ARTICLE IN PRESS YGAME:2460

Games and Economic Behavior ••• (••••) •••–•••

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Optimal crowdsourcing contests

Shuchi Chawla[a,1], Jason D. Hartline[b,2], Balasubramanian Sivan[a,*,3]

[a] Computer Sciences Dept., University of Wisconsin-Madison, United States
[b] Electrical Engineering and Computer Science, Northwestern University, United States

ARTICLE INFO

Article history:
Received 18 May 2012
Available online xxxx

JEL classification:
D44
D47
D82

Keywords:
Crowdsourcing contest
All-pay auction
Bayes–Nash equilibrium
Approximation

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. We show that the total wasted effort is at most the maximum effort which implies that crowdsourcing contests are a 2-approximation to an idealized model of conventional procurement.

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 Neflix algorithm by 10%. More generally, Taskcn.com 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) how inefficient is crowdsourcing over more conventional means of contracting, and (b) what format of crowdsourcing contest induces the highest-quality winning contribution.

Crowdsourcing contests 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 a crowdsourcing contest, the principal announces a reward in return for completing a certain task. Each contestant produces a submission, the quality of which depends both on the skill of the contestant and the level of effort that she puts into it.[4] Effort is a sunk cost for the contestant—it is a "payment" she makes regardless of whether or not she wins the

* Corresponding author.
E-mail addresses: shuchi@cs.wisc.edu (S. Chawla), hartline@eecs.northwestern.edu (J.D. Hartline), balu2901@cs.wisc.edu (B. Sivan).

[1] Supported in part by NSF award CCF-0830494 and in part by a Sloan Foundation fellowship.
[2] Supported in part by NSF award CCF-0830773.
[3] Supported in part by NSF award CCF-0830494.
[4] Importantly, we assume that the effort that a participant puts in to her submission translates linearly into the quality of the submission.

http://dx.doi.org/10.1016/j.geb.2015.09.001
0899-8256/

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

reward; it is therefore analogous to a bid in an all-pay auction. The skill of the contestant, or the rate at which she works, is analogous to her private value. The reward is analogous to the item in an all-pay auction. Whereas in optimal auction design the performance metric is the sum of agent payments, in crowdsourcing contests where bids are submissions, the principal usually only values the winning submission and has no value for lesser submissions. The relevant performance metric for crowdsourcing is therefore the quality of the best submission, i.e., the maximum agent "payment".

Optimal crowdsourcing contests We formalize the problem of designing an optimal crowdsourcing contest as the problem of maximizing the maximum agent payment in a Bayes–Nash Equilibrium (BNE), over the class of all single-round single-item all-pay auctions with a fixed number of participants. If an auction has more than one BNE, we consider the optimal BNE, namely, the one with the largest maximum agent payment. We consider symmetric environments, i.e., where agents' skills or values are identically distributed. We focus on the class of anonymous (a.k.a. symmetric) contests; in other words, the contest does not favor one contestant over another, and permuting the submissions permutes the outcome of the contest. Anonymous contests are not necessarily optimal in asymmetric settings (see, e.g., Che and Gale, 2003). Surprisingly, we show that even in symmetric environments the optimal contest is not necessarily symmetric[5]: handicapping one of the contestants can help improve the outcome of the contest. Nevertheless, restricting attention to symmetric contests is reasonable from a practical viewpoint: in a real world setting it may be socially unacceptable to favor one agent over another in the absence of any distinguishing features. Finally, we assume that the principal fixes the maximum total reward upfront. In real world contests, the principal typically decides the maximum reward amount she is willing to spend on the contest before running the contest, for reasons including liquidity constraints. For instance, Netflix announced a reward size of one million dollars before running its contest. Given that the maximum reward size is fixed, without loss of generality we normalize[6] it to $1.

The designer must determine how to distribute the reward among the agents after submissions have been made. There are several ways of doing so. The principal may, for example, give the entire reward to the entrant with the best submission, or distribute it in a fixed proportion among the top few submissions, or decide the distribution based on the qualities of the submissions, and may or may not set a "reserve" on the submission quality. Which of these formats maximizes the quality of the best submission?

In order to answer this question we develop a theory for the maximum payment objective that is similar to optimal auction theory for the revenue objective (Myerson, 1981). We give a complete characterization of the expected highest payment in symmetric BNE for all single round contests, and use this to derive the crowdsourcing contest that exhibits the optimal symmetric BNE among all single round contests; we refer to this contest henceforth as the optimal crowdsourcing contest. In what would be a familiar result to auction theorists, optimal crowdsourcing contests are "ironed virtual value optimizers" in that the reward is divided evenly among all contestants whose submissions are above a minimum quality threshold, and are tied under a weakly monotone transformation (via an ironed virtual value function) of the submission quality. Importantly, the number of contestants who share the reward is determined based on the quality of their submissions, and each contestant who gets any share at all gets an equal share. Unlike the case of classical auction theory, the transformation to ironed virtual values depends on the number of contestants.

When value distributions have monotone increasing hazard rates, the optimal contest is a winner-takes-all contest with a reserve, that is, no ironing is required. The reserve is an increasing function of the number of contestants. This makes sense intuitively: the reserve bid in an all-pay auction functions as an entry fee, deterring low-value contestants from participating in the competition, and encouraging high-value contestants to bid higher, thereby improving the maximum payment.

The quality minus reward objective Another natural objective in a crowdsourcing contest is to maximize the quality of the best submission minus the monetary reward given out (still with the maximum possible reward normalized to one). Intuitively, in this setting, the principal should be more aggressive about withholding the reward when submissions are of poor quality. Our theory and results extend as-is to this objective, with the optimal contests having somewhat higher reserves than for the maximum payment objective.

Simple contests Optimal crowdsourcing contests, as suggested by our theory, may require ironing and may therefore be complex and impractical. We therefore consider the problem of optimizing a contest over the class of all rank-based-reward contests. Specifically, suppose we fix the reward for the k-th best submission to be a_k (where we normalize the a_k's to $\sum_{k=1}^{n} a_k = 1$) with the restriction that $a_k \geq a_{k+1}$ for all $k < n$. 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 highest quality 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 rank-based-reward contests. Archak and Sundararajan (2009) obtain the same result, albeit only as the number of participants n goes to infinity. Our result holds for every n.

[5] Note that the maximum payment objective is a non-concave function of the allocation rule, and the resulting optimization problem is non-convex.

[6] Because the quality of a submission scales linearly with the effort that a contestant puts into it, the effort that the contestant decides to expend at equilibrium scales linearly with the reward she expects to obtain. In other words, (up to a scaling factor) the outcome of the contest depends not on the total amount of reward offered, but on how it is distributed across contestants.

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

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 next result we consider the loss from not knowing the distribution. For the sum of payments objective, Bulow and Klemperer (1996) proved that for regular distributions recruiting an extra bidder is more profitable to the auctioneer than knowing the distribution. We show that this result implies that a simple highest-bid-wins contest approximates the optimal contest within a factor approaching 2; this limits the benefit of knowing the skill distribution.

