Efficient Crowdsourcing Contests

Ruggiero Cavallo
Yahoo! Research
111 West 40th Street
New York, NY 10018
cavallo@yahoo-inc.com

Shaili Jain
Computer Science Department
Yale University
New Haven, CT 06520
shaili.jain@yale.edu

ABSTRACT

A principal seeks production of a good within a limited time-frame with a hard deadline, after which any good procured has no value. There is inherent uncertainty in the production process, which in light of the deadline may warrant simultaneous production of multiple goods by multiple producers despite there being no marginal value for extra goods beyond the *maximum quality* good produced. This motivates a *crowdsourcing* model of procurement. We address efficient execution of such procurement from a social planner's perspective, taking account of and optimally balancing the value to the principal with the costs to producers (modeled as effort expenditure) while, crucially, contending with self-interest on the part of all players. A solution to this problem involves both an algorithmic aspect that determines an optimal effort level for each producer given the principal's value, and also an incentive mechanism that achieves equilibrium implementation of the socially optimal policy despite the principal privately observing his value, producers privately observing their skill levels and effort expenditure, and all acting selfishly to maximize their own individual welfare. In contrast to popular "winner take all" contests, the efficient mechanism we propose involves a payment to every producer that expends non-zero effort in the efficient policy.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Algorithms, Human Factors

Keywords

mechanism design, crowdsourcing, contests, social welfare

1. INTRODUCTION

An increasingly common model of procurement has multiple agents simultaneously produce versions of a desired good at the behest of a principal who seeks the highest-quality version. This model's popularity has soared as part of the Internet phenomenon known as *crowdsourcing*. In recent years, the number of websites making and facilitating open calls for solutions to tasks such as logo design, software development, and image labeling has grown tremendously; examples include Amazon Mechanical Turk, Taskcn, Topcoder, 99designs, Innocentive, CrowdCloud, and CrowdFlower, to name a few. These developments have reinvigorated a line of research in the field of microeconomics known as *contest design*. The model of a contest matches the standard approach to crowdsourcing: many agents simultaneously exert effort to submit a solution in competitive pursuit of a reward, where the "winner" is dependent on the relative submission qualities.[1] To date, most previous research has focused on maximizing the principal's utility: the goal is to to procure the best submission for the lowest possible price. In contrast, we here consider the crowdsourcing problem from an efficiency standpoint, adopting the perspective of a *social planner* that seeks to maximize social welfare. And in contrast with the standard "winner take all" contest methods, the efficient scheme we derive involves a payment for every agent that expends non-zero effort in the efficient policy.

To motivate the crowdsourcing paradigm from an efficiency standpoint, assuming rational players,[2] uncertainty and deadlines must play a central role. If these factors were not present then the redundant production inherent to the paradigm would be purely wasteful; one could (and should) alternatively order production sequentially rather than simultaneously, stopping further production when the costs are no longer outweighed by the expected gains given the "quality in hand". So we adopt a model in which the principal seeks production of a good within a single unit of time (corresponding to the span required for production of a single good), after which any goods obtained are of no value. There is inherent uncertainty in production, which may warrant simultaneous production of multiple goods. However, if multiple goods are produced there is no marginal value beyond that of the *maximum quality* good produced. Producers (henceforth, "agents") may have varying *skill* and also make a choice about how much *costly effort* to expend on production: higher effort and skill leads to production of a good with greater expected quality, all else equal.

The principal and all agents are presumed to be self-

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

[1] *All-pay auctions* are also highly related, with the key difference being that each agent's cost there is a payment that generates "revenue" enjoyed by the seller, whereas in a typical contest only the highest-quality-producing agent directly benefits the principal.

[2] While we make this assumption, we do not deny that in practice "irrational" behavioral factors may contribute to the success of many crowdsourcing marketplaces.

interested, which gives rise to a problem of incentives. Achieving efficiency thus has two components: determining an optimal effort policy for the agents, given the value of the principal and the various agent skill levels; and designing a payment mechanism in which no individual can improve his expected utility by misreporting private information (or exerting effort other than what the efficient policy prescribes, in the case of the agents). For a setting where agents have no private information about skill, we derive an efficient, individually rational, and budget-balanced solution—a novel achievement to the best of our knowledge. When skill is private we prove that extending this result is impossible, but we show how a result from the recent mechanism design literature can be applied to yield success if agents can be forced to make commitment decisions prior to learning their skill levels, i.e., if *ex ante* individual rationality suffices.

1.1 Related Work

There has recently been work explicitly addressing the theory of crowdsourcing in a model, like ours, where agents have private skill information and choose an effort level. DiPalantino and Vojnovic [2009] make the connection to all-pay auctions and model a market with multiple contests, considering the principal's optimization problem in the limit-case as the number of agents and contests goes to infinity. Archak and Sundararajan [2009] and Chawla, Hartline, and Sivan [2011] focus on the design of a single contest; the problem consists of determining how many prizes should be awarded, and of what value. Chawla et al. make the connection between crowdsourcing contests and optimal auction design, finding that the optimal crowdsourcing contest is a virtual valuation maximizer.

However, these papers consider a *deterministic* model of quality as a function of effort and skill, under which, if the principal's value is proportional to the *maximum* quality over the produced goods, the crowdsourcing paradigm itself is not well-motivated from an efficiency standpoint. And while in this paper we are concerned with *social* welfare, this prior work is geared towards maximizing utility of the principal alone and is unconcerned with the cost to the agents. This focus is characteristic of the broader literature: in both computer science and economics prior work has, for the most part, focused on maximizing submission quality, whether it be the total sum of submission qualities [Moldovanu and Sela, 2001; 2006; Minor, 2011], the top k submissions less the monetary reward [Archak and Sundararajan, 2009], or only the highest quality submission [Moldovanu and Sela, 2006; Chawla *et al.*, 2011].

This ties in with the extensive literature in economics devoted to the design of optimal *contests*. Many of these works consider a contest model where the prize value is known to all players [Tullock, 1980; Moulin, 1986; Baye *et al.*, 1996], while others adopt a model of incomplete information with respect to the prize [Weber, 1985; Hilman and Riley, 1989; Krishna and Morgan, 1997]. There has also been work on research tournaments that award a single prize. For instance in [Fullerton and McAfee, 1999] agents have a cost of production—drawn from a known distribution—that becomes common knowledge after a first round in which agents simultaneously decide whether to participate; then in a second round agents decide how much effort to exert given the common-knowledge costs.

