim will follow by the definition of the interim allocation rule $\hat{x}(\cdot)$ of the direct mechanism. The fact that each z_j is drawn from F, follows by our construction; the distribution of the random variable z_j is the result of the following process: first a value v_j is drawn from F, then a bid $b_j(v_j) = h_j(v_j)$ is generated. Then the random variable z_j is drawn from the distribution of F constrained on the set $S_j(h_j(v_j))$. Let us look at the CDF of z_j: Consider a value $x \in T$ and let $S_j(t) = [L(t), U(t)]$ be the interval containing x in the partition of abilities defined by h_j (i.e. the partition of T such that all abilities in each part make the same contribution under rule $h_j(\cdot)$). The probability that z_j falls below x is the probability that the player j has a true ability v_j that falls below $L(t)$, plus the probability that it falls in interval $S_j(t)$, times the probability that z_j drawn from distribution F conditional on interval $S_j(t)$ falls below x. The latter is $F(L(t)) + \frac{F(x)-F(L(t))}{F(U(t))-F(L(t))}(F(U(t)) - F(L(t))) = F(x)$. This concludes our claim.

Second, observe that if two abilities $w_i > v_i$, have the same deterministic contribution under h_i, i.e. $h_i(w_i) = h_i(v_i)$, then it must be that their interim allocations in the direct mechanism are the same, i.e. $\hat{x}_i(w_i) = \hat{x}_i(v_i)$. The latter follows by the strict monotonicity of the cost functions and the fact that the quantity inside $c_i^{-1}(\cdot)$ in Equation (9) is strictly larger for ability w_i than for ability v_i if $\hat{x}_i(w_i) > \hat{x}_i(v_i)$. Hence, we can conclude that if a player submits a contribution $b_i = h_i(w_i)$ in the all-pay contest for some ability w_i, then no matter what was the draw z_i of the random process of the contest, the expected prize allocation of the player in expectation over the abilities of opponents and the randomness of the contest, will be equal to $\hat{x}_i(w_i)$.

Thus the expected utility of a player with ability v_i, in the all-pay contest, if he submits a contribution $b_i = h_i(w_i)$ is equal to the expected utility of that player in the direct mechanism if he reported an ability level of w_i, i.e. $v_i\hat{x}_i(w_i) - h_i(w_i)$. Thus, by the BIC condition of the direct mechanism, the player in this all-pay contest does not want to deviate and contribute the amount that a player with any other ability would have contributed, since this would be equivalent to reporting that ability in the direct mechanism. Finally, a player does not want to submit a contribution b_i that is not in the range of $h_i(\cdot)$ since that yields zero utillity. Thus each player contributing $b_i(v_i) = h_i(v_i)$ is a BNE of the all-pay contest.

A.2 Proof of Theorem 5.4

We will first prove a Lemma which is an extension of the simplified prophet inequality that we used in the previous section.

Lemma A.1. *For any regular distribution and for any $\hat{q} \geq 1/2$:*

$$R(\hat{q}) \geq (1-\hat{q})R(q^*) \tag{25}$$

Proof. First suppose that $q^* \geq \hat{q}$. Then we know that $R(q)$ is increasing until η and concave. Hence:

$$R(\hat{q}) \geq \frac{R(q^*) - R(0)}{q^*}\hat{q} + R(0) \geq R(q^*)\hat{q} + R(0)(1-\hat{q}) \geq R(q^*)\hat{q}$$

$$\geq R(q^*)\frac{1}{2} \geq R(q^*)(1-\hat{q})$$

Now suppose that $q^* \leq \hat{q}$. Then $R(q)$ is decreasing in the region $(q^*, 1]$ and $R(1) = 0$. Thus by concavity of $R(q)$, since $\hat{q} \in (q^*, 1]$:

$$R(\hat{q}) \geq \frac{R(q^*) - R(1)}{q^* - 1}(\hat{q} - 1) + R(1) = \frac{R(q^*)}{1-q^*}(1-\hat{q}) \geq R(q^*)(1-\hat{q})$$

□

Now we move on to proving the Theorem. For $d \leq 3$, the threshold that we use is the median and hence the bound from Theorem 5.3 applies. So we prove the bound for the case of $d \geq 3$ and $\hat{q} = 1 - \frac{1}{d-1}$. By Lemma 4.2, Lemma 4.3 and Lemma A.1 we have:

$$\text{OPT} \leq n^{\frac{d-1}{d}} \mu^{1/d} \leq n^{\frac{d-1}{d}} (eR(q^*))^{1/d} \leq n^{\frac{d-1}{d}} \left(e \cdot \frac{\hat{q} \cdot v(\hat{q})}{1-\hat{q}}\right)^{1/d}$$

