

Optimal Crowdsourcing Contests

Shuchi Chawla*

Jason D. Hartline[†]

Balasubramanian Sivan[‡]

Abstract

We study the design and approximation of optimal crowdsourcing contests. Crowdsourcing contests can be modeled as all-pay auctions because entrants must exert effort up-front to enter. Unlike all-pay auctions where a usual design objective would be to maximize revenue, in crowdsourcing contests, the principal only benefits from the submission with the highest quality. We give a theory for optimal crowdsourcing contests that mirrors the theory of optimal auction design: the optimal crowdsourcing contest is a virtual valuation optimizer (the virtual valuation function depends on the distribution of contestant skills and the number of contestants). We also compare crowdsourcing contests with more conventional means of procurement. In this comparison, crowdsourcing contests are relatively disadvantaged because the effort of losing contestants is wasted. Nonetheless, we show that crowdsourcing contests are 2-approximations to conventional methods for a large family of “regular” distributions, and 3.2-approximations, otherwise.

1 Introduction

Crowdsourcing contests have become increasingly important and prevalent with the ubiquity of the Internet. For instance, instead of hiring a research team to develop a better collaborative filtering algorithm, Netflix issued the “Netflix challenge” offering a million dollars to the team that develops an algorithm that beats the Netflix algorithm by 10%. More generally, Taskcn allows users to post tasks with monetary rewards, collects submissions by other users, and rewards the best submission; and many Q&A sites allow users to post questions and reward the best answer with much-coveted “points”. We address two questions in this paper, (a) what format of crowdsourcing competition induces the highest-quality winning contribu-

tion, and (b) how inefficient is crowdsourcing over more conventional means of contracting.

Crowdsourcing competitions can be modeled as *all-pay* auctions. In the highest-bid-wins single-item all-pay auction, the auctioneer solicits payments (as bids), awards the item to the agent with the highest payment, and keeps all the agent payments. These auctions are well understood in settings where each agent has an independent private value for obtaining the item. In the connection to crowdsourcing contests, the “item” is the monetary reward, the payments are the submissions, and the private value is the rate at which the contestant works. However, unlike all-pay auctions, in crowdsourcing competitions the principal usually only values the winning submission and has no value for lesser submissions. Therefore, while the performance metric for auctions is usually *revenue* which is the sum of the agent payments, in crowdsourcing contests where payments are submissions, the relevant performance metric is the quality of the best submission, i.e., the maximum agent “payment”.

The *revenue equivalence* principle implies that in equilibrium the revenue of the highest-bid-wins all-pay auction is the same as that of first- and second-price auction formats; however, in these latter auction formats only the winner makes a payment. Since non-winners make payments in all-pay auctions, the maximum agent payment in all-pay auctions is lower than that of first- and second-price auctions. To connect this auction theory back to the setting of procurement, first- and second-price auctions are analogous to conventional procurement mechanisms, e.g., for government contracts; where as, the all-pay format is analogous to crowdsourcing contests. Importantly, the performance of first- and second-price procurement auctions is their revenue, this revenue is equal to the revenue of the all-pay auction, but the principal in crowdsourcing cannot attain this full revenue and therefore suffers a loss relative to conventional methods. Our first result is to show that the maximum agent payment in all-pay auctions is at least half its total revenue in expectation. Consequently the loss in crowdsourcing over conventional procurement is at most a factor of two.

Of course, highest-bid-win auctions are not necessarily revenue optimal. However, for a large class of dis-

*Computer Sciences Dept., University of Wisconsin - Madison.
shuchi@cs.wisc.edu.

[†]EECS, Northwestern University.
hartline@eecs.northwestern.edu.

[‡]Computer Sciences Dept., University of Wisconsin - Madison.
balu2901@cs.wisc.edu.

tributions (termed *regular*), auctions that award the item to the highest bidder that meets a reservation price are optimal [Mye81]. In these settings crowdsourcing contests that require submissions to be of a minimum quality (e.g., the Netflix challenge required submissions to beat the Netflix algorithm by 10%) are 2-approximations. For more general distributional settings reserve pricing gives a $\frac{e}{e-1} = 1.582$ -approximation to the optimal auction revenue [CHMS10] and crowdsourcing contests with minimum quality conditions are, therefore, $\frac{2e}{e-1} = 3.164$ -approximations. Thus, this approximation result implies a “simple versus optimal” style result, i.e., that the gains from precisely optimizing based on the distribution are at most a factor of 3.164.

Our second result derives the optimal static crowdsourcing format that maximizes the quality of the best submission. Specifically, suppose we fix the reward for the k -th best submission to be a_k (where we normalize the a_k ’s to $\sum_{k=1}^n a_k = 1$). What should the a_k ’s be? For instance on computer programming crowdsourcing site TopCoder.com, the best submission receives 2/3rds and the second-best submission receives 1/3rd of the total reward. Is this a better format than awarding the entire amount to the best submission in terms of the quality of that submission? We prove that it is not: $a_1 = 1$ and $a_k = 0$ for $k > 1$, or winner-takes-all, is the optimal choice over all such static contests.

Of course in some settings it may be better to adjust the number of rewards and their distribution across the participants dynamically as functions of the observed submission qualities. Our third result derives the format of crowdsourcing contests that dynamically optimizes the quality of the best submission. We give a complete characterization of optimal crowdsourcing contests. In what would be a familiar result to auction theorists, the optimal crowdsourcing contests are “ironed virtual value optimizers” in that the reward is divided evenly among all contestants whose submissions are tied under a weakly monotone transformation (via the ironed virtual value function) of the submission quality. Importantly, the number of contestants who share the reward is determined dynamically and each contestant’s share is the same. Perhaps surprisingly, and unlike the case of classical auction theory, the transformation to ironed virtual values depends on the number of contestants.

Optimal crowdsourcing contests require the auctioneer to know the distribution of agents’ skills, e.g. in order to pick an appropriate minimum submission quality. In our fourth result we consider the loss from not knowing the distribution. We show that under a regularity assumption on the distributions, the auctioneer is better off recruit-

ing an extra agent and running a simple highest-bid-wins contest than running the optimal contest without the extra agent; This limits the benefit of knowing the skill distribution. This result is reminiscent of and based on a similar theorem by Bulow and Klemperer [BK96] for the revenue objective.

Related Work. This paper follows the connection made between crowdsourcing contests and all-pay auctions from DiPalantino and Vojnovic [DV09] and questions from Archak and Sundararajan [AS09] and Moldovanu and Sela [MS01, MS06] on optimizing the reward structure to improve the quality of the best submissions. [AS09] and [MS01] compare winner-take-all crowdsourcing contests against ones with a statically determined division of the reward among top agents, e.g., as in the TopCoder.com mechanism. The objective in [MS01] is the sum of the qualities of submissions (analogous to revenue in our discussion) and the Archak-Sundararajan objective is the cumulative effort from the top k agents less the monetary reward. Both papers show that when agents’ submission qualities are linear in their effort winner-take-all is optimal over other static divisions. The Moldovanu-Sela result also holds when quality is a convex function of effort, but not generally for concave functions. Minor [Min11] studies a generalization of the problem of Moldovanu and Sela, and derives the optimal crowdsourcing contest via a Myersonian ironing approach (again, to maximize the sum of qualities).