A related line of work uses contests to extract effort under a hidden action [Lazear and Rosen, 1981; Green and Stokey, 1983; Nalebuff and Stiglitz, 1983]. Similar to our work, the output is a stochastic function of the unobservable effort, but the setting is different in that the principal obtains value from the *cumulative* effort of the agents, rather than just the maximum result.

Finally, this work overlaps with the broader agenda of incentives in peer production systems, where there has been work addressing incentives in question and answer forums, human computation, etc. [Jain and Parkes, 2008; Jain *et al.*, 2009; Ghosh and McAfee, 2011; Ghosh and Hummel, 2011].

1.2 Preliminaries

In the crowdsourcing paradigm multiple units of a good are simultaneously produced and submitted to a principal. There is a set of agents $I = \{1, \dots, n\}$ capable of producing goods, where each $i \in I$ has private skill level $s_i \in [0, 1]$. Agents can expend variable effort on production of the good. If an agent attempts production, a good is produced with quality that is a priori uncertain but is a function of the agent's skill and effort expended.

Quality is identified with value to the principal in dollar-terms. The probability distribution over *relative* quality, given any skill and effort levels, is publicly known, but the *absolute* quality in terms of value to the principal is not. The principal has private type $v \in \Re^+$, a scale factor corresponding to his value for the maximum quality good that could possibly be produced, and this, given the known distribution over relative qualities, defines the distribution over absolute quality (henceforth just "quality") corresponding to the principal's value.[3] An effort level δ_i is identified with the *dollar value in costs ascribed to it by agent i*. For simplicity we assume that $\delta_i \in [0, 1]$, $\forall i \in I$.[4] Then, given a $v \in \Re^+$, skill level s_i, and effort level $\delta_i \in [0, 1]$, we denote the p.d.f. and c.d.f. over resulting quality as $f^v_{s_i,\delta_i}$ and $F^v_{s_i,\delta_i}$, respectively. We assume symmetry across bidders in the sense that skill is the only differentiating factor; i.e., for two agents with the same skill level applying the same effort, the distribution over quality is the same (though there is no presumed correlation so the resulting quality may differ).

We will make the natural assumption that for an agent with any given skill level, more effort has first-order stochastic dominance over less effort with respect to quality, i.e.:

$$\forall s_i,\ \forall 0 \leq \delta_i < \delta_i' \leq 1,\ \forall x \in [0, v],\ F^v_{s_i,\delta_i}(x) \geq F^v_{s_i,\delta_i'}(x), \quad (1)$$

and also that, given any effort level, more skill has first-order stochastic dominance over less skill with respect to quality:

$$\forall \delta_i,\ \forall 0 \leq s_i < s_i' \leq 1,\ \forall x \in [0, v],\ F^v_{s_i,\delta_i}(x) \geq F^v_{s_i',\delta_i}(x) \quad (2)$$

Because agents are self-interested there is a problem of incentives: v and s_i (for each i) are private information, and expended effort is privately observed. We adopt a quasi-linear utility model and assume all players are risk-neutral. Given our identification of the quality of the good with the dollar value ascribed to it by the principal, and effort level δ_i with the dollar value in costs ascribed to it by agent i, quasilinearity implies that the principal's utility equals the quality of the good procured minus any payments he must make, and each agent's utility equals any payment he receives minus the effort he expends.

[3] E.g., if quality q ranges in $[0, 1]$, absolute quality is vq.

[4] The specific range of effort levels is not conceptually important; it is the relationship between the effort levels and v that is relevant for determining an optimal policy.

We are concerned with the socially optimal choice of effort level for each agent, where in light of our utility model the appropriate optimization is the *maximum* quality level of the goods produced minus total effort expended; i.e., letting $\mathcal{Q}_i(v, s_i, \delta_i)$ be a random variable representing the quality level produced by $i \in I$ who has skill level s_i and expends effort δ_i (with v the principal's value), we seek to maximize:

$$\mathbb{E}[\max_{i \in I} \mathcal{Q}_i(v, s_i, \delta_i)] - \sum_{i \in I} \delta_i \tag{3}$$

An effort policy is a function of the principal's value and the agents' skill levels. We let $\delta^*(v, s)$ denote an *efficient* policy, i.e., a vector of effort levels that maximizes Eq. (3) given values of v and $s = (s_1, \ldots, s_n)$; when context is clear we will write δ_i^* as shorthand for $\delta_i^*(v, s)$. Given our quasilinear utility model, a policy δ^* that maximizes Eq. (3) maximizes the expected sum of utilities and is Pareto efficient.

At various points we will consider a restricted setting where skill is constant (and publicly known) throughout the population of agents; we call this the *constant skill case*.

2. EFFICIENT EFFORT POLICIES

In this section we address the problem of computing an efficient policy given full knowledge of the principal's value v and agent skill levels s, and given that agents will execute the effort policy that is prescribed. We defer to Section 3 the question of how to implement such a policy in the context of a principal and agents that are self-interested and strategic. We will make heavy use of the following lemma, which demonstrates sufficient conditions under which *extreme-effort* policies—those that involve only total (1) or null (0) effort by each agent—are optimal. In this section we make the technical assumption that the cumulative distribution over quality, evaluated at any particular quality level, is differentiable with respect to effort δ_i.

LEMMA 1. *For arbitrary $i \in I$ with arbitrary skill s_i, for arbitrary skill and effort levels of the other agents and value v for the principal, fixing an arbitrary effort policy for agents other than i, amongst effort levels within arbitrary interval $[a, b] \subseteq [0, 1]$ it is either optimal for i to expend effort $\delta_i = a$ or optimal for i to expend effort $\delta_i = b$ if the following holds:* $\forall \beta \in [0, v]$, $\forall \epsilon \in [a, b)$,

$$-\frac{\partial}{\partial \delta_i}\Big(\int_\beta^v F_{s_i,\delta_i}^v(x)\,dx\Big)\Big|_{\delta_i=\epsilon} \geq 1 \tag{4}$$

$$\Rightarrow -\frac{\partial}{\partial \delta_i}\Big(\int_\beta^v F_{s_i,\delta_i}^v(x)\,dx\Big)\Big|_{\delta_i=k} \geq 1,\ \forall k \in [\epsilon, b] \tag{5}$$