On the other hand by Lemma 5.2, we can bound the contributions of the quantile threshold mechanism, with quantile $\hat{q} = 1 - \frac{1}{d-1} \geq \frac{1}{2}$:

$$\text{APX} \geq \frac{n \cdot \hat{q} \cdot v(\hat{q})^{1/d}}{(1-\hat{q}+n\cdot\hat{q})^{1/d}} \geq \frac{n \cdot \hat{q} \cdot v(\hat{q})^{1/d}}{(\hat{q}+n\cdot\hat{q})^{1/d}} = \frac{n}{(n+1)^{1/d}} \frac{\hat{q} \cdot v(\hat{q})^{1/d}}{\hat{q}^{1/d}}$$

Thus the ratio of the two mechanisms is at least:

$$\frac{\text{APX}}{\text{OPT}} = \left(\frac{n}{n+1}\right)^{1/d} \frac{1}{e^{1/d}} \left(\hat{q}^{d-1} \frac{1-\hat{q}}{\hat{q}}\right)^{1/d}$$
$$= \left(\frac{n}{n+1}\right)^{1/d} \frac{1}{e^{1/d}} \left(\left(1 - \frac{1}{d-1}\right)^{d-1} \frac{1}{d-2}\right)^{1/d}$$

Since $(1-1/x)^x \geq \frac{1}{4}$ for $x \geq 2$, simplifying the latter yields the theorem.

A.3 Proof of Theorem 6.1

The fact that the given bidding function is an equilibrium follows from the fact that it is a monotone function of the valuation of each player. Therefore, it gives rise to the interim allocation described in the theorem. Moreover, the pair of the bid function and interim allocation satisfy Myerson's payment identity. Finally, no player wants to bid outside of the support of the bid distribution. Hence, from standard analysis that can be found in [18], this implies that the proposed pair of bid function and interim allocation constitute a Bayes-Nash equilibrium. Uniqueness of this equilibrium follow from the results of [8]. We now move on to analyzing the contributions of this equilibrium of the all-pay contest. For simplicity we will assume that the number of players is a multiple of 4. The result easily extends to the general case, albeit with more complex notation.

We first prove upper and lower bounds on the interim allocation of players as a function of their quantile $q_i = 1 - F(v_i)$. For any quantile q, denote with $X_j(q) = 1\{q_j \leq q\}$. Observe that: $\mathbb{E}[X_j(q)] = \Pr[q_j \leq q] = q$, since quantiles are distributed uniformly in $[0,1]$. Moreover, if we denote with $\hat{x}_i(q_i)$ the interim allocation of player i when he has quantile q_i, then:

$$\hat{x}_i(q_i) = \tfrac{4}{n} \Pr\left[\textstyle\sum_{j\neq i} X_j(q_i) \leq \tfrac{n}{4} - 1\right] \tag{26}$$

Let $S_{n-1}(q_i) = \sum_{j\neq i} X_j(q_i)$. Since S_{n-1} is the sum of $n-1$ independent 0/1 random variables, each with success probability q, we get by the Chernoff bound that, for any $\epsilon > 0$:

$$\Pr[(q_i - \epsilon)(n-1) \leq S_{n-1}(q_i) \leq (q_i+\epsilon)(n-1)] \geq 1 - 2\exp\{-2\epsilon^2 n\}$$

Thus we have that for q_i, such that: $(q_i+\epsilon)(n-1) \leq \frac{n}{4} - 1$:

$$\hat{x}_i(q_i) = \frac{4}{n}\Pr[S_{n-1}(q_i) \leq \frac{n}{4} - 1] = \frac{4}{n}\left(1 - \Pr[S_{n-1}(q_i) > \frac{n}{4} - 1]\right)$$
$$\geq \frac{4}{n}\left(1 - \Pr[S_{n-1}(q_i) > (q_i+\epsilon)(n-1)]\right) \geq \frac{4}{n}\left(1 - 2e^{-2\epsilon^2 n}\right)$$

Re-arranging the condition on q_i, we get that the latter bound on the interim allocation holds for $q_i \leq \frac{1}{4} - \frac{3}{4}\frac{1}{n-1} - \epsilon$.

Similarly, we have that for q_i, such that: $(q_i - \epsilon)(n-1) > \frac{n}{4} - 1$:

$$\hat{x}_i(q_i) = \frac{4}{n}\Pr[S_{n-1}(q_i) \leq \frac{n}{4} - 1]$$
$$\leq \frac{4}{n}\Pr[S_{n-1}(q_i) \leq (q_i - \epsilon)(n-1)] \leq \frac{8}{n}\exp\{-2\epsilon^2 n\}$$

Re-arranging the condition on q_i, we get that the latter bound on the interim allocation holds for $q_i > \frac{1}{4} - \frac{3}{4}\frac{1}{n-1} + \epsilon$. Finally, we know that the interim allocation $\hat{x}_i(q_i)$ is non-increasing in q_i.