Moldovanu and Sela [MS06] study both the highest quality submission objective and the revenue (sum of all submission qualities) objective. They compare the performance of two-stage contests against one-stage contests. Among one-stage contests they consider both a single grand contest, as well as many sub-contests in parallel, with the winner of each sub-contest receiving a prize. In two-stage contests, the winners of the first stage sub-contests compete in a final round. These are all static contests in the sense that the division of reward among winners of different sub-contests is predetermined. For the sum-of-qualities objective, [MS06] prove that a single grand contest is best among these contest formats. For the highest quality objective, if there are sufficiently many competitors then it is optimal to split the competitors in two divisions and to have a final among the two divisional winners. Further, as the number of competitors tends to infinity, [MS06] show that the optimal highest quality objective is at least half of the optimal sum-of-qualities objective — this is a result we generalize in this paper by proving that this factor two ratio holds for any number of competitors.

In this paper our goal is to optimize the quality of the best submission (unlike [MS01, Min11] which consider total quality) with a total reward normalized to one (unlike [AS09] which optimizes the quality of the best submissions less the monetary reward). We study optimal crowdsourcing contests over all single-stage all-pay formats, unlike [MS06] which limits the format of one-stage contests but also studies two-stage contests. Our main results are to show that the wasted effort is not large and to characterize optimal crowdsourcing contests that can potentially divide the award between agents dynamically depending on the qualities of submissions. We also show that our model is consistent with that of [MS01, AS09] in that the optimal static allocation of the award is winner-take-all¹.

The following other results relating to crowdsourcing contests are technically unrelated to ours. DiPalantino and Vojnovic [DV09] study crowdsourcing websites as a matching market. They discuss equilibria where contestants first choose which contest to participate in and then their level of effort. Yang et al. [YAA08] and DiPalantino and Vojnovic [DV09] empirically study bidder behavior from crowdsourcing website Taskcn and conclude that experienced contestants strategize better than others and their strategies match the BNE predictions fairly well.

There have been a number of studies of all-pay auctions in complete information settings (e.g., Baye et al. [BKdV96]), but these works are also technically unrelated to ours.

2 Preliminaries

Auction Theory. Consider the standard auction-theoretic problem of selling a single item to n agents. Each agent i has a private value v_i for receiving the object and is risk-neutral with linear utility $u_i = v_i x_i - p_i$ for receiving the item with probability x_i and making payment p_i . An auction \mathcal{A} solicits bids and determines the outcome which consists of an allocation $\mathbf{x} = (x_1, \dots, x_n)$ and payments $\mathbf{p} = (p_1, \dots, p_n)$.

Suppose that the agents values are drawn i.i.d. from continuous distribution F (that is, having no point-masses) with distribution function $F(z) = \Pr[v_i \leq z]$ and density function $f(z)$. A Bayes-Nash equilibrium (BNE) in auction \mathcal{A} is a profile of strategies for mapping values to bids in the auction that are a mutual best response, i.e., when the values are drawn from F and other agents follow their equilibrium strategies then each agent (weakly) prefers to also follow the prescribed strategy over taking

any other action. On valuation profile $\mathbf{v} = (v_1, \dots, v_n)$, denote the composition of an auction and a strategy profile by allocation rule $\mathbf{x}(\mathbf{v})$ and payment rule $\mathbf{p}(\mathbf{v})$.

When agent i is bidding in the auction she knows her own value v_i and assumes the other agent values are drawn from the distribution F . Denote her *interim* allocation and payment rules as $x_i(v_i) = \mathbf{E}_{\mathbf{v}}[x_i(\mathbf{v}) \mid v_i]$ and $p_i(v_i) = \mathbf{E}_{\mathbf{v}}[p_i(\mathbf{v}) \mid v_i]$, respectively. Bayes-Nash equilibrium requires that $v_i x_i(v_i) - p_i(v_i) \geq v_i x_i(z) - p_i(z)$ for all z and from this constraint is derived the standard characterization of BNE.

Theorem 1 [Mye81] *Allocation and payment rules $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$ are in BNE if and only if for all i*

1. $x_i(v_i)$ is monotone non-decreasing in v_i and
2. $p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z) dz + p_i(0)$

where usually $p_i(0) = 0$.

A simple consequence of this characterization is the *revenue equivalence* principle which states that two mechanisms with the same equilibrium allocation have the same equilibrium revenue—in fact each agent’s expected interim payment is the same.

There are three standard formats for highest-bid-wins single-item auctions: first-price, second-price, and all-pay. In the first-price variant the highest bidder wins and pays her bid, in the second-price variant (a.k.a. Vickrey auction) the highest bidder wins and pays the second highest bid, and in the all-pay variant the highest bidder wins and all bidders pay their bids. These auction formats all have BNE in which the agent with the highest valuation wins²; Therefore, revenue equivalence implies that they have the same expected revenue (sum of payments) in equilibrium.

The highest-bid-win auction formats do not always yield the highest expected revenue. To solve for optimal auctions, Myerson [Mye81] defined *virtual valuations for revenue* as $\phi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$ and proved that the expected payment of an agent, $\mathbf{E}_{v_i}[p_i(v_i)]$ is equal to her expected virtual value $\mathbf{E}_{v_i}[\phi(v_i)x_i(v_i)]$. The distribution F is said to be *regular* if the virtual valuation function is monotone. For regular distributions, maximizing virtual values point-wise is a monotone allocation rule, and therefore can be implemented in BNE. This auction serves the agent with the highest positive virtual value. By symmetry, this agent is identically the agent with the highest value that meets a reserve price of $\phi^{-1}(0)$.

¹Moldovanu and Sela in [MS06] also state that this should be true but do not provide a reference or a proof.

²Furthermore, these symmetric BNE are unique in first-price auctions, as shown in [Leb06].

Theorem 2 [Mye81] *When the virtual valuation function $\phi(\cdot)$ is monotone, the optimal auction format is highest-bid-wins with a reservation value of $\phi^{-1}(0)$, a.k.a., the monopoly price.*

It will be useful to be able to solve for the equilibrium strategies in all-pay auctions with reserves. Revenue equivalence makes this easy: the expected payment of an agent with value v_i is the same in both all-pay and the second-price auction formats. Of course in the all-pay format the agent always pays her bid; therefore, her bid $b(v_i)$ must be equal to her expected payment in the second-price auction. Let $v_{(j)}$ denote the j th largest value. Agent i 's expected payment in the second-price auction when $v_i \geq r$, is exactly $\mathbf{E}_{\mathbf{v}_{-i}}[\max(v_{(2)}, r) \mid v_i = v_{(1)}] \Pr_{\mathbf{v}_{-i}}[v_i = v_{(1)}]$, so her bid in the all-pay auction must be equal to this expectation.