Comparison to conventional procurement For a principal looking to contract out work, the following model of conventional procurement is an alternative to crowdsourcing. The principal issues a request for quotes for undertaking a project for a given reward; agents submit quotes in the form of a quality guarantee; the principal chooses the highest quality quote and pays the winning agent the prespecified reward; and the winning agent completes the project at the quoted quality. This method of procurement corresponds to a first-price auction (see Section 2 for a definition) via the same correspondence by which crowdsourcing corresponds to an all-pay auction. As there may be additional efficiency losses associated with implementing this method of procurement in practice we refer to it as *idealized procurement*.

Idealized procurement leads to a better quality outcome in comparison to crowdsourcing. To compare these two kinds of contests it is equivalent to compare the first-price and the highest-bid-wins all-pay auctions. For the objective of sum of payments, the *revenue equivalence* principle of Myerson (1981) implies that the two corresponding auction formats are equivalent: in the unique BNE with i.i.d. agents,[7] the sum of payments of the highest-bid-wins all-pay auction is the same as that of the first-price auction. When the objective is the maximum payment, however, the latter format is better: in the first-price format only the winner makes a payment, so the maximum payment is equal to the auction's sum of payments; on the other hand, non-winners also make payments in all-pay auctions, so therein the maximum payment is smaller than the sum of payments. We conclude that there is a loss in the quality of the winning submission in crowdsourcing relative to idealized procurement. This loss is exactly the sum of payments of the losers.

On the other hand, conventional procurement may be impractical in settings where it is difficult for the principal to judge a proposal without seeing the end product. For example, for tasks requiring some amount of innovation or creativity (such as logo design or the Netflix challenge), conventional procurement would overlook candidates that have a short résumé but an unusual skillset and that may produce excellent submissions.

In weighing the relative benefits and drawbacks of crowdsourcing and conventional procurement, it is important to be able to quantify the loss in the quality of the outcome that a principal suffers by choosing the former instead of the latter. We use approximation to compare crowdsourcing contests to idealized procurement. We ask whether there are simple, practical crowdsourcing contests that approximate the performance of ideal procurement. We quantify this approximation in terms of the ratio between the maximum payment achieved by the optimal first-price auction and that achieved by a candidate all-pay auction. Our goal is to show that this ratio is small.

We show that the ratio of the expected sum of payments and the expected highest payment in any highest-bid-wins all-pay auction is at most two.[8] This means that for a fixed cost the principal can obtain at least half as good a submission through the crowdsourcing format as through idealized procurement. Since the quality of a submission is linearly related to the amount of effort that an agent spends on it, this in turn implies that procuring the same quality good costs at most twice as much with crowdsourcing as with idealized procurement. This comparison is, of course, a bit unfair as our model of idealized procurement ignored the costs of submitting quotes, evaluating quotes, and enforcing that the procured work meets the quoted quality. Therefore, while crowdsourcing costs twice as much as idealized procurement, once these practical constraints are considered, its relative performance is only better. The fact that it is a good approximation to ideal procurement and is robust to many of the implementation issues in conventional procurement gives some justification for its gaining popularity.

For a large class of skill (value) distributions (termed *regular)*, the contests that achieve the factor-of-two approximation to idealized procurement are extremely simple: the principal announces a certain reserve on quality, and of the submissions that cross this quality threshold, the best one is rewarded. The Netflix challenge, which required submissions to beat the Netflix algorithm by 10% is an example of such a contest. When the distribution over agents' skills is irregular, we similarly show that crowdsourcing contests with a quality reserve are generally at worst a four-approximation to the optimal idealized procurement.

Related work There is an extensive literature on contest architecture where the objective is to maximize the quality of the best submission. This literature has applications in political contests, R & D races, trade wars, and many other settings where agents expend resources before a winner is determined. Whereas our work characterizes the optimal contest among the class of all single-round contests, most other works optimize over a smaller restricted class of competitions. Taylor (1995) focuses on winner-takes-all contests and considers the problem of restricting entry into the contest by setting an appropriate entry price. Fullerton and McAfee (1999) show that in a large class of settings, the optimal number of contestants is two,

[7] See Chawla and Hartline (2013) for proof of uniqueness for both first-price and highest-bid-wins all-pay auctions.

[8] Recently Gavious and Minchuk (2014) extended this result to show that in any highest-bid-wins all-pay auction, for any $k < n$, where n is the number of agents, the expected sum of all but the top k payments is at most the expected kth highest payment. Similar results hold also for highest-m-bids-win all-pay auctions.

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

and consider using auctions to select the contestants. Che and Gale (2003) consider a setting where the designer can specify the number of participants that are allowed to compete and each contestant bids for a prize from a given menu. Megidish and Sela (2013) consider contests where the entry fee is exogenously given. They show that when the entry fee is too high, winner-takes-all is not necessarily the optimal contest format, and can be improved upon by distributing the prize equally among all participants. Moldovanu and Sela (2006) study multi-round contests, and compare the performance of two-round contests with that of single-round ones. They show that 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, they show that the optimal highest quality objective is at least half of the optimal sum-of-qualities objective—we generalize this result to a broader set of contests, namely winner-takes-all with reserve, and show that the factor of two ratio holds for any number of competitors.

Other objectives that have been studied in the context of contest design include maximizing the sum of the qualities of submissions with a normalized reward (Moldovanu and Sela, 2001, 2006; Minor, 2011), and maximizing the cumulative effort from the top k agents less the monetary reward (Archak and Sundararajan, 2009). In many of these models (as we also show for the maximum payment objective), the optimal *rank-based-reward* allocation of the award turns out to be winner-take-all.

Following up on our work, Gavious and Minchuk (2014) recently studied the "sum-of-top-k-payments" objective for highest-bid-wins all-pay auctions. Extending our Theorem 4.1, they showed that in this auction for any k, the expected sum of all but the top k payments is at most the expected kth highest payment; They obtained similar results for highest-m-bids-win all-pay auctions as well.

The following other results relating to crowdsourcing contests are technically unrelated to ours. DiPalantino and Vojnovic (2009) 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. (2008) and DiPalantino and Vojnovic (2009) empirically study bidder behavior from crowdsourcing website Taskcn and conclude that experienced contestants strategize better than others and their strategizes match the BNE predictions fairly well. This paper follows the connection made between crowdsourcing contests and all-pay auctions in DiPalantino and Vojnovic (2009).

There have been a number of studies of all-pay auctions in complete information settings (e.g., Baye et al., 1996; Barut and Kovenock, 1998); however, these works are also technically unrelated to ours.

2. Preliminaries

Model for a single item auction Here we formalize the standard auction-theoretic problem of selling a single item to n agents. Each agent i has a non-negative private value v_i for receiving the object and is risk-neutral with linear utility: the utility for receiving the item with probability x_i and making payment p_i is given by $u_i = v_i x_i - p_i$. Agents' values are assumed to be drawn i.i.d. from continuous distribution F over $[\underline{v}, \bar{v}]$ where $\underline{v} \geq 0$, with distribution function $F(z) = \mathbf{Pr}[v_i \leq z]$ and strictly positive density function $f(z)$. An auction $\mathcal{A}$ solicits bids and determines the outcome which consists of allocation and payments. Formally, an auction $\mathcal{A}$ is a pair of functions $(\mathbf{x}^{\mathcal{A}}, \mathbf{p}^{\mathcal{A}})$, where

1. $\mathbf{x}^{\mathcal{A}} : {\mathbb{R}^+}^n \to [0,1]^n$ is the allocation function determining the probability of allocation $x_i^{\mathcal{A}}(\mathbf{b})$ for agent i at bid profile $\mathbf{b}$ such that $\sum_{i\in[n]} x_i^{\mathcal{A}}(\mathbf{b}) \leq 1$;
2. $\mathbf{p}^{\mathcal{A}} : {\mathbb{R}^+}^n \to \mathbb{R}^n$ is the payment function determining the payment $p_i^{\mathcal{A}}(\mathbf{b})$ for agent i at bid profile $\mathbf{b}$.