PROOF. For arbitrary agent $i \in I$, consider arbitrary β representing the maximum quality that is to be realized by the production of the other agents—this is a priori unknown, but a result for *arbitrary* β will demonstrate that regardless of its realization the result holds. Then the expected marginal impact on efficiency from i exerting effort δ_i equals:

$$\int_\beta^v f_{s_i,\delta_i}^v(x)(x - \beta)\,dx - \delta_i \tag{6}$$

$$= (x - \beta)F_{s_i,\delta_i}^v(x)\Big|_{x=\beta}^{x=v} - \int_\beta^v F_{s_i,\delta_i}^v(x)\,dx - \delta_i \tag{7}$$

$$= v - \beta - \int_\beta^v F_{s_i,\delta_i}^v(x)\,dx - \delta_i \tag{8}$$

The first step above is integration by parts. To find the maximum with respect to effort, we consider the derivative with respect to δ_i, i.e.,

$$\frac{\partial}{\partial \delta_i}\Big(v - \beta - \int_\beta^v F_{s_i,\delta_i}^v(x)\,dx - \delta_i\Big) \tag{9}$$

$$= -\frac{\partial}{\partial \delta_i}\Big(\int_\beta^v F_{s_i,\delta_i}^v(x)\,dx\Big) - 1 \tag{10}$$

If as δ_i increases this derivative never changes from positive to negative (this is the condition of the Lemma, in Eqs. (4) and (5)), then the maximum lies at one of the extremes, $\delta_i = a$ or $\delta_i = b$, which completes the proof. □

COROLLARY 1. *For environments where the quality distribution functions satisfy the relationship of Eqs. (4–5) in Lemma 1 over the full range of effort levels ($[a, b] = [0, 1]$), an efficient effort policy consists of full-effort participation by a subset of the agents and non-participation by the others.*

Note that the lemma holding for the interval $[0, 1]$ is sufficient *but not necessary* for the optimal policy to involve only extreme-effort (i.e., effort 0 or 1 by all agents). For instance if the condition of the lemma (Eqs. (4) and (5)) does not hold for some β, yet $\forall \delta_i$ the expected quality output for i is less than $\beta - \delta_i$, then the optimal policy would involve non-participation (0 effort) by i when other agents achieve quality β, and so it may still be the case that an optimal policy never involves intermediate effort.

For any constant skill environment (say skill equals $\hat{s}$ for each agent) where we can establish that an extreme-effort policy is optimal, fully determining an efficient policy is easy. We simply need to compute:

$$m^* = \operatorname*{arg\,max}_{m \in \{0,\ldots,n\}} \Big[m \int_0^v F_{\hat{s},1}^v(x)^{m-1} f_{\hat{s},1}^v(x) x\,dx - m\Big] \tag{11}$$

m^* agents will participate with full-effort and the other $n - m^*$ will not participate (i.e., will apply 0 effort).

When skill is not constant, by Eq. (2) having an agent participate who has *less* skill than one who does not participate could never be optimal. So more generally, for any setting where extreme-effort has been established as efficient we can determine a precise optimal policy by iteratively considering each agent in *decreasing order of skill*, accepting agents for (full-effort) participation until stopping and accepting no more in the ordered list.

In the rest of the section we will show that extreme-effort policies are optimal in important canonical settings, but we first observe that this is not universally the case. Imagine that effort $\delta_i \in [0, 0.05)$ yields quality 0 (with certainty), $\delta_i \in [0.05, 0.3)$ yields quality 0 with probability 0.8 and quality 0.9 with probability 0.2, and $\delta_i \in [0.3, 1]$ yields quality 0.6 with probability 0.8 and quality 0.9 with probability 0.2. An optimal policy for two agents has one agent expend effort 0.3 and the other expend effort 0.05.

2.1 Uniformly distributed quality

We now look at specific distributional settings, starting with one in which quality is uniformly distributed between 0 and the product of the principal's value and the agent's skill and effort. That is, F_{s_i,δ_i}^v for each $i \in I$ is the uniform distribution over $[0, \delta_i s_i v]$, i.e.,

$$f_{s_i,\delta_i}^v(x) = \begin{cases} \frac{1}{\delta_i s_i v} & \text{if } x \in [0,\ \delta_i s_i v] \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

We call this the *uniformly distributed quality case*. The range of possible qualities (and the expected quality) increases linearly with skill and effort. We can use Lemma 1 to show that an extreme-effort policy is optimal here.

LEMMA 2. *For the uniformly distributed quality case, for arbitrary skill levels, there is an efficient policy in which each agent $i \in I$ exerts either no effort $\delta_i = 0$ or full effort $\delta_i = 1$.*

PROOF. Consider arbitrary agent $i \in I$ and arbitrary $\beta \in [0, s_i v]$. If we can show that effort $\delta_i = 0$ or $\delta_1 = 1$ for i would yield optimal expected marginal efficiency even if we knew that the other agents would achieve quality β, then the theorem follows. Note that in the case of uniformly distributed quality no effort level less than or equal to $\frac{\beta}{s_i v}$ could possibly be optimal because there would be 0 probability of improving the maximum quality over β (and so if $\beta > s_i v$ then 0 effort by i is clearly optimal). We will now prove the result by using Lemma 1. For arbitrary $\delta_i \in [\frac{\beta}{s_i v}, 1]$:

$$\int_\beta^v F^v_{\delta_i}(x)\,dx = \int_\beta^{\delta_i s_i v} \frac{x}{\delta_i s_i v}\,dx + \int_{\delta_i s_i v}^v 1\,dx \quad (13)$$

$$= \frac{x^2}{2\delta_i s_i v}\Big|_{x=\beta}^{x=\delta_i s_i v} + v - \delta_i s_i v \quad (14)$$

$$= \frac{\delta_i^2 s_i^2 v^2}{2\delta_i s_i v} - \frac{\beta^2}{2\delta_i s_i v} + v - \delta_i s_i v \quad (15)$$

$$= v - \frac{\delta_i s_i v}{2} - \frac{\beta^2}{2\delta_i s_i v} \quad (16)$$

Then:

$$-\frac{\partial}{\partial \delta_i}\Big(\int_\beta^v F^v_{s_i,\delta_i}(x)\,dx\Big) = \frac{s_i v}{2} - \frac{\beta^2}{2 s_i v \delta_i^2}, \quad (17)$$

and this is at least 1 if and only if:

$$\delta_i^2 \geq \frac{\beta^2}{2 s_i v - (s_i v)^2} \quad (18)$$

If this holds for $\delta_i = \epsilon \geq \frac{\beta}{s_i v}$ it holds for $\delta_i = k$ for all $k > \epsilon$. Therefore the distribution satisfies Lemma 1, which tells us that the optimum over the range $[\frac{\beta}{s_i v}, 1]$ occurs at either $\frac{\beta}{s_i v}$ or 1. Since no effort level on the interval $(0, \frac{\beta}{s_i v}]$ could yield positive expected utility and thus be better than effort level 0, the global optimum lies at either $\delta_i = 0$ or $\delta_i = 1$, which completes the proof. □

Then in the constant skill case (with skill normalized to 1) we can compute the optimal number of participants as follows.