Lemma 3 *In a highest-bid-wins all-pay auction with value reserve r an agent with value v_i bids*

$$b(v_i) = \mathbf{E}_{\mathbf{v}_{-i}}[\max(v_{(2)}, r) \mid v_i = v_{(1)}] \Pr_{\mathbf{v}_{-i}}[v_i = v_{(1)}] \text{ if } v_i \geq r \text{ and } 0 \text{ otherwise.}$$

The reserve specified above is in value-space. To implement such a reserve in an auction, one must translate it to a reserve in bid-space. For first- and second-price auctions this transformation is the identity function. For all-pay auctions, it can be calculated as follows. An agent with value equal to the reserve r in the second price auction pays the reserve if she wins, i.e., her expected payment is $r \Pr_{\mathbf{v}_{-i}}[r = v_{(1)}] = rF(r)^{n-1}$. By revenue equivalence the same agent in the equivalent all-pay auction must bid this expected payment; as this bid is the minimum bid that should be accepted, it is the reserve.

Lemma 4 *The highest-bid-win all-pay auction with reserve bid $rF(r)^{n-1}$ implements the highest-value-wins allocation rule with a reserve value of r .*

For irregular distributions, the optimal auction is not reserve-price based. Instead it selects the highest virtual value subject to monotonicity of the allocation rule. This optimization can be simplified by a very general *ironing* technique.

Theorem 5 [Mye81, HR08] *There is an ironing procedure that converts any virtual valuation function $\phi(\cdot)$ to a ironed virtual valuation function $\bar{\phi}(\cdot)$ that is monotone and has the property that maximizing $\phi(\cdot)$ subject to monotonicity (of the allocation rule) is equivalent to maximizing $\bar{\phi}(\cdot)$ point-wise, with ties broken randomly. The BNE with this outcome is optimal.*

Crowdsourcing. The following model for crowdsourcing contests and its connection to all-pay auctions was proposed in [DV09]. To outsource a task to the crowd a principal announces a monetary reward (normalized to 1). Each of n agents (the crowd) enters a submission. Agent i 's *skill* is denoted by v_i and with *effort*, e_i , she can produce a submission with *quality* $p_i = v_i e_i$, i.e., her skill can be thought of as a rate of work and her effort the amount of work. Each agent's skill is her private information. If x_i fraction of the reward is awarded to agent i then her utility is $u_i = x_i - e_i$. From her perspective v_i is a constant so maximizing utility is equivalent to maximizing $v_i u_i = v_i x_i - p_i$. Notice that this latter formulation of the agent's objective mirrors that from the single-item auction setting discussed previously; furthermore, as the agents exert effort up-front, crowdsourcing contests intrinsically have all-pay semantics. Because of this connection, it will be convenient to refer interchangeably to skills as values, submission qualities as payments, and rewards as allocations.

The objective for crowdsourcing contests \mathcal{A} is to maximize the quality of the best submission. Because of the connection to all-pay auctions we refer to this objective as the *maximum payment* objective and denote its value for an auction \mathcal{A} as $\mathbf{MP}[\mathcal{A}] = \mathbf{E}_{\mathbf{v}}[\max_i p_i(\mathbf{v})]$. This objective is quite different from the standard revenue maximization objective $\mathbf{Rev}[\mathcal{A}] = \mathbf{E}_{\mathbf{v}}[\sum_i p_i(\mathbf{v})]$

One aim of this paper is to quantify the loss the principal incurs from running an all-pay auction versus a more conventional means of contracting. For instance, standard formats for procurement auctions are first- or second-price. Importantly, in first- and second-price auctions \mathcal{A} all the payment comes from the highest bidder, therefore $\mathbf{MP}[\mathcal{A}] = \mathbf{Rev}[\mathcal{A}]$ and the principal is able to extract quality workmanship with no loss. In contrast, in all-pay auctions which are revenue equivalent to first- and second-price auctions the maximum payment is not equal to the total revenue and thus the efforts of non-winners constitute a loss in performance. We thus quantify the *utilization ratio* of an auction \mathcal{A} as $\frac{\mathbf{Rev}[\mathcal{A}]}{\mathbf{MP}[\mathcal{A}]}$.

We will see that the all-pay auction that optimizes maximum payment is not the same as the auction (all-pay or otherwise) that maximizes revenue. We define the *approximation ratio* of an all-pay auction to quantify its maximum payment relative to the revenue of the optimal auction, i.e., \mathcal{A} 's approximation ratio is $\frac{\mathbf{Rev}[\mathbf{OPT}]}{\mathbf{MP}[\mathcal{A}]}$. The *cost of crowdsourcing* (over conventional procurement) is then the approximation ratio of the best all-pay auction, i.e., $\inf_{\mathcal{A}} \frac{\mathbf{Rev}[\mathbf{OPT}]}{\mathbf{MP}[\mathcal{A}]}$.

Assumptions. For the rest of this paper, the space of valuations V is assumed to be an interval and the density function $f(\cdot)$ is assumed to be non-zero everywhere in V .

3 Utilization and approximation ratios

As noted previously, the maximum payment of a second- or first-price auction is equal to its total payment or revenue. On the other hand, in all-pay auctions the payment made by non-winners leads to a loss in performance. In this section we quantify this loss for a special class of all-pay auctions, namely those that always reward the highest bidder subject to an anonymous reserve price. This further allows us to find a simple all-pay auction that approximates optimal procurement.

In this section we consider highest-bidder-wins reserve-price auctions under either second-price or all-pay semantics. It is easy to see that under all-pay semantics these auctions induce symmetric continuous increasing bid functions at BNE and therefore their allocation function is identical to a second-price auction with an appropriate reserve price.

Theorem 6 *Let \mathcal{A} be any highest-bidder-wins reserve-price all-pay auction. Then $\mathbf{Rev}[\mathcal{A}] \leq 2\mathbf{MP}[\mathcal{A}]$. That is, its utilization ratio is bounded by 2.*

Proof: Let \mathbf{x} denote the allocation function of the auction and suppose that the bid function that it induces in BNE is given by $b(v)$. We can write the expected revenue of the auction as the sum of the contribution from the winning agent (i.e. the agent with the maximum payment), and the contribution from other agents. Call the first term A and the second B .

$$\mathbf{Rev}[\mathcal{A}] = \underbrace{\sum_i \int_v b(v) \Pr_{\mathbf{v}_{-i}}[v = v_{(1)}] f(v) dv}_A + \underbrace{\sum_i \int_v b(v) (1 - \Pr_{\mathbf{v}_{-i}}[v = v_{(1)}]) f(v) dv}_B$$

Note that A is precisely $\mathbf{MP}[\mathcal{A}]$. We will now show that $A \geq B$, or $A - B \geq 0$. By the revenue equivalence principle, $b(v)$ is equal to the expected payment that an agent with value v makes in a second-price auction with the same allocation rule. Let $g(v)$ denote the expected payment (the maximum of the second highest value and the reserve price) in the second-price auction given that v is the highest value. Then we get that