Consider the interim cost of a player as a function of his quantile, which by transforming Myerson's identity to quantile space and observing that $v_i'(z) = -|v_i'(z)|$, takes the form:

$$c_i(b_i(q_i)) = v_i(q_i)\hat{x}_i(q_i) - \int_{q_i}^1 \hat{x}_i(z)|v_i'(z)|dz$$

For ease of notation, let $\theta_+ = \frac{1}{4} - \frac{3}{4}\frac{1}{n-1} + \epsilon$ and $\theta_- = \frac{1}{4} - \frac{3}{4}\frac{1}{n-1} - \epsilon$. We will lower bound the interim cost of a player with $q_i \leq \theta_-$:

$$c_i(b_i(q_i)) = v_i(q_i)\hat{x}_i(q_i) - \int_{q_i}^{\theta_+} \hat{x}_i(z)|v_i'(z)|dz - \int_{\theta_+}^1 \hat{x}_i(z)|v_i'(z)|dz$$
$$\geq v_i(q_i)\hat{x}_i(q_i) - \hat{x}_i(q_i)\int_{q_i}^{\theta_+}|v_i'(z)|dz - \frac{8e^{-2\epsilon^2 n}}{n}\int_{\theta_+}^1 |v_i'(z)|dz$$
$$\geq v_i(\theta_+)\,\hat{x}_i(q_i) - \frac{8e^{-2\epsilon^2 n}}{n}\int_{\theta_+}^1|v_i'(z)|dz$$
$$\geq v_i(\theta_+)\frac{4}{n}(1 - 2e^{-2\epsilon^2 n}) - \frac{8e^{-2\epsilon^2 n}}{n}\bar{v}$$
$$\geq \frac{4}{n}v_i\left(\frac{1}{4}+\epsilon\right) - \frac{16}{n}e^{-2\epsilon^2 n}\bar{v}$$

Since a player has quantile smaller than θ_- with probability equal to θ_-, and since interim cost is non-increasing in quantile, we can lower bound the ex-ante expected contribution of each player:

$$\mathbb{E}[b_i(q_i)] \geq \theta_- \cdot c_i^{-1}\left(\frac{4}{n}v_i\left(\frac{1}{4}+\epsilon\right) - \frac{16}{n}e^{-2\epsilon^2 n}\bar{v}\right)$$

By picking $\epsilon = \sqrt{\frac{\log(16\bar{v}/\kappa)}{2n}}$ (where κ is the median), we get:

$$\mathbb{E}[b_i(q_i)] \geq \theta_- \cdot c_i^{-1}\left(\frac{4}{n}v_i\left(\frac{1}{4}+\epsilon\right) - \frac{\kappa}{n}\right)$$

Assuming that $\sqrt{\frac{\log(16\bar{v}/\kappa)}{2n}} \leq \frac{1}{8}$, which happens if $n \geq 32\log(16\bar{v}/\kappa)$, we have that: $v_i\left(\frac{1}{4}+\epsilon\right) \geq v_i(1/2) = \kappa$ and $\theta_- \geq \frac{1}{8} - \frac{3}{4}\frac{1}{n-1}$. Hence:

$$\mathbb{E}[b_i(q_i)] \geq \left(\frac{1}{8} - \frac{3}{4}\frac{1}{n-1}\right)\cdot c_i^{-1}\left(\frac{3}{n}\kappa\right)$$

Further assuming that $\frac{3}{4}\frac{1}{n-1} \leq \frac{1}{16} \Leftrightarrow n \geq 13$ (which holds whenever $n \geq 32\log(16\bar{v}/\kappa)$), we get that: $\mathbb{E}[b_i(q_i)] \geq \frac{1}{16}\cdot c_i^{-1}\left(\frac{3}{n}\kappa\right)$.

Combining the above we get for $n \geq 32\log(16\bar{v}/\kappa)$, the contributions of the all-pay contest is lower bounded by: $\text{APX} \geq \frac{n}{16}\cdot c_i^{-1}\left(\frac{3}{n}\kappa\right)$. If $c_i(x) = x^d$, then: $\text{APX} \geq n \cdot \frac{1}{16}\cdot\left(\frac{3}{n}\kappa\right)^{1/d}$. On the other hand by

Lemma 4.4, we have that for MHR distributions: OPT $\leq n\left(\frac{e\kappa}{n}\right)^{1/d}$. Combining the lower bound on APX and the upper bound on OPT yields the theorem.

REFERENCES

[1] Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. 2013. The simple economics of approximately optimal auctions. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 628–637.