THEOREM 1. *For the constant skill, uniformly distributed quality case, a mechanism that elicits maximum-effort participation by m^* arbitrary agents (and 0-effort participation by others) is efficient, where:*

$$m^* = \begin{cases} \lfloor\sqrt{v}\rfloor - 1 & \textit{if } \lfloor\sqrt{v}\rfloor^2 + \lfloor\sqrt{v}\rfloor > v \\ \lfloor\sqrt{v}\rfloor & \textit{otherwise} \end{cases} \quad (19)$$

PROOF. By Lemma 2 an optimal policy will involve full effort by some number $m \in \{0, \dots, n\}$ agents and 0 effort by the other $n - m$ agents. If m agents participate with full-effort the expected efficiency equals the expected maximum of m draws from $U[0, v]$ minus m, i.e.:

$$\int_0^v \frac{1}{v} m x \Big(\frac{x}{v}\Big)^{m-1} dx - m \;=\; \frac{m}{m+1} v - m \quad (20)$$

Though m is discrete, imagining it as a continuous variable yields a parabola with a single maximum. Taking the derivative with respect to m, we get $\frac{v}{(m+1)^2} - 1$, which has a single positive root at $m = \sqrt{v} - 1$. But since m can only take integer values, when $\sqrt{v}$ is not an integer we have to consider both $\lfloor\sqrt{v} - 1\rfloor$ and $\lceil\sqrt{v} - 1\rceil$ (i.e., $\lfloor\sqrt{v}\rfloor$). The expected value increase of adding a $\lfloor\sqrt{v}\rfloor^{th}$ agent to a group of $\lfloor\sqrt{v} - 1\rfloor$ equals, considering Eq. (20):

$$\Big(\frac{\lfloor\sqrt{v}\rfloor}{\lfloor\sqrt{v}\rfloor + 1} v - \lfloor\sqrt{v}\rfloor\Big) - \Big(\frac{\lfloor\sqrt{v}\rfloor - 1}{\lfloor\sqrt{v}\rfloor} v - (\lfloor\sqrt{v}\rfloor - 1)\Big) \quad (21)$$

$$= \frac{1}{\lfloor\sqrt{v}\rfloor(\lfloor\sqrt{v}\rfloor + 1)} v - 1 \quad (22)$$

This is at least 0 if and only if $\lfloor\sqrt{v}\rfloor^2 + \lfloor\sqrt{v}\rfloor \leq v$, and so the maximum is as characterized by Eq. (19). □

Figure 1 provides a graphic depiction of how the optimal number of agents that should participate relates to the principal's value, for $v \in [0, 100]$.

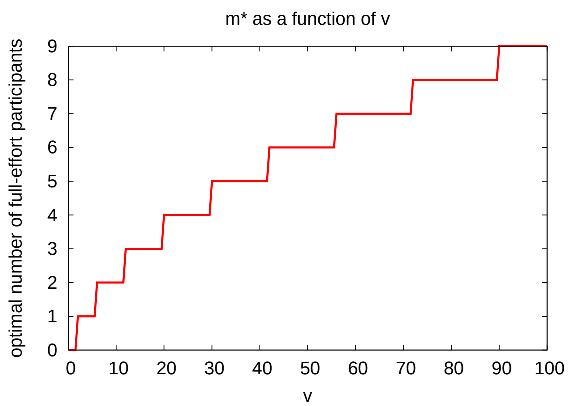

Figure 1: Optimal number of agents to produce a good (with full effort) in the uniformly distributed quality case, as a function of the principal's value v.

2.2 Normally distributed quality

While for more complex distributions beyond the uniform case we would have difficulty demonstrating similar results analytically, we can query whether Lemma 1 holds experimentally. We now consider quality that is normally distributed, over a bounded interval, with mean increasing proportional to effort and skill. Specifically, consider the truncated normal distribution over the interval $[0, v]$, with location parameter μ equal to $\delta_i s_i v$ and scale parameter σ equal to $v/8$; this distribution is illustrated in Figure 2 for various effort levels.

For arbitrary v, β, and δ_i (assuming a fixed $s_i = 1$) we can computationally approximate $\int_\beta^v F^v_{s_i,\delta_i}(x)\,dx$ and accordingly evaluate whether the conditions of Lemma 1 are satisfied. But a more direct approach is equally tractable here: we can simply compute the expected marginal efficiency, given any β, of effort for arbitrarily fine discretizations of the effort space $\delta_i \in [0, 1]$. Then, the next step is to use this approach to check whether extreme-effort is optimal for *all* β in the range $[0, v]$, as this would imply that regardless of the quality obtained by agents other than arbitrary $i \in I$, for i extreme-effort (0 or 1) is optimal. We

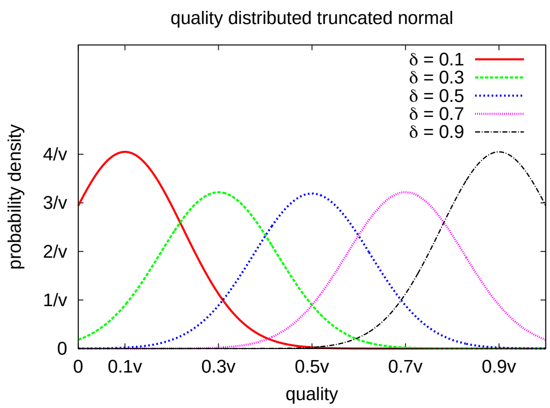

Figure 2: Probability densities over quality for varying degrees of effort, for an agent with skill level 1. Truncated normal with $\mu = \delta_i v$ and $\sigma = v/8$.

checked this by again discretizing the search space; this time the space in question is that of possible (v, β) pairs where β is constrained to fall within $[0, v]$. The results suggest that for all values of v above a very low threshold (2.8), there is no possibility that anything other than extreme effort could be optimal. And we emphasize that with this approach we are only checking certain *sufficient conditions* for optimality of extreme-effort. Finally, given that an extreme-effort policy is optimal, we can easily compute an optimal set of full-effort participants. For the constant skill case this is done according to Eq. (11); see Figure 3 for the results.