$b(v) = g(v) \Pr_{\mathbf{v}_{-i}}[v = v_{(1)}]$. We note that g is a weakly increasing function, but, $\Pr_{\mathbf{v}_{-i}}[v = v_{(1)}] = F(v)^{n-1}$ is a strictly increasing function since $f(v) \neq 0$ for all $v \in V$. Now we can write $A - B$ as

$$\begin{aligned} A - B &= \sum_i \int_v b(v) (2\Pr_{\mathbf{v}_{-i}}[v = v_{(1)}] - 1) f(v) dv \\ &= \sum_i \int_v g(v) F(v)^{n-1} (2F(v)^{n-1} - 1) f(v) dv \\ &= \sum_i \int_v g(F^{-1}(t)) t^{n-1} (2t^{n-1} - 1) dt \end{aligned}$$

where, in the third inequality, we substituted t for $F(v)$. Note that since $F(v)$ is a strictly increasing function, $F^{-1}(t)$ is well defined, and is strictly increasing in t .

Next we note that ignoring the g term, the integral is non-negative:

$$\int_0^1 t^{n-1} (2t^{n-1} - 1) dt = \frac{2}{2n-1} - \frac{1}{n} > 0$$

Let us consider the effect of the g term. The function $t^{n-1} (2t^{n-1} - 1)$ vanishes for two values of t namely 0 and $(1/2)^{\frac{1}{n-1}}$. Between these two values the function is negative, and for $t > (1/2)^{\frac{1}{n-1}}$, the function is positive. So when the function is multiplied by $g(F^{-1}(t))$, a non-decreasing function of t , the negative portion of the integral is magnified to a smaller extent than the positive portion, implying that the integral stays positive. This completes the proof. \blacksquare

Tightness of Theorem 6. We now exhibit an example where the utilization ratio of 2 is tight. Consider a setting with n agents, with each agent's value distributed independently according to the $U[0, 1]$ distribution. Consider the second-price auction with reserve price of $1/2$. The expected revenue of this auction can be computed to be $(n-1)/(n+1) - \Theta(1/2^n)$. The corresponding all-pay auction uses a reserve price of $1/2^n$ and induces a bid function $b(v) = 0$ for $v < 1/2$ and $b(v) = \frac{n-1}{n}v^n + \frac{1}{2n}v^{n-1} - \frac{n-1}{n2^{n-1}}v + 2^{-n}(1 - 1/n)$ otherwise. The expected revenue of the all-pay auction is the same as that of the second-price auction, roughly $(n-1)/(n+1)$, which approaches 1 as n increases. On the other hand, the expected maximum payment can be computed to be $(2n^2 - 2n + 1)/(4n^2 - 2n) + \Theta(1/2^n)$ which approaches $1/2$ as n increases. In Section 5 we will revisit this example and show that even the optimal all-pay auction (which

is slightly better) only achieves an expected maximum payment approaching $1/2$ for this setting.

Utilization ratio for other all-pay auctions. We note that the bound on utilization ratio does not hold for arbitrary symmetric all-pay auctions. For example, the all-pay auction corresponding to a revenue-optimal auction that requires ironing over large intervals of values induces a bidding function that is constant over those intervals. This results in many agents being tied for the reward, all making the same (low) payments but only one contributing to the maximum payment.

Approximation ratio. Recall that the approximation ratio of an all-pay auction \mathcal{A} is $\frac{\text{Rev}[\text{OPT}]}{\text{MP}[\mathcal{A}]}$, where OPT is the revenue optimal auction. We now use the bound on utilization ratio to prove that all-pay auctions achieve good approximation ratios. In particular, we use the fact that for regular distributions highest-bidder-wins reserve-price auctions are revenue-optimal (Theorem 2) whereas for irregular distributions they are within a factor of $\frac{e}{e-1} = 1.582$ of optimal [CHMS10]. The following corollaries then follow from Theorem 6 upon applying the revenue equivalence principle.

Corollary 7 *When agents' value distributions are regular, there exists an α such that the highest-bid-wins all-pay auction with reserve bid α achieves an approximation ratio at most 2.*

Corollary 8 *For all i.i.d. value distributions, there exists an α such that the highest-bid-wins all-pay auction with reserve bid α achieves an approximation ratio at most 3.164.*

These corollaries imply that the cost of crowdsourcing is always small — no more than 3.164. The above example with uniform distributions shows that the approximation factor in Corollary 7 is tight. An extension of the same example in Section 5 shows that the worst-case cost of crowdsourcing can be no smaller than 2.

4 Optimal static crowdsourcing contests

Consider the class of static contests that predetermine the division of the reward into a_1, \dots, a_n , with $\sum_i a_i = 1$. Agents are ordered by their submission qualities and awarded the corresponding fraction of reward, i.e., the i th best submission gets an a_i fraction of the reward. Note

that the Topcoder.com example mentioned in the introduction, where the best submission receives $2/3$ rd and the second-best submission receives $1/3$ rd of the total reward, falls under this class of contests.

In this section we characterize the static auctions that are optimal with respect to the maximum payment objective. Note that the class of static auctions is symmetric, i.e., a permutation of bids results in the same permutation of the allocation and payments. Since agents' private values are identically distributed, any such symmetric allocation rule induces a symmetric equilibrium in which all agents use an identical bidding function. This in turn implies that the allocation as a function of agents' values is also symmetric across agents. The following theorem shows that in fact it is optimal over all static contests to allocate the entire reward to the best submission. (See the appendix for a proof.)

Theorem 9 *When the bidders' valuations are drawn i.i.d., the optimal static all-pay auction is a highest-bid-wins auction.*

5 Optimal crowdsourcing contests

In this section, we characterize symmetric all-pay auctions that are optimal with respect to the maximum payment objective. As discussed in Section 4, a symmetric auction induces a symmetric equilibrium which makes the allocation as a function of agents' values symmetric across agents.

We first present a characterization of the expected maximum payment of any symmetric all-pay auction in terms of an appropriately defined virtual value function. This characterization immediately implies that the optimal mechanism is a virtual value maximizer.