Basic auction theory We now present some known results that will be useful in our analysis of crowdsourcing. Readers familiar with the basics of auction theory may safely skip the first part of this subsection up to Theorem 2.3.

A Bayes–Nash equilibrium (BNE) in a single-item auction $\mathcal{A}$ is a profile of agent strategies $\mathbf{s} = (s_1, \ldots, s_n)$, each $s_i : \mathbb{R}^+ \to \mathbb{R}^+$ mapping a value to a bid that is a best response to the other strategies and the common knowledge that values are drawn i.i.d. from F. Let $\mathbf{s}(\mathbf{v}) = (s_1(v_1), \ldots, s_n(v_n))$. Let $x_i(v_i) = \mathbf{E}_{\mathbf{v}}\left[x_i^{\mathcal{A}}(\mathbf{s}(\mathbf{v})) \mid v_i\right]$ and $p_i(v_i) = \mathbf{E}_{\mathbf{v}}\left[p_i^{\mathcal{A}}(\mathbf{s}(\mathbf{v})) \mid v_i\right]$ denote the *interim* allocation and payment rules, respectively, for an agent i. The following is a standard characterization of BNE outcomes:

Theorem 2.1. *(See Myerson, 1981.) Interim allocation and payment rules $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$ of an auction $\mathcal{A}$ are in BNE only if for all i*[9]*:*

1. *$x_i(v_i)$ is monotone non-decreasing in v_i and*
2. *$p_i(v_i) = v_i x_i(v_i) - \int_{\underline{v}}^{v_i} x_i(z)dz + h_i$, where h_i is a constant independent of v_i.*

In this paper, we focus on the BNE of auctions where $x_i(\underline{v}) = 0$ and $h_i = 0$ for all i.

We consider auctions in three standard formats: first-price, second-price and all-pay.

[9] Furthermore, the "if" direction holds under the additional assumption that strategies map "onto" actions.

Definition 1. A single-item auction $\mathcal{A}$:

1. has all-pay semantics if every agent pays his bid: $p_i^{\mathcal{A}}(\mathbf{b}) = b_i$ for all i.
2. has first-price semantics if every *winner* pays his bid: $p_i^{\mathcal{A}}(\mathbf{b}) = b_i \cdot x_i^{\mathcal{A}}(\mathbf{b})$ for all i.
3. has second-price semantics if the agents' payments follow the payment identity in Theorem 2.1: $x_i^{\mathcal{A}}(\mathbf{b})$ is non-decreasing in b_i, and[10] $p_i^{\mathcal{A}}(\mathbf{b}) = b_i x_i^{\mathcal{A}}(\mathbf{b}) - \int_0^{b_i} x_i^{\mathcal{A}}(z, \mathbf{b}_{-i})\, dz + h_i(\mathbf{b}_{-i})$ where $h_i(\mathbf{b}_{-i})$ does not depend on b_i. This auction has a BNE in which every agent bids his true value.

We will be particularly interested in the highest-bid-wins version of first-price, second-price and all-pay auctions, with a reserve $r^{\mathcal{A}}$. In all three auctions, when there is a tie for the highest bid, it is broken uniformly at random. Let $H = \text{argmax}_{i \in [n]} b_i$ be the set of agents with the highest bids. In all three auction formats, if $i \in H$, then $x_i^{\mathcal{A}}(\mathbf{b}) = \frac{1}{|H|}$ if $b_i \geq r^{\mathcal{A}}$ and 0 otherwise, and, $x_i^{\mathcal{A}}(\mathbf{b}) = 0$ for all $i \notin H$. Only the payment functions differ, as defined in Definition 1.

A simple consequence of the characterization in Theorem 2.1 is the *revenue equivalence* principle which states that two mechanisms with the same interim allocation rules in a BNE have the same interim expected payments for every agent in that BNE. In particular, this implies that with i.i.d. agents, the three standard highest-bid-wins single-item auctions (with or without reserve)—first-price, second-price, and all-pay—all have the same interim expected payments, in the BNE that awards the item to the agent with the highest value subject to reserve price.[11]

Highest-bid-wins auction formats do not always exhibit the BNE with the highest revenue, i.e., highest sum of payments $\sum_i \mathbf{E}_{v_i}[p_i(v_i)]$. To solve for the auction that exhibits the BNE with the highest revenue, Myerson (1981) defined *virtual valuations for sum of payments* as $\phi(v_i) = v_i - \frac{1-F(v_i)}{f(v_i)}$ and proved that the expected payment $\mathbf{E}_{v_i}[p_i(v_i)]$ of an agent in any BNE is equal to her expected virtual surplus $\mathbf{E}_{v_i}[\phi(v_i)x_i(v_i)]$ in the same BNE. The distribution F is said to be *regular* if the virtual valuation function $\phi(\cdot)$ is monotone. For regular distributions, an interim allocation rule that allocates the item to the agent with the highest non-negative virtual value (and allocates to no agent if all agents' virtual values are negative) is a monotone allocation rule and therefore can be implemented in BNE. Let $H = \text{argmax}_{i \in [n]} \phi(b_i)$ be the set of agents with the highest virtual bids. Since $\phi(\cdot)$ is monotone H is also the set of agents with the highest bids. The auction that implements the above interim allocation rule has $x_i^{\mathcal{A}}(\mathbf{b}) = \frac{1}{|H|}$ if $i \in H$ and $b_i \geq \phi^{-1}(0)$, and, $x_i^{\mathcal{A}}(\mathbf{b}) = 0$ otherwise.

Theorem 2.2. *(See Myerson, 1981.) When the virtual valuation function $\phi(\cdot)$ is monotone, the sum-of-payments-optimal auction format is highest-bid-wins with a reservation value of $\phi^{-1}(0)$, a.k.a., the* monopoly *price.*

For irregular distributions, Chawla et al. (2010) show that highest-bid-wins auctions with reserve are close to optimal.

Theorem 2.3. *(See Chawla et al., 2010.) The highest-bid-wins second-price auction with an appropriate reserve obtains at least half the total sum of payments as the sum-of-payments-optimal auction.*

We now focus on symmetric all-pay auctions since we model crowdsourcing contests as one.

Definition 2. An auction $\mathcal{A}$ is symmetric if under a permutation of bids, its allocation and payments are permuted. I.e., given a permutation σ, we have $\mathbf{x}^{\mathcal{A}}(\sigma(\mathbf{b})) = \sigma(\mathbf{x}^{\mathcal{A}}(\mathbf{b}))$ and $\mathbf{p}^{\mathcal{A}}(\sigma(\mathbf{b})) = \sigma(\mathbf{p}^{\mathcal{A}}(\mathbf{b}))$.