[2] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. 2013. Steering User Behavior with Badges. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13)*. ACM, New York, NY, USA, 95–106. DOI: http://dx.doi.org/10.1145/2488388.2488398

[3] Nikolay Archak and Arun Sundararajan. 2009. Optimal design of crowdsourcing contests. *ICIS 2009 proceedings* (2009), 200.

[4] Richard E. Barlow, Albert W. Marshall, and Frank Proschan. 1963. Properties of Probability Distributions with Monotone Hazard Rate. *Ann. Math. Statist.* 34, 2 (06 1963), 375–389. DOI: http://dx.doi.org/10.1214/aoms/1177704147

[5] Kim C. Border. 1991. Implementation of Reduced Form Auctions: A Geometric Approach. *Econometrica* 59, 4 (1991), pp. 1175–1187. http://www.jstor.org/stable/2938181

[6] Jeremy Bulow and Paul Klemperer. 1996. Auctions Versus Negotiations. *The American Economic Review* 86, 1 (1996), 180–194.

[7] Y. Cai, C. Daskalakis, and S. M. Weinberg. 2013. Understanding Incentives: Mechanism Design Becomes Algorithm Design. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. 618–627. DOI: http://dx.doi.org/10.1109/FOCS.2013.72

[8] Shuchi Chawla and Jason D. Hartline. 2013. Auctions with Unique Equilibria. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. ACM, New York, NY, USA, 181–196. DOI: http://dx.doi.org/10.1145/2482540.2483188

[9] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. 2012. Optimal Crowdsourcing Contests. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '12)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 856–868. http://dl.acm.org/citation.cfm?id=2095116.2095185

[10] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. 2015. Revenue maximization with a single sample. *Games and Economic Behavior* 91 (2015), 318–333.

[11] Dominic DiPalantino and Milan Vojnovic. 2009. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 119–128.

[12] David Easley and Arpita Ghosh. 2013. Incentives, Gamification, and Game Theory: An Economic Approach to Badge Design. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. ACM, New York, NY, USA, 359–376. DOI: http://dx.doi.org/10.1145/2482540.2482571

[13] Hu Fu, Jason Hartline, and Darrell Hoy. 2013. Prior-independent Auctions for Risk-averse Agents. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. ACM, New York, NY, USA, 471–488. DOI: http://dx.doi.org/10.1145/2482540.2482551

[14] Arpita Ghosh and Patrick Hummel. 2011. A Game-theoretic Analysis of Rank-order Mechanisms for User-generated Content. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC '11)*. ACM, New York, NY, USA, 189–198. DOI: http://dx.doi.org/10.1145/1993574.1993603

[15] Arpita Ghosh and Patrick Hummel. 2012. Implementing Optimal Outcomes in Social Computing: A Game-theoretic Approach. In *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*. ACM, New York, NY, USA, 539–548. DOI: http://dx.doi.org/10.1145/2187836.2187910

[16] Arpita Ghosh and Patrick Hummel. 2015. Cardinal Contests. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 377–387. DOI: http://dx.doi.org/10.1145/2736277.2741652

[17] Arpita Ghosh and Preston McAfee. 2011. Incentivizing High-quality User-generated Content. In *Proceedings of the 20th International Conference on World Wide Web (WWW '11)*. ACM, New York, NY, USA, 137–146. DOI: http://dx.doi.org/10.1145/1963405.1963428

[18] Jason D Hartline. 2015. Mechanism design and approximation. *Book draft. October* (2015).

[19] Jason D. Hartline and Tim Roughgarden. 2009. Simple Versus Optimal Mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC '09)*. ACM, New York, NY, USA, 225–234. DOI: http://dx.doi.org/10.1145/1566374.1566407

[20] Nicole Immorlica, Greg Stoddard, and Vasilis Syrgkanis. 2015. Social Status and Badge Design. In *Proceedings of the 24th International Conference on World Wide Web (WWW '15)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 473–483. DOI: http://dx.doi.org/10.1145/2736277.2741664

[21] Eric Maskin and John Riley. 1984. Optimal Auctions with Risk Averse Buyers. *Econometrica* 52, 6 (1984), pp. 1473–1518. http://www.jstor.org/stable/1913516

[22] Benny Moldovanu and Aner Sela. 2001. The Optimal Allocation of Prizes in Contests. *The American Economic Review* 91, 3 (2001), 542–558. http://www.jstor.org/stable/2677878

[23] Roger B Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.

[24] Mallesh M. Pai and Rakesh Vohra. 2014. Optimal auctions with financially constrained buyers. *Journal of Economic Theory* 150 (2014), 383 – 425. DOI: http://dx.doi.org/10.1016/j.jet.2013.09.015

[25] Tim Roughgarden, Inbal Talgam-cohen, and Qiqi Yan. 2016. Robust Auctions for Revenue via Enhanced Competition. (2016).

[26] Ester Samuel-Cahn. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability* (1984), 1213–1216.