Definition 1 *For a given distribution F with density function f and an integer n , we define the virtual value for maximum payment, $\psi_n(z)$ as*

$$\psi_n(z) = zF(z)^{n-1} - \frac{1 - F(z)^n}{n.f(z)}$$

Lemma 10 *Consider a setting with n agents and values distributed i.i.d. according to distribution F . Let \mathcal{A} be a symmetric all-pay auction implementing the allocation function \mathbf{x} . Then $\text{MP}[\mathcal{A}] = \mathbf{E}[\sum_i x_i(\mathbf{v})\psi_n(v_i)]$.*

Proof: Suppose that the allocation function \mathbf{x} induces a symmetric bid function $b(\cdot)$ on the agents. Recall that by the revenue equivalence principle, $b(v)$ is equal to the expected payment that an agent with value v makes under

$\mathbf{x}(\cdot)$. From Theorem 1 we get the following expression for $b(v)$ where x_i is the expected allocation to agent i in expectation over \mathbf{v}_{-i} .

$$b(v_i) = v_i x_i(v_i) - \int_{z=0}^{v_i} x_i(z) dz$$

Because the equilibrium is symmetric, one of the agents with the highest bid is the agent with the highest value³, i.e., with $v_i = v_{(1)}$. We attribute the maximum payment received by the mechanism to this agent. We can now use the above formulation of the bid function to calculate the expected contribution of agent i to the maximum payment objective.

$$\begin{aligned} \mathbf{MP}_i[\mathcal{A}] &= \int_{v_i} b(v_i) \mathbf{Pr}_{\mathbf{v}_{-i}}[v_i = v_{(1)}] f(v_i) dv_i \\ &= \int_{v_i} \left[v_i x_i(v_i) - \int_{z=0}^{v_i} x_i(z) dz \right] F(v_i)^{n-1} f_i(v_i) dv_i \end{aligned}$$

In order to simplify the second term in the integral we interchange the order of integration over z and v_i , integrate over v_i , and then rename z as v_i . We get:

$$\begin{aligned} \mathbf{MP}_i[\mathcal{A}] &= \int_{v_i} v_i x_i(v_i) F(v_i)^{n-1} f_i(v_i) dv_i \\ &\quad - \int_{v_i} x_i(v_i) \left(\frac{1 - F(v_i)^n}{n} \right) dv_i \\ &= \int_{v_i} x_i(v_i) \left\{ v_i F(v_i)^{n-1} - \frac{1 - F(v_i)^n}{n f(v_i)} \right\} f_i(v_i) dv_i \\ &= \int_{v_i} x_i(v_i) \psi_n(v_i) f_i(v_i) dv_i = \mathbf{E}_{\mathbf{v}}[x_i(\mathbf{v}) \psi_n(v_i)] \end{aligned}$$

Summing over i implies the lemma. \blacksquare

Optimal allocation rules and regularity. The characterization of Lemma 10 immediately implies that in order to maximize the expected maximum payment, we should maximize the virtual surplus of the mechanism for maximum payment. In other words, we should allocate the entire reward to the agent that has the maximum virtual value $\psi_n(v_i)$ (subject to this value being non-negative). However, this results in a monotone allocation function only if the virtual value function is monotone non-decreasing. To this end, we define regularity for maximum payment as follows.

³Note that the bid function need only be weakly increasing, so there may be ties for the highest bid.

Definition 2 A distribution F is said to be n -regular with respect to maximum payment if $\psi_n(\cdot)$ is a monotone non-decreasing function. The distribution is said to be regular w.r.t. maximum payment if $\psi_n(\cdot)$ is monotone non-decreasing for all positive integers n .

For distributions that are regular w.r.t. maximum payment, allocating to the agent with the highest non-negative virtual value is monotone and therefore can be implemented in BNE. Since agents have i.i.d. values, this outcome corresponds to allocating to the agent with the highest value, who is in turn the agent with the highest bid. Therefore, the optimal mechanism is a highest-bid-wins reserve-price mechanism. The reserve value for the mechanism is given by $\psi_n^{-1}(0)$ and the reserve bid can be computed by applying Lemma 4 to this value. We note that generally the reserve price is a function of n and decreases with n , even for distributions that are regular for all n .

Theorem 11 Let F be a distribution that is n -regular w.r.t. maximum payment. Then the optimal all-pay auction for n agents with values distributed independently according to F is a highest-bid-wins auction with a reserve price.

Two examples. We now revisit the example with n agents and values distributed according to $U[0, 1]$ that was discussed in Section 3. The following expression defines the virtual value for maximum payment in this case:

$$\psi_n(z) = z^n(1 + 1/n) - 1/n \text{ for } z \in [0, 1]$$

This is an increasing function for all n . Therefore, the $U[0, 1]$ distribution is regular. The optimal reserve value is given by $\psi_n^{-1}(0) = (n + 1)^{-1/n}$, and the optimal reserve bid is $1/(n + 1)$. Therefore, the optimal all-pay auction serves the highest bidder subject to her bid being at least $1/(n + 1)$. The expected maximum payment of this auction can be calculated to be $\frac{n}{2(n+1)}$.

Next consider a setting with two agents and values distributed i.i.d. according to the exponential distribution. That is, $F(v) = 1 - e^{-v}$ for $v \geq 0$. We can calculate the virtual value function as $\psi_2(z) = (z - 1) + e^{-z}(1/2 - z)$. This function is negative below 1.21 and positive thereafter. Furthermore, it is non-decreasing above 0.24, particularly throughout the range where it is non-negative. So although the exponential distribution is not regular w.r.t. maximum payment, the optimal all-pay auction still turns out to be a highest-bid-wins auction with a reserve price of 1.21 and a corresponding reserve bid of 0.85.

An interesting point to note about the above example is that distributions that are regular with respect to the

usual notion of virtual value for revenue, are not necessarily regular with respect to maximum payment even for $n = 2$. However, for a large subset of such distributions, namely those that satisfy the monotone hazard rate condition (Definition 3 below), the optimal all-pay auction continues to have the simple form given in Theorem 11.

Regularity and MHR. A frequently used assumption in mechanism design literature is that value distributions satisfy the *monotone hazard rate* (MHR) condition defined below. Many common distributions such as the uniform, Gaussian, and exponential distributions satisfy this property. Distributions that satisfy MHR are regular and therefore do not require ironing in the context of revenue maximization. As our example above shows, MHR distributions are not necessarily regular with respect to maximum payment.

Definition 3 *The hazard rate of a distribution F with density function f is defined as $h(x) = \frac{f(x)}{1-F(x)}$. A distribution is said to have a monotone hazard rate (MHR) if the hazard rate function is monotone non-decreasing.*

Lemma 12 *Let F be a distribution satisfying the MHR condition. Then for any n and any interval of values over which ψ_n is non-negative, ψ_n is monotone non-decreasing.*

Proof: We can rewrite the virtual value function in terms of the hazard rate $h(z)$ of the distribution as follows.

$$\begin{aligned} \psi_n(z) &= zF(z)^{n-1} - \frac{1}{nh(z)} \sum_{j=0}^{n-1} F(z)^j \\ &= F(z)^{n-1} \left\{ z - \frac{1}{nh(z)} \sum_{j=0}^{n-1} F(z)^{-j} \right\} \end{aligned}$$

The function $h(z)$ is a non-negative non-decreasing function. Therefore, $(-1/nh(z))$ is a negative non-decreasing function. On the other hand, $\sum_{j=0}^{n-1} F(z)^{-j}$ is a decreasing function of z . The product of a negative non-decreasing function and a decreasing function is a non-decreasing function. Therefore, the term within brackets is a non-decreasing function of z . The term outside brackets, $F(z)^{n-1}$, is also an always positive increasing function. Therefore, the product of the two terms is an increasing function over any interval where it is positive. ■

We obtain the following corollary.