Henceforth, in our discussion of symmetric all-pay auctions, we will generally focus on the direct implementation of the auction and specify the allocation rule in value space, i.e., the agents reveal their true values to the auction, and the auction charges each agent i the interim expected price $p_i(v_i)$ that depends only on the agent's own value v_i, and not on the outcome of the auction $\mathbf{x}(\mathbf{v})$. In the indirect implementation of an all-pay auction (the one used in practice), where every agent pays his bid, the bid of an agent is equal to his expected payment in the direct implementation. This correspondence will allow us to translate an allocation rule in a direct implementation to that in an indirect implementation.

It will be useful to be able to solve for the equilibrium strategies in highest-bid-wins 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 the all-pay and the second-price auction formats (in the BNE where no agent apart from the highest valued agent ever gets allocated). Of course in the all-pay format the agent always pays her bid; therefore, her bid $b(v_i)$ must be equal to her interim expected payment in the second-price auction. Consider, for example, the highest-bid-wins all-pay auction with a value reserve (reserve in the value space) denoted by r. Let $v_{(j)}$ denote the jth largest value. Agent i's expected payment in the second-price auction when $v_i \geq r$, is exactly $\mathbf{E}_{\mathbf{v}_{-i}}\left[\max(v_{(2)}, r) \mid v_i = v_{(1)}\right] \mathbf{Pr}_{\mathbf{v}_{-i}}\left[v_i = v_{(1)}\right]$, so her bid in the all-pay auction must be equal to this expectation.

[10] The lower bound of the integral in the expression for $p_i^{\mathcal{A}}(\mathbf{b})$ is 0 even though the smallest value is $\underline{v}$ because the smallest possible bid is still 0.

[11] With i.i.d. agents, all three auction formats are guaranteed to have a BNE whose interim allocation rule awards the item to the agent with the highest value subject to reserve price. Although there is a unique BNE for the first-price and the highest-bid-wins all-pay auctions (Chawla and Hartline, 2013), uniqueness does not hold for second-price auctions.

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

Lemma 2.4. *In a highest-bid-wins all-pay auction with value reserve r an agent with value v_i bids*

$$b(v_i) = \mathbf{E}_{\mathbf{v}_{-i}}\big[\max(v_{(2)}, r) \mid v_i = v_{(1)}\big]\,\mathbf{Pr}_{\mathbf{v}_{-i}}\big[v_i = v_{(1)}\big]$$

if $v_i \geq r$ and 0 otherwise.

The reserve specified above is in value-space. To implement such a reserve in an indirect implementation all-pay 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 $r\,\mathbf{Pr}_{\mathbf{v}_{-i}}\big[r = v_{(1)}\big] = rF(r)^{n-1}$. 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 2.5. *The highest-bid-win all-pay auction with a bid reserve of $rF(r)^{n-1}$ implements the highest-value-wins allocation rule with a value reserve of r.*

For irregular distributions, i.e., when $\phi(\cdot)$ is non-monotone, the sum-of-payments-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 2.6. *(See Myerson, 1981; Hartline and Roughgarden, 2008.) There is an* ironing procedure *that converts any virtual valuation function $\phi(\cdot)$ to an* ironed virtual valuation *function $\bar{\phi}(\cdot)$ that is monotone and has the property that the allocation rule that awards the item to the agent i with the highest non-negative $\bar{\phi}(v_i)$ (and awarding to no agent if all $\bar{\phi}(v_i)$'s are negative), with ties broken uniformly at random, is sum-of-payments optimal.*

The final ingredient from auction theory that is necessary for our analysis is the uniqueness of equilibrium. For most of the symmetric all-pay auctions discussed in this paper, the symmetric Bayes–Nash equilibria where each agent follows the same bid function are in fact the only equilibria. This uniqueness of equilibrium theorem follows from a general result of Chawla and Hartline (2013) stated below. The only auctions we consider that are not captured by this result are those that require ironing; In this case, it is unknown whether asymmetric equilibria exist in addition to the canonical symmetric one.

Theorem 2.7. *(See Chawla and Hartline, 2013.) In the all-pay auction parameterized by a reserve price and a non-increasing sequence of rewards $a_1, \ldots, a_n$ where the agents whose bids meet the reserve are assigned to the rewards in decreasing order of bid, a symmetric Bayes–Nash equilibrium exists and is the unique equilibrium.*[12] *In this equilibrium, bids are in the same order as the values, that is, the highest bidder is the highest value agent, and so on.*

Model for crowdsourcing The following model for crowdsourcing contests and their connection to all-pay auctions was proposed in DiPalantino and Vojnovic (2009). To outsource a task to the crowd a principal announces a monetary reward, which we normalize to \$1 (i.e., the maximum size of the total reward is \$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 u_i 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 contests as auctions, skills as values, submission qualities as payments, and rewards as allocations.

The objective for a crowdsourcing contest $\mathcal{A}$ 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 $\mathbf{MP}[\mathcal{A}] = \max_{\text{All BNEs of } \mathcal{A}} \mathbf{E}_{\mathbf{v}}[\max_i p_i(\mathbf{v})]$ if $\mathcal{A}$ has at least one BNE, and 0 otherwise. The central design question in this paper is to identify the auction $\mathcal{A}$, among the class of all symmetric all-pay auctions, with the maximum $\mathbf{MP}[\mathcal{A}]$. The objective $\mathbf{MP}[\mathcal{A}]$ is quite different from the standard *sum of payments* objective $\mathbf{SP}[\mathcal{A}] = \max_{\text{All BNEs of } \mathcal{A}} \mathbf{E}_{\mathbf{v}}\big[\sum_i p_i(\mathbf{v})\big]$ if $\mathcal{A}$ has at least one BNE, and 0 otherwise. Therefore, the optimal auction for sum of payments, $\mathcal{A}^{\mathrm{SP}}$, and optimal all-pay auction for maximum-payment, $\mathcal{A}^{\mathrm{MP}}$, are generally distinct.

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 have first- or second-price semantics. Importantly, in the highest-bid-wins first- and second-price auctions $\mathcal{A}$ all the payment comes from the highest bidder, therefore $\mathbf{MP}[\mathcal{A}] = \mathbf{SP}[\mathcal{A}]$ and the principal is able to extract quality workmanship with no loss. In contrast, in the highest-bid-wins all-pay auction which is revenue equivalent to highest-bid-wins first- and second-price auctions the maximum

[12] Notice that the highest-bid-wins all-pay auction sets $a_1 = 1$ and $a_j = 0$ for all $j \neq 1$.

payment is not equal to the total sum of payments and thus the efforts of non-winners constitute a loss in performance. We thus quantify the *utilization ratio* of an auction $\mathcal{A}$ as $\frac{\mathbf{SP}[\mathcal{A}]}{\mathbf{MP}[\mathcal{A}]}$ (if $\mathcal{A}$ does not have a BNE, we define its utilization ratio to be 0).