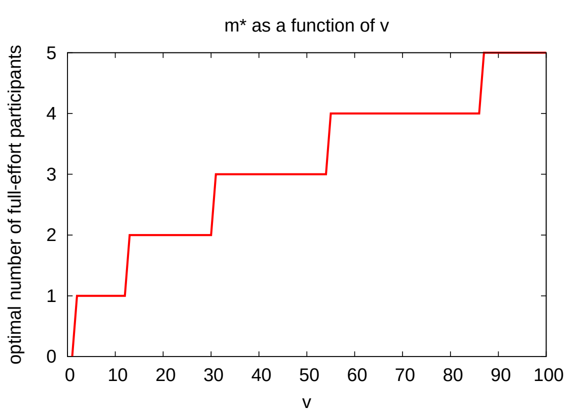

Figure 3: Optimal number of agents that should produce a good (with full effort) as a function of the principal's value v. For the constant skill (equal to 1), truncated ($[0, v]$) normally distributed quality case with $\mu = \delta_i v$ and $\sigma = v/8$.

3. INCENTIVES

We now consider the problem of implementing an efficient policy in a context of selfish players. Since utility in our setting is quasilinear and thus transferable, we can use monetary payments as a tool. Through payments we seek to establish an equilibrium where no agent can gain by doing anything other than what the mechanism asks, as follows:

Definition 1 (Incentive compatibility). *A mechanism is incentive compatible if and only if for each player i, given that all other players abide by the mechanism's prescriptions, i's expected utility can never be improved by doing other than what the mechanism prescribes.*

This definition is a generalization of the standard "truthful reporting" definition that is sufficient for mechanisms that only involve sharing of private information. In our setting, one player (the principal) must share private information *truthfully*, while others (the agents) must behave *faithfully* according to what the mechanism prescribes and will also have to share private information truthfully if skill is variable and private knowledge.[5] An incentive compatible mechanism gives us reason to believe that the outcome the mechanism prescribes will occur, given rational agents. But for a mechanism that makes payments there are additional constraints: the mechanism should be weakly *budget-balanced*, in expectation never paying out more than it takes in, since otherwise external subsidies would be required for its implementation. The mechanism should also be *individually rational*, meaning no agent should have negative expected utility in equilibrium from truthfully participating.

3.1 Constant skill

We first consider a context of *constant skill*, where two agents exerting the same effort produce quality according to the same (known) distribution. Recall that an efficient effort policy $\delta^* = (\delta_1^*, \ldots, \delta_n^*)$ is a function of the principal's value and the vector of agents' skill levels s; so in constant skill settings the only "variable" relevant to computation of δ^* is v, and we omit s from all notation. The mechanism we propose is efficient and incentive compatible in such settings, without running a deficit or violating individual rationality. It defines payments that, in some cases, depend on a priori *expected* quality for a given effort level. Recall notation $\mathcal{Q}_i(v, \delta_i)$ for the random variable representing the quality produced when agent i expends effort δ_i, given the principal's value v; $Q_i(v)$ denotes an actual quality level *realized* by i. $Q(v)$ denotes the vector $(Q_1(v), \ldots, Q_n(v))$ and $Q_{-i}(v)$ denotes $(Q_1(v), \ldots, Q_n(v))$ with $Q_i(v)$ excluded; analogously $\mathcal{Q}(v, \delta)$ denotes $(\mathcal{Q}_1(v, \delta_1), \ldots, \mathcal{Q}_n(v, \delta_n))$ and $\mathcal{Q}_{-i}(v, \delta_{-i})$ denotes the same excluding $\mathcal{Q}_i(v, \delta_i)$. For any vector x we let $x^{(k)}$ denote the k^{th} highest element of x.

Definition 2. (Constant skill efficient crowdsourcing (CSEC) mechanism) *The principal reports v and then efficient effort levels $\delta_1^*, \ldots, \delta_n^*$ are computed. Each agent i is instructed to expend effort δ_i^* on production, and goods are produced with quality levels $Q(v) = (Q_1(v), \ldots, Q_n(v))$. The principal is charged:*

$$\sum_{i \in I} \delta_i^*, \tag{23}$$

agent $h = \arg\max_{i \in I} Q_i(v)$ is paid:

$$\delta_h^* + Q_h(v) - Q^{(2)}(v) - \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*) - \mathcal{Q}_{-h}^{(1)}(v, \delta_{-h}^*)], \tag{24}$$

and each other agent $i \in I \setminus \{h\}$ is paid:

$$\delta_i^* - \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*) - \mathcal{Q}_{-i}^{(1)}(v, \delta_{-i}^*)] \tag{25}$$

[5] This dual-nature incentive situation appears in many other scenarios; see [Shneidman and Parkes, 2004] and [Cavallo and Parkes, 2008] for precedents in the literature.

The principal pays the sum of the prescribed effort levels; each agent is paid his prescribed effort minus the *expected* difference between the highest quality level overall and the highest quality level achieved by the other agents; each agent is also paid the difference between the *actual* highest quality level produced overall and the highest quality level produced by the other agents—this value is 0 for all agents except he who produces the highest quality good, and thus that agent (h) ends up with a "bonus" $(Q_h(v) - Q^{(2)}(v))$.

Since to compute payments *actual* quality must be known, one can either assume the quality of a produced good given any v is publicly observable or, alternatively, in settings where this is unrealistic the mechanism can be slightly (and harmlessly) modified to have the principal report the quality level of each produced good.