Corollary 13 *Let F be a distribution that satisfies MHR. Then the optimal all-pay auction for values distributed independently according to F is a highest-bid-wins auction with a reserve price.*

Irregular distributions and ironing. For distributions that are not regular according to the definition above, we can apply an ironing procedure from Theorem 5 to ψ_n to obtain an ironed virtual value function $\tilde{\psi}_n$. This function is monotone non-decreasing and by Theorem 5 the BNE that optimizes it pointwise optimizes the maximum payment objective.

The optimal mechanism in this case allocates the entire reward to the agent with the maximum ironed virtual value, in the case of ties distributing the reward equally among the tied agents⁴. Since the ironed virtual value function is a weakly increasing function, the induced bid function is constant in the intervals where the ironed virtual value is constant, and discontinuous at the ends of those intervals. In effect, this creates intervals of bids that are suboptimal to make at any value; call these intervals “forbidden”. In order to implement the mechanism as an all-pay auction, we identify the forbidden intervals; then we round every bid in a forbidden interval down to the closest “allowed” bid, and distribute the reward equally among the highest bidders (subject to an appropriate reserve price defined by $\tilde{\psi}_n^{-1}(0)$). We therefore get the following theorem:

Theorem 14 *For any setting with i.i.d. values, the optimal all-pay auction is defined by a reserve price and a subset of bids called forbidden bids, that has the following format: the auction solicits bids and rounds them down to the nearest non-forbidden bids; it then distributes the reward equally among the highest bidders subject to the bids being above the reserve price.*

An example of ironing. We now present a simple example of a distribution that is irregular w.r.t. maximum payment, and derive its ironed virtual value and as well as forbidden bids. There are two agents, each with a value drawn independently from $U[1, 2]$ with probability 3/4 and from $U(2, 3]$ with probability 1/4. Figure 1 below shows the virtual value function ψ_2 and its integral with respect to $q = F(v)$ using thick grey lines; their ironed counterparts are shown in thin red lines. The integral of the virtual value function as a function of q is given by the expression $\frac{1}{2}F^{-1}(q)(1 - q^2)$. We iron this function by

⁴An equivalent way of resolving ties in the maximum ironed virtual value is to allocate the reward to a random tied agent.

taking its convex envelope; $\tilde{\psi}_2$ is then the derivative with respect to q of that convex envelope.

The ironed virtual value is constant in the interval $[1.918, 2.167]$. The probability of allocation (not plotted), and therefore the bid function, are also constant over this interval. The corresponding bid function is plotted with a thin black line below; there are two forbidden bid intervals, namely $[1.10, 1.199]$ and $(1.199, 1.31]$, with the intermediate value of 1.199 being allowed. The two forbidden intervals correspond to the two discontinuities in the probability of allocation at the end points of the ironed interval.




Figure 1: The Ironing Procedure

Irregularity as a function of n . An interesting point to note is that irregularity increases with n . Specifically, the intervals of values that require ironing under ψ_n increase with n .⁵ This does not necessarily imply that as n increases a larger and larger number of agents are tied for the reward, for two reasons: (1) reserve value (not the reserve bid) increases with n , and (2), due to the form of the virtual value function, ironing is typically necessary at low values rather than at high values.

6 Contests without reserve prices

As we show above, optimal crowdsourcing contests depend on knowing the agents' value distribution. To what

⁵This happens because the intervals requiring ironing are precisely those where the integral of the virtual value function is non-concave; Increasing n amounts to multiplying the integral with a convex function resulting in non-concave intervals continuing to stay non-concave.

extent is it important to know the distribution? In particular, under what conditions does the simple highest-bidder-wins contest without any reserve bid approximate the optimal one? We now show a result in the style of Bulow and Klemperer's work [BK96] on the standard goal of maximizing the expected revenue. Bulow and Klemperer showed that for i.i.d. value distributions that are regular w.r.t. revenue, it is better to run a Vickrey auction with no reserve price on $n + 1$ agents than to run an optimal auction on only n agents. That is, the ability to recruit an extra agent in the auction is more profitable to the auctioneer than knowing the distribution.

We first note that Bulow and Klemperer's result implies that for distributions that are regular w.r.t. revenue, the highest-value-wins auction with no reserve price on n agents is within a factor of $(1 - 1/n)$ of the optimal mechanism in terms of revenue. This combined with Theorem 6 gives us the following theorem.

Theorem 15 *For i.i.d. distributions that are regular w.r.t. revenue, the highest-bid-wins all-pay auction without a reserve bid obtains an approximation ratio of $2n/(n - 1)$.*

We remark that for the highest-value-wins auction without reserve prices, the revenue converges to the optimal as more and more agents are added. However for all-pay auctions adding more and more agents does not improve the approximation ratio beyond 2. Next in the style of Bulow and Klemperer, we compare the benefit of knowing the distribution to that of adding an extra bidder for the maximum payment objective.

Theorem 16 *For i.i.d. distributions that are regular w.r.t. maximum payment, the expected maximum payment of the highest-bidder-wins all-pay auction with no reserve bid on $n + 1$ agents is at least as large as the expected maximum payment of the optimal all-pay auction on n agents.*

Proof: Consider all all-pay auction on $n + 1$ agents that must always allocate the entire reward to the agents. Lemma 10 implies that the optimal such auction allocates the entire reward to the agent with the highest virtual value, who by symmetry and regularity is the agent with the highest value.

On the other hand, another way to always allocate the entire reward is to run the optimal all-pay auction on the first n agents; if the reward goes unallocated in that auction, then allocate it to the $n + 1$ st agent. Clearly, in this mechanism, in any BNE the $n + 1$ st agent makes a 0 payment regardless of his value because his value does not effect the outcome. Therefore the expected maximum payment of the auction is precisely the expected maximum payment of the optimal all-pay auction on n agents. ■

In Section 5 we showed that distributions that satisfy the MHR condition are essentially regular w.r.t. maximum payment in the sense that virtual values are non-decreasing wherever they are nonnegative (and so, no ironing is required for the optimal all-pay auction). We remark that for the above theorem this property is not sufficient and we need complete regularity, i.e. virtual values should be non-decreasing over the entire range, in order for the result to hold.