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 sum of payments. We define the *approximation ratio* of an all-pay auction to quantify its maximum payment relative to the sum of payments of the optimal auction (among all auctions), i.e., $\mathcal{A}$'s approximation ratio is $\frac{\mathbf{SP}[\mathcal{A}^{\mathrm{SP}}]}{\mathbf{MP}[\mathcal{A}]}$. The *cost of crowdsourcing* (over conventional procurement) is then the approximation ratio of the best all-pay auction, i.e., $\frac{\mathbf{SP}[\mathcal{A}^{\mathrm{SP}}]}{\mathbf{MP}[\mathcal{A}^{\mathrm{MP}}]}$.

Utilization and approximation ratios can be calculated in worst-case over distributions from a given class. We will consider the class of all distributions as well as the restriction to distributions that satisfy Myerson's regularity condition (i.e., with monotone virtual valuation functions; defined in the following subsection).

3. Optimal crowdsourcing contests

In this section we characterize optimal crowdsourcing contests over all symmetric single-round contests.

Rank-based-reward contests Consider the class of contests that predetermine the division of the reward into $a_1, \ldots a_n$, with $\sum_i a_i = 1$, and $a_k \geq a_{k+1}$ for all $k < n$. Agents are ordered by their submission qualities and awarded the corresponding fraction of reward, i.e., the ith 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 2/3rds and the second-best submission receives 1/3rd of the total reward, falls under this class of contests. For this class, the following theorem shows that the optimal contest allocates the entire reward to the entrant with the best submission; we defer the proof to Appendix A.

Theorem 3.1. *When the bidders' valuations are i.i.d., the optimal rank-based-reward all-pay auction is a highest-bid-wins auction.*

Symmetric contests A symmetric auction is one where a permutation of contestants' bids results in the same permutation of the allocation and payments. We will now derive the optimal symmetric all-pay auction for the maximum payment objective. We begin by characterizing the expected maximum payment in symmetric Bayes–Nash equilibria. This characterization will allow us to solve for the all-pay auction that exhibits the optimal symmetric equilibrium. The resulting optimal auction will turn out have a unique BNE that is also symmetric.

We characterize the expected maximum payment of any symmetric all-pay equilibrium in terms of an appropriately defined virtual value function. This characterization immediately implies that the optimal all-pay auction is a virtual value maximizer (cf. Myerson, 1981).

Definition 3. For a given distribution F with density function f and an integer n, we define the *maximum-payment virtual value*, $\psi_n(z)$, as

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

Lemma 3.2. *Let x be the interim allocation rule for each agent induced by a symmetric Bayes–Nash equilibrium in a symmetric all-pay auction $\mathcal{A}$ over n agents with agent values drawn i.i.d. from F. Then the expected maximum payment satisfies* $\mathbf{MP}[\mathcal{A}] = n \cdot \mathbf{E}_{v\sim F}[\psi_n(v)x(v)]$.

Proof. In an all-pay auction an agent's bid is her payment therefore, from Theorem 2.1, we get the following expression for $b(v)$ in terms of $x(v)$, the probability that an agent with value v wins when the other agents' values are drawn from the distribution.

$$b(v) = vx(v) - \int_{\underline{v}}^{v} x(z)\,dz.$$

Because the equilibrium is symmetric, the agent with the highest value, i.e., with $v_i = v_{(1)}$, is one of the agents with the highest bid.[13] 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.

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

$$\mathbf{MP}_i[\mathcal{A}] = \int_{\underline{v}}^{\bar{v}} b(v_i)\mathbf{Pr}_{\mathbf{v}_{-i}}\left[v_i = v_{(1)}\right] f(v_i)\,dv_i$$
$$= \int_{\underline{v}}^{\bar{v}} \left[v_i x(v_i) - \int_{\underline{v}}^{v_i} x(z)\,dz \right] F(v_i)^{n-1} f(v_i)\,dv_i .$$

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:

$$\mathbf{MP}_i[\mathcal{A}] = \int_{\underline{v}}^{\bar{v}} v_i x(v_i) F(v_i)^{n-1} f(v_i)\,dv_i - \int_{\underline{v}}^{\bar{v}} x(v_i)\left(\frac{1-F(v_i)^n}{n}\right) dv_i$$
$$= \int_{\underline{v}}^{\bar{v}} \left(v_i F(v_i)^{n-1} - \frac{1-F(v_i)^n}{nf(v_i)} \right) \times x(v_i) f(v_i)\,dv_i$$
$$= \int_{\underline{v}}^{\bar{v}} \psi_n(v_i) x(v_i) f(v_i)\,dv_i$$
$$= \mathbf{E}_{v_i \sim F}[\psi_n(v_i) x(v_i)] .$$

Summing over i implies the lemma. □

Maximum payment less the monetary reward whenever given out Given an all-pay auction $\mathcal{A}$, the objective $\mathbf{MPLMR}[\mathcal{A}] = \max_{\text{All BNEs of } \mathcal{A}} \mathbf{E}_{\mathbf{v}}\left[\max_i p_i(\mathbf{v}) - \sum_i x(\mathbf{v})\right]$ (defined to be -1 if $\mathcal{A}$ does not have any BNE) captures the principal's expected utility. By the characterization in Lemma 3.2, we have

$$\mathbf{E}_{\mathbf{v}}\left[\max_i p_i(\mathbf{v}) - \sum_i x(\mathbf{v})\right] = \sum_i \mathbf{E}_{v_i \sim F}[(\psi_n(v_i) - 1) x(v_i)] .$$

Thus, all our discussions and conclusions in this section for the maximum payment objective carry over to this objective upon replacing the maximum-payment virtual value $\psi_n(\cdot)$, by $\psi_n(\cdot) - 1$. From now on, we just focus on the maximum payment objective.

Optimal allocation rules and regularity The characterization of Lemma 3.2 immediately implies that in order to maximize the expected maximum payment, we should maximize the maximum-payment virtual surplus of the mechanism. In other words, we should allocate the entire reward to the agent who has the largest $\psi_n(v_i)$ (subject to this value being non-negative). However, this results in a monotone allocation function only if the maximum-payment virtual value function is monotone non-decreasing. We define the notion of regularity for maximum payment to capture this monotonicity requirement.

Definition 4. A distribution F is said to be *n-maximum-payment-regular* if $\psi_n(\cdot)$ is a monotone non-decreasing function wherever it is non-negative, that is, $\psi_n(v) > 0$ implies $\psi_n'(v) \geq 0$. The distribution is said to be *maximum-payment regular* if for all positive integers n, $\psi_n(\cdot)$ is monotone non-decreasing wherever it is non-negative.

For distributions that are maximum-payment regular, allocating to the agent with the highest non-negative maximum-payment 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 calculated by applying Lemma 2.5 to this value. We note that the reserve value and the reserve bid are increasing functions of n. This is because the virtual value can be rewritten as

$$F(z)^{n-1} \cdot \left[z - \frac{1-F(z)^n}{nF(z)^{n-1} f(z)} \right] = F(z)^{n-1} \cdot \left[z - \frac{1-F(z)}{f(z)} \frac{\sum_{j=0}^{j=n-1} F(z)^{-j}}{n} \right],$$

which can be seen to be a decreasing function in n, whenever it is positive.