THEOREM 2. *The CSEC mechanism is efficient, incentive compatible, individually rational, and budget-balanced in expectation for constant skill settings.*

PROOF. We start by showing incentive compatibility. The expected utility of the principal, given that he announces $\hat{v}$ and that the agents abide by the mechanism, is:

$$\mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*(\hat{v}))] - \sum_{i \in I} \delta_i^*(\hat{v}) \tag{26}$$

By efficiency of the computed effort levels (see Eq. (3)), this quantity is maximized with truthful report $\hat{v} = v$.

Now consider arbitrary agent $i \in I$, assume that the principal is truthful and other agents abide by the mechanism and expend effort δ_{-i}^*, and let δ_i denote i's chosen effort level. i is paid the aggregate utility of the other agents minus a quantity completely independent of his behavior; i.e., omitting v from the notation with truthful v understood:

$$\Big[Q^{(1)} - \sum_{j \in I \setminus \{i\}} \delta_j^*\Big] - \tag{27}$$

$$\Big[Q_{-i}^{(1)} - \sum_{j \in I \setminus \{i\}} \delta_j^* - \delta_i^* + \mathbb{E}[\mathcal{Q}^{(1)}(\delta^*) - \mathcal{Q}_{-i}^{(1)}(\delta_{-i}^*)]\Big] \tag{28}$$

Recall that δ_i^* is computed independent of the behavior of i (or any other agent). Therefore i's *expected* utility equals a quantity independent of his control (−Eq. (28)), plus:

$$\mathbb{E}\Big[\mathcal{Q}^{(1)}(\delta_i, \delta_{-i}^*)\Big] - \sum_{j \in I \setminus \{i\}} \delta_j^* - \delta_i \tag{29}$$

By efficiency of δ^*, this is maximized by exerting effort $\delta_i = \delta_i^*$ as prescribed by the mechanism.

Now we consider individual rationality. We now omit δ^* from the $\mathcal{Q}$ notation as well since abiding by the mechanism is understood. The principal's expected utility in the truthful equilibrium equals: $\mathbb{E}[\mathcal{Q}^{(1)}] - \sum_{i \in I} \delta_i^*$, and this is non-negative by efficiency of the policy. Each agent i's expected utility in the truthful equilibrium equals his expected payment minus his expended effort, i.e.:

$$\mathbb{E}\Big[\mathcal{Q}^{(1)} - \sum_{j \in I \setminus \{i\}} \delta_j^*\Big] - \tag{30}$$

$$\mathbb{E}\Big[\mathcal{Q}_{-i}^{(1)} - \sum_{j \in I \setminus \{i\}} \delta_j^* - \delta_i^* + \mathbb{E}[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}]\Big] - \delta_i^* \tag{31}$$

$$= \mathbb{E}\Big[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}\Big] - \mathbb{E}\Big[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}\Big] = 0 \tag{32}$$

Finally, consider the expected aggregate payments received by the social planner. Noting that in the truthful equilibrium Eq. (27) minus Eq. (28) reduces to $\delta_i^* + Q^{(1)} - Q_{-i}^{(1)} - \mathbb{E}[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}]$, incorporating payments received from the principal in expectation this equals:

$$\sum_{i \in I} \delta_i^* - \sum_{i \in I} \Big(\delta_i^* + \mathbb{E}[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}] - \mathbb{E}[\mathcal{Q}^{(1)} - \mathcal{Q}_{-i}^{(1)}]\Big) = 0 \tag{33}$$

And so the budget is exactly balanced in expectation. □

Let us consider an example. Imagine there are three agents (with constant skill equal to 1), a principal with value $v = 8$, and uniformly distributed quality. We can use Theorem 1 to determine an optimal policy: since $\lfloor\sqrt{v}\rfloor^2 + \lfloor\sqrt{v}\rfloor = 6 < v = 8$, the optimal policy calls for $\lfloor\sqrt{v}\rfloor = 2$ agents—say agents 1 and 2—to expend effort $\delta_1 = \delta_2 = 1$ and the third to expend effort $\delta_3 = 0$. Imagine that the realized quality levels turn out to be $Q_1 = 3$ and $Q_2 = 5$. The mechanism requires that the principal pay: $\delta_1 + \delta_2 + \delta_3 = 2$. Noting that $\mathbb{E}[\mathcal{Q}^{(1)}(8, (1, 1, 0))] = 16/3$ and $\mathbb{E}[\mathcal{Q}^{(1)}(8, (1, 0))] = 4$, agent 1 is paid: $1 - (16/3 - 4) = -1/3$, i.e., he is charged 1/3. Agent 2, the maximum quality-producing agent, is paid: $1 + (5 - 3) - (16/3 - 4) = 5/3$. Finally, agent 3 is paid: $0 - (0 - 0) = 0$. Each agent's utility equals his payment minus effort (−4/3 for agent 1, 2/3 for agent 2, and 0 for agent 3); the principal's utility equals $5 - (1 + 1) = 3$; and revenue to the mechanism designer equals $2 + 1/3 - 5/3 = 2/3$. No agent could have gained in expectation from deviating from the mechanism's prescriptions, and although agent 1 was worse off for having participated, in expectation he was not so participation is rational given risk-neutrality.

Perhaps it could be considered a flaw of the mechanism that agents do not have *strict* incentive to participate: their expected utility from doing so is 0. First of all we note that the effort cost δ_i for an agent can be understood to incorporate *opportunity costs*, and can thus be construed as the difference in cost between the given effort level and the value of the agent's "outside option" (which will equal 0 if the agent has no other options).

But we can go further. Note that the principal *does* obtain positive surplus from the mechanism; in fact he is the only player (including the social planner) that does so in expectation. We can seek to distribute this more broadly. Let $\underline{v}$ be the minimum value the principal could *possibly* have, i.e., the greatest value that—independent of the principal's announcement—the mechanism designer *knows* is no greater than the true v (in the worst case $\underline{v} = 0$, but it may be greater). Let:

$$G = \mathbb{E}[\mathcal{Q}^{(1)}(\underline{v}, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) \tag{34}$$

We can amend the CSEC mechanism by charging the principal G and paying each agent G/n. Since G is completely independent of the principal's report, charging him thus will not change his incentives. Because quality is monotonically increasing in his value, G is a lower bound (guarantee) on the expected surplus the principal obtains in equilibrium under the CSEC mechanism, and so individual rationality will still hold in the amended mechanism. For any v in the space of possible values, the principal's expected utility will equal:

$$\mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*(v))] - \sum_{i \in I} \delta_i^*(v) - G \tag{35}$$

$$\geq \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) - G \tag{36}$$

$$\geq \mathbb{E}[\mathcal{Q}^{(1)}(\underline{v}, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) - G = 0, \tag{37}$$

where the first inequality holds by efficiency of δ^*. In expectation the mechanism remains perfectly budget-balanced, while now each agent may obtain positive utility and so will the principal (assuming $v \neq \underline{v}$).