References

- [AS09] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *International Conference on Information Systems (ICIS)*, 2009.
- [BK96] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86:180–194, 1996.
- [BKdV96] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8:291–305, 1996. 10.1007/BF01211819.
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 311–320, New York, NY, USA, 2010. ACM.
- [DV09] Dominic DiPalantino and Milan Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, EC '09, pages 119–128, New York, NY, USA, 2009. ACM.
- [HR08] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 75–84, New York, NY, USA, 2008. ACM.
- [Leb06] Bernard Lebrun. Uniqueness of the equilibrium in first-price auctions. *Games and Economic Behavior*, 55(1):131–151, 2006.
- [Min11] Dylan Minor. Increasing efforts through incentives in contests. Manuscript, <http://faculty.haas.berkeley.edu/dminor/Increasing%20Effort-Minor,%20Dylan.pdf>, Jan. 2011.
- [MS01] Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001.
- [MS06] Benny Moldovanu and Aner Sela. Contest architecture. *Journal of Economic Theory*, 126(1):70–97, 2006.
- [Mye81] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [YAA08] Jiang Yang, Lada A. Adamic, and Mark S. Ackerman. Crowdsourcing and knowledge sharing: strategic user behavior on taskcn. In *Proceedings of the 9th ACM conference on Electronic commerce*, EC '08, pages 246–255, New York, NY, USA, 2008. ACM.

A Appendix

Proof of Theorem 9. Let the agent values be distributed independently according to distribution function F , with density function f . Consider the static allocation rule $\mathcal{A} = (a_1, \dots, a_k, 0, \dots, 0)$, i.e, the agent with the i -th highest bid gets a_i fraction of the reward if $i \leq k$, an 0 otherwise. We have $\sum_{i=1}^k a_i = 1$. We focus on the symmetric bid-function $b(\cdot)$ induced by this allocation rule.

In a truthful auction with allocation rule \mathcal{A} , the expected payment made by the r -th highest bidder is $p_r(z) = \sum_{j=r+1}^{k+1} v_{jr}(z)(a_{j-1} - a_j)$, where $v_{jr}(z)$ is the expectation of the j -th highest bid (=value) given the r -th highest bid is z .

Let $g(j, n, z)$ denote the expectation of the j -th highest draw among n draws from F , given that the maximum draw is at most z . Then we have $v_{jr}(z) = g(j - r, n - r, z)$.

The contribution of bidder i to the maximum payment objective is

$$\begin{aligned} \mathbf{MP}_i[\mathcal{A}] &= \int_{v_i} b(v_i) \Pr_{\mathbf{v}_{-i}}[v_i = v_{(1)}] f(v_i) dv_i \\ &= \int_{v_i} b(v_i) F(v_i)^{n-1} f_i(v_i) dv_i \end{aligned}$$

Since agents values are drawn i.i.d. from F , we have $\mathbf{MP}[\mathcal{A}] = n\mathbf{MP}_i[\mathcal{A}]$.

Because the bid functions are symmetric, by the revenue equivalence principle, $b(z)$ equals the expected payment made by an agent with value z in a truthful auction with the same allocation rule. So,

$$\begin{aligned} b(z) &= \sum_{r=1}^k \Pr_{\mathbf{v}_{-i}}[z = v_{(r)}] \cdot p_r(z) \\ &= \sum_{r=1}^k \binom{n-1}{r-1} (1 - F(z))^{r-1} F(z)^{n-r} \cdot \left\{ \sum_{j=1}^{k+1-r} g(j, n-r, z)(a_{j+r-1} - a_{j+r}) \right\} \end{aligned}$$

We prove the theorem by showing that $\frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k}$ is negative. When we change a_k we assume that all the mass is transferred to (or drawn from) a_1 . This will prove that the optimal allocation rule is to put all the mass on a_1 , i.e., $a_1 = 1$.

Using the formula for $b(z)$, it is easy to observe that for $r = 2$ to $r = k - 1$, terms corresponding to that specific r in $\frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k}$ will be an integral with an integrand of

$$\binom{n-1}{r-1} (1 - F(z))^{r-1} F(z)^{2n-r-1} \cdot \{-g(k-r, n-r, z) + g(k-r+1, n-r, z)\}$$

This integrand is negative because g is a decreasing function in its first argument.

The term corresponding to $r = 1$ in $\frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k}$ will be an integral with an integrand of

$$F(z)^{2n-2} \cdot \{-g(1, n-1, z) - g(k-1, n-1, z) + g(k, n-1, z)\}$$

Note that the above integrand is negative even if $g(1, n-1, z)$ term were not there.

The term corresponding to $r = k$ in $\frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k}$ will be an integral with a positive integrand of

$$\binom{n-1}{k-1} (1 - F(z))^{k-1} F(z)^{2n-k-1} \cdot \{g(1, n-k, z)\}$$

Our proof is going to upper bound $\frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k}$ by ignoring certain negative terms in it, and show that even the upper bound is negative. In particular, we only consider terms corresponding to $r = k - 1$, $r = k$ and one term of $r = 1$,

namely $F(z)^{2n-2} \cdot \{-g(1, n-1, z)\}$. Let this upper bound be denoted by Q .

$$\begin{aligned} \frac{d\mathbf{MP}_i[\mathcal{A}]}{da_k} &\leq Q = - \int_z F(z)^{2n-2} g(1, n-1, z) dF(z) \\ &\quad - \binom{n-1}{k-2} \int_z (1-F(z))^{k-2} F(z)^{2n-k} g(1, n-k+1, z) dF(z) \\ &\quad + \binom{n-1}{k-2} \int_z (1-F(z))^{k-2} F(z)^{2n-k} g(2, n-k+1, z) dF(z) \\ &\quad + \binom{n-1}{k-1} \int_z (1-F(z))^{k-1} F(z)^{2n-k-1} g(1, n-k, z) dF(z) \end{aligned}$$

We derive the expressions for $g(1, n, z)$ and $g(2, n, z)$ below.

$$\begin{aligned} g(1, n, z) &= n \int_0^z y \frac{f(y)}{F(z)} \left(\frac{F(y)}{F(z)} \right)^{n-1} dy \\ &= z - \frac{\int_0^z F(t)^n dt}{F(z)^n} \end{aligned}$$

$$\begin{aligned} g(2, n, z) &= n(n-1) \int_0^z y \frac{f(y)}{F(z)} \left(1 - \frac{F(y)}{F(z)} \right) \left(\frac{F(y)}{F(z)} \right)^{n-2} dy \\ &= z - \left[n \frac{\int_0^z F(t)^{n-1} dt}{F(z)^{n-1}} - (n-1) \frac{\int_0^z F(t)^n dt}{F(z)^n} \right] \end{aligned}$$

We substitute the expression for g into Q .

$$\begin{aligned} Q &= - \int_z F(z)^{n-2} \left[zF(z)^n - \int_z F(t)^n dt \right] dF(z) \\ &\quad + \binom{n-1}{k-1} \int_z (1-F(z))^{k-1} F(z)^{2n-k-1} z dF(z) \\ &\quad + \binom{n-1}{k-2} \int_z (1-F(z))^{k-2} F(z)^{n-1} \left(\int_0^z F(t)^{n-k+1} dt \right) dF(z) \\ &\quad + \binom{n-1}{k-2} (n-k) \int_z (1-F(z))^{k-2} F(z)^{n-1} \left(\int_0^z F(t)^{n-k+1} dt \right) dF(z) \\ &\quad - \binom{n-1}{k-2} (n-k+1) \int_z (1-F(z))^{k-2} F(z)^n \left(\int_0^z F(t)^{n-k} dt \right) dF(z) \\ &\quad - \binom{n-1}{k-1} \int_z (1-F(z))^{k-1} F(z)^{n-1} \left(\int_0^z F(t)^{n-k} dt \right) dF(z) \end{aligned}$$