Theorem 3.3. *When agent values are distributed i.i.d. from a distribution that is n-maximum-payment-regular then the optimal all-pay auction is highest-bid-wins with a reserve price.*

Two examples We now give two examples to illustrate the derivation of the maximum-payment virtual values and the optimal contest.

Consider a setting with n agents, with each agent's value distributed independently according to the $U[0, 1]$ distribution. To calculate the expected maximum payment, we begin by writing the expression for the maximum-payment virtual value:

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

This is an increasing function for all n. Therefore, the $U[0, 1]$ distribution is maximum-payment 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 $\frac{n}{2(n+1)}$ which approaches $1/2$ as n increases.

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 maximum-payment virtual value function as $\psi_2(z) = (z - 1) + e^{-z}(1/2 - z)$. This function is negative below $z = 1.21$ and positive thereafter. Furthermore, it is non-decreasing above $z = 0.24$, particularly throughout the range where it is non-negative. The exponential distribution is therefore 2-maximum-payment-regular, and the optimal all-pay auction is a highest-bid-wins auction with a reserve price of 1.21 and a corresponding reserve bid of 0.85.

We now show that a large class of distributions, namely those that satisfy the monotone hazard rate condition (Definition 5 below), are maximum-payment regular.

Regularity and MHR A common assumption in mechanism design literature is that value distributions satisfy the *monotone hazard rate* (MHR) condition defined below. Many 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 sum of payments maximization. We will show that they are also maximum-payment regular.

Definition 5. 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 3.4. *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.

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

The function $h(z)$ is a non-negative non-decreasing function, and thus $\frac{1}{nh(z)}$ is a non-negative non-increasing function. On the other hand, $\sum_{j=0}^{n-1} F(z)^{-j}$ is a decreasing function of z. The product of two non-negative non-increasing functions is a non-increasing 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 3.5. *When agent values are distributed i.i.d. from a monotone hazard rate distribution then the optimal all-pay auction 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 2.6 to ψ_n to obtain an ironed virtual value function $\bar{\psi}_n$. This function is monotone non-decreasing, and along the lines of Theorem 2.6, the BNE that allocates the entire reward to the agent i with the highest non-negative $\bar{\psi}_n(v_i)$ (and rewarding no agent if all $\bar{\psi}_n(v_i)$'s are negative), breaking ties uniformly at random, or equivalently, distributing the reward equally among the tied agents, optimizes the maximum payment objective.

Having determined the optimal allocation rule as a function of agents' values (direct implementation), we will now determine the allocation rule in bid space (indirect implementation). In order to do so, we first derive the BNE strategies corresponding to the optimal allocation rule, and then translate the allocation into a function of the bids. Once again we use revenue equivalence to derive the BNE bidding strategies. 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

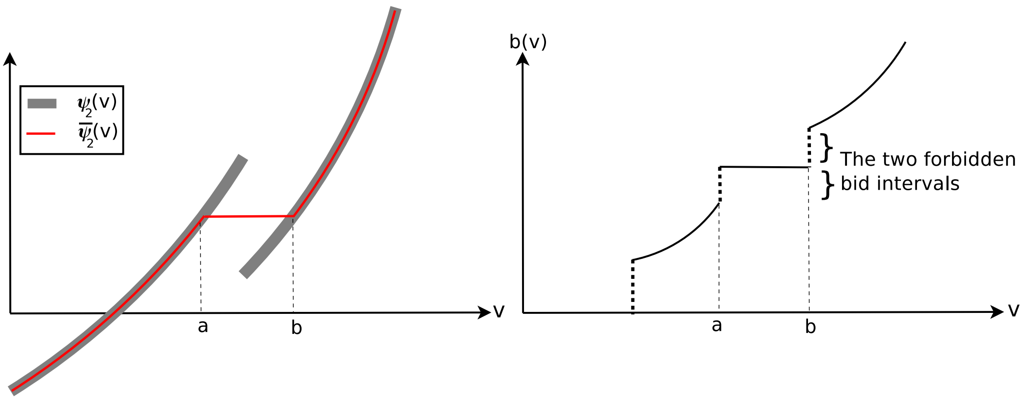

Fig. 1. The ironing procedure.

at the ends of those intervals. In effect, this creates intervals of bids that are suboptimal to make at any value; call these bid intervals "forbidden". In order to implement the mechanism as an all-pay auction, we identify the forbidden bid intervals; then we round every bid in a forbidden bid interval down to the closest "allowed" bid, and distribute the reward equally among the highest bidders (subject to an appropriate reserve value given by $\bar{\psi}_n^{-1}(0)$). We therefore get the following theorem:

Theorem 3.6. *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 bid intervals, 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. All agents pay their bids.*

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 bid intervals. 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 $1/4$. Fig. 1 shows the virtual value function ψ_2 and its ironed counterpart. 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 has two forbidden bid intervals, namely $[1.100, 1.199)$ and $(1.199, 1.310]$, with the intermediate value of 1.199 being allowed. The two forbidden bid intervals correspond to the two discontinuities in the probability of allocation at the end points of the ironed interval.

Asymmetric contests Recall that in deriving the optimal all-pay auction above, we focused on the class of all symmetric contests. For objectives such as the sum of payments or social welfare, symmetric environments (i.e. i.i.d. values) admit optimal auctions that are symmetric or anonymous. This is because the set of monotone allocation functions is convex and the objective is linear. The maximum payment objective is not a concave objective, and so the resulting optimization problem is not a convex optimization problem. So even in symmetric settings asymmetric all-pay auctions can obtain higher expected maximum payment than symmetric all-pay auctions.

We now present an example that exhibits this. Consider two contestants with values drawn i.i.d. from the interval $[0, 1]$ according to the distribution $F(x) = x^{1.5}$. The optimal symmetric auction described above sets a reserve value of $(0.25)^{1/3} = 0.63$ (which translates to a reserve bid of $(0.25)^{5/6} = 0.315$) and serves the highest bidder who exceeds this reserve bid. This gives an expected maximum payment of 0.396.

We now define a better auction that favors contestant 1 over contestant 2. The rules of the contest in value space are as follows. When contestant 1's value is more than 0.75 we serve him irrespective of contestant 2's value, otherwise we serve the contestant with the higher value subject to a reserve value of 0.63. This allocation rule creates a discontinuous increase in the expected allocation probability of player 1 at 0.75, and hence a discontinuity in his bid function at 0.75. In bid space, this corresponds to the following contest:

1. We set a reserve bid of 0.315 as before.
2. All bids of contestant 1 in the range $[0.418, 0.681)$ get rounded down to 0.418.
3. When contestant 1 bids at least 0.681 he wins irrespective of 2's bid.
4. Otherwise, the highest bidder wins with ties broken in favor of contestant 2.

By guaranteeing victory for contestant 1 beyond a certain bid, the above auction encourages contestant 1 to bid higher, thus boosting maximum payment. Since the objective is maximum payment, this type of bias is useful: the asymmetric auction obtains a smaller sum of payments but a larger expected maximum payment of 0.397. We remark that in a real-world setting with a priori identical agents, favoring one agent over another may be socially unacceptable.