In using this approach we are essentially adopting the technique of [Cavallo, 2006] in which revenue is "redistributed" to the agents in an effort to maintain wealth within the group rather than in the hands of the mechanism designer, without distorting incentives. Here we are seeking to redistribute surplus to the agents from the principal, but the technique is identical to that of [Cavallo, 2006] except instead of redistributing revenue we redistribute surplus.

3.2 Privately known skill

In the more general case where the principal has private value information, the agents have privately observed effort, *and* the agents have private skill information, the incentives problem is significantly more challenging. In fact, we can use the Myerson-Satterthwaite impossibility theorem [Myerson and Satterthwaite, 1983] to demonstrate the impossibility of achieving an efficient, incentive compatible, individually rational, and budget-balanced mechanism.

THEOREM 3. *There exists no mechanism that—for unrestricted quality distributions, private value for the principal, and private agent skill levels—is efficient, incentive compatible, individually rational, and budget-balanced.*

PROOF. We can prove the result via a "reduction" to efficient crowdsourcing from bilateral trade, for which we know by [Myerson and Satterthwaite, 1983] that there is no efficient, incentive compatible, individually rational, and budget-balanced mechanism. Assume for contradiction that the theorem fails. Then for any bilateral trade setting where the seller has value $\theta_s \in [0, 1]$ and the buyer has value $\theta_b \in [0, 1]$ consider the following crowdsourcing problem: the principal has value $v = \theta_b$, there is a single agent i who has skill $s_i = 1 - \theta_s$, and quality is (deterministically) distributed as follows:

$$Q_i(v, s_i, \delta_i) = \begin{cases} v & \text{if } \delta_i \geq 1 - s_i \\ 0 & \text{otherwise} \end{cases} \tag{38}$$

Assume that if the social planner chooses a policy in which quality v would be realized with certainty (given announced skill), the agent can be *compelled* to exert effort until quality v results (despite the planner not being able to observe the actual effort level expended). This only makes solving the crowdsourcing problem *easier*, and so a solution to the full crowdsourcing problem implies a solution to this variant.

Note that in any efficient policy i exerts effort either 0 or $1 - s_i = 1 - (1 - \theta_s) = \theta_s$. If the policy calls for production then the principal gains utility θ_b and the agent loses utility θ_s. Therefore in the efficient policy production occurs if and only if $\theta_b \geq \theta_s$. Moreover note that the strategic situations of the principal and agent in this crowdsourcing problem are *identical* to that of the buyer and seller in the bilateral trade problem, in the sense that the expected utility of each, given any strategy they play, for any strategy played by the other, is identical in either problem. So a payment scheme that is effective for the crowdsourcing problem would constitute a solution for the bilateral trade problem, and this would contradict the Myerson-Satterthwaite theorem. □

This negative result notwithstanding, there is a way forward with a weaker individual rationality concept. Though skill is private, it is not implausible to imagine that agents only learn their skill levels *after* the nature of the project is announced, in which case *ex ante* individual rationality would be sufficient to achieve participation if we force agents to make participation decisions before announcing the nature of the task. And [Cavallo, 2011] provides an efficient mechanism—which here we will call the *ex-ante-commitment mechanism*—that is incentive compatible for arbitrary private values settings and achieves individual rationality ex post of type realization for one player, while achieving individual rationality (and budget-balance) *ex ante* of type realizations for all others. This fits our setting perfectly, since the principal will know his value from the outset, but commitment by the agents may potentially be arranged to occur prior to realization of their types. Our setting has private *action* (effort) as well as private information, but the incentives provided by the mechanism extend.

The ex-ante-commitment mechanism makes payments based on expectations with respect to a prior distribution over agent types, assumed to be shared (a priori) by the mechanism designer and all agents. We use notation $\mathbb{E}_{\tilde{s}}[\cdot]$ to denote the expected value of a quantity with respect to the prior distribution over agents' skill levels, with $\tilde{s}$ a random variable representing the skill vector. Again letting $\underline{v}$ denote the lowest value in the principal's value space, let:

$$G(s) = \mathbb{E}[\mathcal{Q}^{(1)}(\underline{v}, s, \delta^*(\underline{v}, s))] - \sum_{i \in I} \delta_i^*(\underline{v}, s) \tag{39}$$

A derivative of the ex-ante-commitment mechanism for our setting takes the following form:

DEFINITION 3. (PRIVATE SKILL EFFICIENT CROWDSOURCING (PSEC) MECHANISM) *The principal reports v, each agent $i \in I$ reports s_i, and then the efficient effort levels $\delta_1^*, \ldots, \delta_n^*$ are computed. Each $i \in I$ is instructed to expend effort δ_i^*, and goods are produced with quality levels $Q(v) = (Q_1(v), \ldots, Q_n(v))$. The principal is charged:*

$$\sum_{i \in I} \delta_i^* + G(s), \tag{40}$$

and each agent $i \in I$ is paid:

$$Q^{(1)}(v) - \sum_{j \in I \setminus \{i\}} \delta_j^* - \tag{41}$$

$$\mathbb{E}_{\tilde{s}}\left[\mathcal{Q}^{(1)}(v, \tilde{s}, \delta^*(v, \tilde{s})) - \sum_{j \in I} \delta_j^*(v, \tilde{s}) - \frac{1}{n} G(\tilde{s})\right] \tag{42}$$

The following theorem is essentially a consequence of the main theorem of [Cavallo, 2011], although in that work the strategic element is solely related to private information, while for us there is also a hidden action (effort). But the basic logic extends: Groves mechanisms achieve straightforward *behavior*, which typically consists of truthful reporting but can also encompass other actions, e.g., production effort.