We now factor the term $(1 - F(z))^{k-1}$ as $(1 - F(z))^{k-2} \cdot (1 - F(z))$ and then group terms. We get

$$\begin{aligned}
Q &= - \int_z F(z)^{n-2} \left[zF(z)^n - \int_z F(t)^n dt \right] dF(z) \\
&\quad - \binom{n-1}{k-1} \int_z (1 - F(z))^{k-2} F(z)^{2n-k} z dF(z) \\
&\quad + \binom{n-1}{k-1} \int_z (1 - F(z))^{k-2} F(z)^{2n-k-1} z dF(z) \\
&\quad + \binom{n-1}{k-2} (n-k+1) \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left(\int_0^z F(t)^{n-k+1} dt \right) dF(z) \\
&\quad - \binom{n-1}{k-2} \left[(n-k+1) - \frac{n-k+1}{k-1} \right] \int_z (1 - F(z))^{k-2} F(z)^n \left(\int_0^z F(t)^{n-k} dt \right) dF(z) \\
&\quad - \binom{n-1}{k-1} \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left(\int_0^z F(t)^{n-k} dt \right) dF(z)
\end{aligned}$$

We have to prove that $Q \leq 0$. This is equivalent to proving that

$$\begin{aligned}
&\int_z F(z)^{n-2} \left[zF(z)^n - \int_0^z F(t)^n dt \right] dF(z) \\
&\quad + \binom{n-1}{k-1} \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left[zF(z)^{n-k+1} - \int_0^z F(t)^{n-k+1} dt \right] dF(z) \\
&\quad - \binom{n-1}{k-1} \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left[zF(z)^{n-k} - \int_0^z F(t)^{n-k} dt \right] dF(z) \\
&\quad \geq \\
&\quad \binom{n-1}{k-1} (k-2) \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left(\int_0^z F(t)^{n-k+1} dt \right) dF(z) \\
&\quad - \binom{n-1}{k-1} (k-2) \int_z (1 - F(z))^{k-2} F(z)^n \left(\int_0^z F(t)^{n-k} dt \right) dF(z)
\end{aligned}$$

The RHS can be seen to be negative. Thus it is enough to prove that the LHS is positive. Rewriting the terms in the square bracket via integration by parts,

$$\begin{aligned}
&n \int_z F(z)^{n-2} \left(\int_0^z tF(t)^{n-1} dF(t) \right) dF(z) \\
&\quad + \binom{n-1}{k-1} (n-k+1) \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left(\int_0^z tF(t)^{n-k} dF(t) \right) dF(z) \\
&\quad - \binom{n-1}{k-1} (n-k) \int_z (1 - F(z))^{k-2} F(z)^{n-1} \left(\int_0^z tF(t)^{n-k-1} dF(t) \right) dF(z)
\end{aligned}$$

Changing the order of integration, we have the LHS as,

$$\int_{t=0}^{\infty} tF(t)^{n-k-1} f(t) \left\{ \begin{aligned} &\binom{n-1}{k-1} (n-k+1) \left(\int_{F(t)}^1 (1 - F(z))^{k-2} F(z)^{n-1} dF(z) \right) \left[F(t) - \frac{n-k}{n-k+1} \right] \\ &+ n \left(\int_{F(t)}^1 F(z)^{n-2} dF(z) \right) F(t)^k \end{aligned} \right\} dt$$

Applying integration by parts again, (this time taking t as one term and the rest as the differential part) we get the LHS

as,

$$\int_{t=0}^{\infty} \left(\int_{F(t)}^1 u^{n-k-1} \left\{ \binom{n-1}{k-1} (n-k+1) \left(\int_u^1 (1-F(z))^{k-2} F(z)^{n-1} dF(z) \right) \left[u - \frac{n-k}{n-k+1} \right] + n \left(\int_u^1 F(z)^{n-2} dF(z) \right) u^k \right\} du \right) dt$$

Rewrite the above integral as $\int_{t=0}^{\infty} H_n(F(t)) dt$ where

$$H_n(x) = \int_x^1 u^{n-k-1} \left\{ \binom{n-1}{k-1} (n-k+1) \left(\int_u^1 (1-v)^{k-2} v^{n-1} dv \right) \left[u - \frac{n-k}{n-k+1} \right] + n \left(\int_u^1 v^{n-2} dv \right) u^k \right\} du$$

If we prove that $H_n(x)$ is always non-negative for $x \in [0, 1]$ we are done. We have

$$-H'_n(x) = x^{n-k-1} \left\{ \binom{n-1}{k-1} (n-k+1) \left(\int_x^1 (1-v)^{k-2} v^{n-1} dv \right) \left[x - \frac{n-k}{n-k+1} \right] + n \left(\int_x^1 v^{n-2} dv \right) x^k \right\}$$

Observe that $-H'_n(x)$ is negative for small values of x and positive for large values of x and never becomes negative after it has become positive. Thus, $H_n(x)$ is first increasing and then decreasing. We know that $H_n(1) = 0$. If we prove that $H_n(0) \geq 0$, we would have proven that $H_n(x)$ is always non-negative.

$$\begin{aligned} H_n(0) &= \binom{n-1}{k-1} (n-k+1) \int_0^1 u^{n-k-1} \left(\int_u^1 (1-v)^{k-2} v^{n-1} dv \right) \left[u - \frac{n-k}{n-k+1} \right] du + n \int_0^1 \left(\int_u^1 v^{n-2} dv \right) u^{n-1} du \\ &= \binom{n-1}{k-1} (n-k+1) \int_0^1 (1-v)^{k-2} v^{n-1} \left(\int_0^v u^{n-k-1} \left[u - \frac{n-k}{n-k+1} \right] du \right) dv + n \int_0^1 v^{n-2} \left(\int_0^v u^{n-1} du \right) dv \\ &= \binom{n-1}{k-1} \int_0^1 (1-v)^{k-2} v^{2n-k-1} (v-1) dv + \frac{1}{2n-1} \\ &= -2 \binom{n-1}{k-1} \int_0^{\pi/2} \cos^{4n-2k-1}(\theta) \sin^{2k-1}(\theta) d\theta + \frac{1}{2n-1} \end{aligned}$$

The integral $\int_0^{\pi/2} \cos^m(\theta) \sin^n(\theta) d\theta = \frac{\Gamma(\frac{m+1}{2})\Gamma(\frac{n+1}{2})}{2\Gamma(\frac{m+n+2}{2})}$ Accordingly, we have

$$\begin{aligned} H_n(0) &= -\binom{n-1}{k-1} \frac{\Gamma(2n-k)\Gamma(k)}{\Gamma(2n)} + \frac{1}{2n-1} \\ &> 0 \end{aligned}$$

■