4. Utilization and approximation of the highest-bid-wins auction

As noted previously, the maximum payment of a second- or first-price auction is equal to its sum of payments. 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 highest-bid-wins all-pay auctions with reserve. When the value distributions are sum-of-payments regular, this implies a bound on the utilization ratio of the optimal all-pay auction. It further allows us to find a simple all-pay auction that approximates ideal procurement.

As discussed in Section 2 the unique Bayes–Nash equilibrium of highest-bid-wins all-pay auction can be solved for by analyzing interim expected payments of the revenue-equivalent second-price variant of the auction. This analysis enables our proof that the utilization ratio of these auctions is at most two.

Theorem 4.1. *Any highest-bidder-wins reserve-price all-pay auction $\mathcal{A}$ satisfies $\mathbf{SP}[\mathcal{A}] \leq 2\mathbf{MP}[\mathcal{A}]$, i.e., its utilization ratio is bounded by two.*

Proof. Theorem 2.7 implies that there is a unique BNE with symmetric bid function $b(v)$. We can write the expected sum-of-payments 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. (Recall that the values are drawn with density f on the range $[\underline{v}, \bar{v}]$.)

$$\mathbf{SP}[\mathcal{A}] = \underbrace{\sum_i \int_{\underline{v}}^{\bar{v}} b(v)\mathbf{Pr}_{\mathbf{v}_{-i}}\left[v = v_{(1)}\right] f(v)\,dv}_{A} + \underbrace{\sum_i \int_{\underline{v}}^{\bar{v}} b(v)(1 - \mathbf{Pr}_{\mathbf{v}_{-i}}\left[v = v_{(1)}\right]) f(v)\,dv}_{B}.$$

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 (Lemma 2.4). Let $g(v)$ denote the expected payment in the second-price auction with reserve, given that v is the highest agent value. Then we get that $b(v) = g(v)\mathbf{Pr}_{\mathbf{v}_{-i}}\left[v = v_{(1)}\right]$.

Now we can write $A - B$ as

$$\begin{aligned} A - B &= \sum_i \int_{\underline{v}}^{\bar{v}} b(v)(2\mathbf{Pr}_{\mathbf{v}_{-i}}\left[v = v_{(1)}\right] - 1) f(v)\,dv \\ &= n \cdot \int_{\underline{v}}^{\bar{v}} g(v)F(v)^{n-1}(2F(v)^{n-1} - 1) f(v)\,dv \\ &= n \cdot \int_0^1 g(F^{-1}(t))t^{n-1}(2t^{n-1} - 1)\,dt, \end{aligned}$$

where, in the third equality, we substituted t for $F(v)$.

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. □

Tightness of Theorem 4.1 We now revisit the example in Section 3 with n agents, and each agent' value distributed independently according to the $U[0, 1]$ distribution, and show that the utilization ratio of 2 is tight for this example. In this setting, the second-price auction with no reserve price obtains an expected sum of payments of $\frac{n-1}{n+1}$, and thus the highest-bid-wins all-pay auction (that induces a bid function $b(v) = \frac{n-1}{n} v^n$) obtains the same expected sum of payments. On the other hand,

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

the expected maximum payment in the highest-bid-wins auction can be calculated to be $\frac{n-1}{2n}$ which approaches 1/2 as n increases.

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 sum-of-payments-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.

Nevertheless, we can often bound the utilization ratio of optimal all-pay auctions. For maximum-payment regular distributions, the optimal all-pay auction is a highest-bid-wins auction with a reserve, so Theorem 4.1 applies. For sum-of-payments regular distributions, the sum-of-payments-optimal auction is a highest-bid-wins auction with a reserve price; the optimal all-pay auction has a higher expected maximum payment and a lower expected sum of payments in comparison, so its utilization ratio is once again bounded by Theorem 4.1.

Approximation ratio Recall that the "cost of crowdsourcing" is $\frac{\mathbf{SP}[\mathcal{A}^{\mathrm{SP}}]}{\mathbf{MP}[\mathcal{A}^{\mathrm{MP}}]} \leq \frac{\mathbf{SP}[\mathcal{A}^{\mathrm{SP}}]}{\mathbf{MP}[\mathcal{A}]}$, where $\mathcal{A}^{\mathrm{SP}}$ is the sum-of-payments-optimal auction, $\mathcal{A}^{\mathrm{MP}}$ is the all-pay auction optimal for maximum payment, and $\mathcal{A}$ is any all-pay auction. We now use the bound on utilization ratio to prove that the cost of crowdsourcing is always small—no more than 4 in general, and no more than 2 when the value distribution is sum-of-payments regular. In particular, we recall that for regular distributions highest-bidder-wins reserve-price auctions are sum-of-payments-optimal (Theorem 2.2). Corollary 4.2 below then follows from Theorem 4.1 and revenue equivalence. For irregular distributions, Chawla et al. (2010) show that highest-bidder-wins auctions with an anonymous reserve price are within a factor of 2 of optimal (Theorem 2.3), and this gives Corollary 4.3.

Corollary 4.2. *For sum-of-payments-regular value distributions, the highest-bid-wins all-pay auction with an appropriate reserve achieves an approximation ratio at most* 2.

Corollary 4.3. *The highest-bid-wins all-pay auction with an appropriate reserve achieves an approximation ratio at most* 4.

The example with uniform distributions that was described earlier shows that the approximation factor in Corollary 4.2 is tight.

5. Prior-independent approximation

As we show above, optimal crowdsourcing contests depend on knowing the agents' value distribution. To what 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 that for distributions that are sum-of-payments regular, the simple highest-bidder-wins contest obtains an approximation ratio of $2\left(1+\frac{1}{n-1}\right)$, thus limiting the power of distributional knowledge.

For the standard objective of maximizing the expected sum of payments, Bulow and Klemperer (1996) showed that for i.i.d. value distributions that are sum-of-payments regular, it is better to run a second-price 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 sum-of-payments regular, the highest-value-wins auction with no reserve price on n agents is within a factor of $(1-1/n)$ of the optimal mechanism for the sum of payments objective. This combined with Theorem 4.1 gives us the following theorem.

Theorem 5.1. *For i.i.d. distributions that are sum-of-payments regular, the highest-bid-wins all-pay auction without a reserve bid obtains an approximation ratio of* $2\left(1+\frac{1}{n-1}\right)$.

We remark that for the highest-value-wins auction without reserve prices, the sum of payments converges to the optimal as more and more agents are added. However for all-pay auctions with the maximum-payment objective adding more and more agents does not improve the approximation ratio beyond 2.

Appendix A. Proof of Theorem 3.1

Let the agent values be distributed independently according to distribution function F, with density function f. Consider the rank-based-reward allocation rule $\mathcal{A}=(a_1,\ldots,a_k,0,\ldots,0)$, i.e., the agent with the i-th highest bid gets a_i fraction of the reward if $i \leq k$, and 0 otherwise. The rewards are normalized so that $\sum_{i=1}^{k} a_i = 1$. Since $a_k \geq a_{k+1}$ for all $k<n$, Theorem 2.7 implies that there is a unique symmetric equilibrium wherein agents are assigned to rewards in the same order as their values; let $b(\cdot)$ denote the bid function of the equilibrium.