THEOREM 4. *The PSEC mechanism is incentive compatible, individually rational for the principal, individually rational for each agent ex ante of skill level realizations, and budget-balanced in expectation ex ante of skill realizations.*

We demonstrate the workings of the mechanism on the following uniformly distributed quality example: there are two agents, where the prior distribution over each's skill level assigns probability 0.5 to skill 0 and probability 0.5 to skill 1; assume the realized skills are $s_1 = 0$ and $s_2 = 1$. Assume the principal's value space is $[5, 50]$ (so $\underline{v} = 5$), and that the principal's actual value is $v = 10$. In the optimal policy agent 2 participates with full effort and agent 1 does not participate. $G(s) = 2.5-1 = 1.5$ and $\mathbb{E}_{\tilde{s}}\big[\mathcal{Q}^{(1)}(v, \tilde{s}, \delta^*(v, \tilde{s})) - \sum_{i\in I} \delta_i^*(v, \tilde{s}) - \frac{1}{n} G(\tilde{s})\big] \approx 2.6$. Imagine that quality 6 is realized. Then the principal pays $1 + 1.5 = 2.5$, obtaining a net utility of 3.5. Agent 1 is paid $6 - 1 - 2.6 = 2.4$, for a net utility of 2.4. Agent 2 is paid $6 - 0 - 2.6 = 3.4$, for a net utility of 2.4. Revenue equals $2.5 - 2.4 - 3.4 = -3.3$.

While the social planner or some agent may end up worse off ex post, as in the example, in expectation given the distribution over types each will gain from participation.

4. CONCLUSION

While most prior work seeks to maximize the utility of the principal alone, in this paper we pursue a crowdsourcing scheme that is *socially* optimal, maximizing the aggregate efficiency to all stakeholders in the system; we believe this holds the potential to bring significant added value to crowdsourcing marketplaces. Our findings and proposals may be of interest from both a theoretical and practical standpoint. From a theoretical perspective, we provide an efficient, individually rational, budget-balanced mechanism in the constant skill case; while we show that this is not possible in the general case, there we describe a mechanism that is efficient, budget-balanced, IR for the principal, and *ex-ante* IR for the producers. The results inform a designer of a crowdsourcing contest how to compute the optimal number of participants, given the principal's value and the agents' distribution over quality, and also tell the designer how to award the payments or prizes. An interesting facet of the mechanism we propose is that if the optimal number of participants is k, then the mechanism should award k payments (or more in the private skill case). This is in contrast with the winner-take-all schemes currently prevalent in crowdsourcing, where the participant who submits the highest quality good is the sole recipient of a lump sum prize.

5. ACKNOWLEDGMENTS

We thank John Langford for a helpful discussion about computing efficient policies.

6. REFERENCES

[Archak and Sundararajan, 2009] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the Thirtieth International Conference on Information Systems (ICIS-09)*, 2009.

[Baye *et al.*, 1996] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8:291–305, 1996.

[Cavallo and Parkes, 2008] Ruggiero Cavallo and David C. Parkes. Efficient metadeliberation auctions. In *Proceedings of the 26th Annual Conference on Artificial Intelligence (AAAI-08)*, pages 50–56, 2008.

[Cavallo, 2006] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-06)*, pages 882–889, 2006.

[Cavallo, 2011] Ruggiero Cavallo. Efficient allocation through delayed information revelation. working paper, 2011.

[Chawla *et al.*, 2011] Shuchi Chawla, Jason Hartline, and Balusubramanian Sivan. Optimal crowdsourcing contests. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA-12)*, pages 856–868, 2011.

[DiPalantino and Vojnovic, 2009] Dominic DiPalantino and Milan Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC-09), Stanford, CA*, pages 119–128, 2009.

[Fullerton and McAfee, 1999] Richard Fullerton and R. Preston McAfee. Auctioning entry into tournaments. *Journal of Political Economy*, 107:573 – 605, 1999.

[Ghosh and Hummel, 2011] Arpita Ghosh and Patrick Hummel. A game-theoretic analysis of rank order mechanisms for user generated content. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC-11), San Jose, CA, USA*, pages 189–198, 2011.

[Ghosh and McAfee, 2011] Arpita Ghosh and Preston McAfee. Incentivizing high quality user generated content. In *Proc. of the 20th International Conference on World Wide Web (WWW-11), Hyderabad, India*, pages 137–146, 2011.

[Green and Stokey, 1983] J. Green and N. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–364, 1983.

[Hilman and Riley, 1989] A. Hilman and J. Riley. Politically contestable rents and transfers. *Economics and Politics*, 1:17–39, 1989.

[Jain and Parkes, 2008] Shaili Jain and David C. Parkes. A game-theoretic analysis of games with a purpose. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE-08)*, pages 342–350, 2008.

[Jain *et al.*, 2009] Shaili Jain, Yiling Chen, and David C. Parkes. Designing incentives for online question and answer forums. In *Proc. of the 10th ACM Conference on Electronic Commerce (EC-09), Stanford, CA*, pages 129–138, 2009.

[Krishna and Morgan, 1997] V. Krishna and J. Morgan. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory*, 72:343 – 362, 1997.

[Lazear and Rosen, 1981] E. Lazear and S. Rosen. Rank order tournaments as optimum labor contracts. *Journal of Political Economy*, 89:841–864, 1981.

[Minor, 2011] Dylan Minor. Increasing efforts through incentives in contests. Manuscript, 2011.

[Moldovanu and Sela, 2001] Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91:542–558, June 2001.

[Moldovanu and Sela, 2006] Benny Moldovanu and Aner Sela. Contest architecture. *Journal of Economic Theory*, 126(1):70–96, January 2006.

[Moulin, 1986] Herve Moulin. *Game Theory for the Social Sciences*. New York University Press, 1986.

[Myerson and Satterthwaite, 1983] Roger Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[Nalebuff and Stiglitz, 1983] B. Nalebuff and J. Stiglitz. Prizes and incentives: Towards a general theory of compensation and competition. *Bell Journal of Economics*, 14:21–43, 1983.

[Shneidman and Parkes, 2004] Jeffrey Shneidman and David C. Parkes. Specification faithfulness in networks with rational nodes. In *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing*, pages 88–97, 2004.

[Tullock, 1980] Gordon Tullock. Efficient rent-seeking. In James M. Buchanan, Robert D. Tollison, and Gordon Tullock, editors, *Towards a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, 1980.

[Weber, 1985] Robert Weber. Auctions and competitive bidding. In H.P. Young, editor, *Fair Allocation*, pages 143–170. American Mathematical Society, 1985.