The contribution of bidder i to the maximum payment objective is

$$\mathbf{MP}_i[\mathcal{A}] = \int_{\underline{v}}^{\bar{v}} b(v_i)\mathbf{Pr}_{\mathbf{v}_{-i}}\left[v_i = v_{(1)}\right] f(v_i)\,dv_i$$
$$= \int_{\underline{v}}^{\bar{v}} b(v_i)F(v_i)^{n-1}\,dF(v_i)$$

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 the truthful direct revelation mechanism with the same allocation rule. In the truthful direct revelation mechanism with the same allocation rule, 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 agent value given the r-th highest value 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 $v_{jr}(z) = g(j-r, n-r, z)$. So we get the following expression for $b(z)$.

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

We prove the theorem by showing that increasing a_k to $a_k + h$ and decreasing a_1 to $a_1 - h$ decreases the maximum payment. Formally, let $\tau_k^i(h) = \mathbf{MP}_i[a_1 - h, \ldots, a_k + h, \ldots]$. We show that $\frac{d\tau_k^i(h)}{dh}$ is negative for any k and i. 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 $\frac{d\tau_k^i(h)}{dh}$ 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 $\frac{d\tau_k^i(h)}{dh}$ will be an integral with an integrand of

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

As above, $g(k, n-1, z) < g(k-1, n-1, z)$, and $g(1, n-1, z) > 0$. So the above integrand is negative.

The term corresponding to $r = k$ in $\frac{d\tau_k^i(h)}{dh}$ 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 $\frac{d\tau_k^i(h)}{dh}$ 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.

$$\frac{d\tau_k^i(h)}{dh} \le Q = -\int_0^1 F(z)^{2n-2} g(1, n-1, z)\,dF(z)$$
$$- \binom{n-1}{k-2}\int_0^1 (1-F(z))^{k-2}F(z)^{2n-k} g(1, n-k+1, z)\,dF(z)$$
$$+ \binom{n-1}{k-2}\int_0^1 (1-F(z))^{k-2}F(z)^{2n-k} g(2, n-k+1, z)\,dF(z)$$
$$+ \binom{n-1}{k-1}\int_0^1 (1-F(z))^{k-1}F(z)^{2n-k-1} g(1, n-k, z)\,dF(z)$$

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

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

$$g(2,n,z) = n(n-1)\int_0^z y\frac{f(y)}{F(z)}\left(1-\frac{F(y)}{F(z)}\right)\left(\frac{F(y)}{F(z)}\right)^{n-2} dy$$

$$= z - \left[n\frac{\int_0^z F(t)^{n-1}\,dt}{F(z)^{n-1}} - (n-1)\frac{\int_0^z F(t)^n\,dt}{F(z)^n}\right]$$

We substitute the expression for g into Q.

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

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

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

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

$$\int_0^1 F(z)^{n-2}\left[zF(z)^n - \int_0^z F(t)^n\,dt\right] dF(z)$$

$$+\binom{n-1}{k-1}\int_0^1 (1-F(z))^{k-2}F(z)^{n-1}\left[zF(z)^{n-k+1} - \int_0^z F(t)^{n-k+1}\,dt\right] dF(z)$$

$$-\binom{n-1}{k-1}\int_0^1 (1-F(z))^{k-2}F(z)^{n-1}\left[zF(z)^{n-k} - \int_0^z F(t)^{n-k}\,dt\right] dF(z)$$

$$\ge \binom{n-1}{k-1}(k-2)\int_0^1 (1-F(z))^{k-2}F(z)^{n-1}\left(\int_0^z F(t)^{n-k+1}\,dt\right) dF(z)$$

$$-\binom{n-1}{k-1}(k-2)\int_0^1 (1-F(z))^{k-2}F(z)^{n}\left(\int_0^z F(t)^{n-k}\,dt\right) dF(z)$$

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,

$$n\int_0^1 F(z)^{n-2}\left(\int_0^z tF(t)^{n-1}\,dF(t)\right) dF(z)$$

$$+\binom{n-1}{k-1}(n-k+1)\int_0^1 (1-F(z))^{k-2}F(z)^{n-1}\left(\int_0^z tF(t)^{n-k}\,dF(t)\right) dF(z)$$

$$-\binom{n-1}{k-1}(n-k)\int_0^1 (1-F(z))^{k-2}F(z)^{n-1}\left(\int_0^z tF(t)^{n-k-1}\,dF(t)\right) dF(z)$$

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

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

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

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

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.

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

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

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

References

Archak, N., Sundararajan, A., 2009. Optimal design of crowdsourcing contests. In: International Conference on Information Systems. ICIS.

Barut, Y., Kovenock, D., 1998. The symmetric multiple prize all-pay auction with complete information. Europ. J. Polit. Economy 14 (4), 627–644.

Baye, M.R., Kovenock, D., de Vries, C.G., 1996. The all-pay auction with complete information. Econ. Theory 8, 291–305.

Bulow, J., Klemperer, P., 1996. Auctions versus negotiations. Amer. Econ. Rev. 86, 180–194.

Chawla, S., Hartline, J.D., 2013. Auctions with unique equilibria. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce. EC '13. ACM, New York, NY, USA, pp. 181–196.

Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B., 2010. Multi-parameter mechanism design and sequential posted pricing. In: Proceedings of the 42nd ACM Symposium on Theory of Computing. STOC '10. ACM, New York, NY, USA, pp. 311–320.

Che, Y.-K., Gale, I., 2003. Optimal design of research contests. Amer. Econ. Rev. 93 (3), 646–671.

DiPalantino, D., Vojnovic, M., 2009. Crowdsourcing and all-pay auctions. In: Proceedings of the 10th ACM Conference on Electronic Commerce. EC '09. ACM, New York, NY, USA, pp. 119–128.

Fullerton, R.L., McAfee, R.P., 1999. Auctioning entry into tournaments. J. Polit. Economy 107 (3), 573–605.

Gavious, A., Minchuk, Y., 2014. Revenue in contests with many participants. Oper. Res. Lett. 42 (2), 119–122.

Hartline, J.D., Roughgarden, T., 2008. Optimal mechanism design and money burning. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing. STOC '08. ACM, New York, NY, USA, pp. 75–84.

Megidish, R., Sela, A., 2013. Allocation of prizes in contests with participation constraints. J. Econ. Manag. Strategy 22 (4), 713–727.

Minor, D., 2011. Increasing effort through rewarding the best less (or not at all). Manuscript.

Moldovanu, B., Sela, A., 2001. The optimal allocation of prizes in contests. Amer. Econ. Rev. 91 (3), 542–558.

Moldovanu, B., Sela, A., 2006. Contest architecture. J. Econ. Theory 126 (1), 70–97.

Please cite this article in press as: Chawla, S., et al. Optimal crowdsourcing contests. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.09.001

Myerson, R., 1981. Optimal auction design. Math. Oper. Res. 6, 58–73.

Taylor, C.R., 1995. Digging for golden carrots: an analysis of research tournaments. Amer. Econ. Rev. 85 (4), 872–890.

Yang, J., Adamic, L.A., Ackerman, M.S., 2008. Crowdsourcing and knowledge sharing: strategic user behavior on Taskcn. In: Proceedings of the 9th ACM Conference on Electronic Commerce. EC '08. ACM, New York, NY, USA, pp. 246–255.