AAAI Technical Report SS-12-06
Wisdom of the Crowd

Efficient Crowdsourcing With Stochastic Production

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 timeframe 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 only 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.

Introduction

An increasingly common procurement model has multiple agents simultaneously produce versions of a desired good at the behest of a principal who seeks the highest-quality version. The popularity of this paradigm 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 subfield of microeconomics known as *contest design*. The model of a contest matches the standard approach to crowdsourcing: many agents simultaneously exert effort towards individual submissions 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 obtained 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-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 doing other than what the efficient policy prescribes. We address both components, focusing in this

[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" to the seller.

[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.

short version of the paper on the case where skill is constant across agents, for which we derive an efficient, individually rational, and budget-balanced solution.

Related Work

There has recently been work explicitly addressing the theory of crowdsourcing, in a model where agents have private information about a skill parameter 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 equals the *maximum* value over the produced goods, the crowdsourcing paradigm itself is not well-motivated from an efficiency standpoint. 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, Hartline, and Sivan, 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, Kovenock, and de Vries, 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. The authors find that the optimal contest should include only the two most skilled agents and propose using an all-pay auction to "select" these two contestants before the actual contest.

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, Chen, and Parkes, 2009; Jain and Parkes, 2009; Ghosh and McAfee, 2011; Ghosh and Hummel, 2011].

Preliminaries

In the crowdsourcing paradigm multiple units of a good are produced and submitted simultaneously to a principal. There is a set of agents $I = \{1, \ldots, n\}$ 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.:

$$\forall s_i,\ \forall 0 \le \delta_i < \delta_i' \le 1,\ \forall x \in [0, v],\ F^v_{s_i,\delta_i}(x) \ge F^v_{s_i,\delta_i'}(x), \tag{1}$$

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

$$\forall \delta_i,\ \forall 0 \le s_i < s_i' \le 1,\ \forall x \in [0, v],\ F^v_{s_i,\delta_i}(x) \ge F^v_{s_i',\delta_i}(x) \tag{2}$$

Because agents are self-interested there is a problem of incentives: v and s_i are private information, and expended effort is privately observed. We adopt a quasilinear 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

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

the good procured minus any payments he must make, and each agent's utility equals any payment he receives minus the effort he expends.

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 the 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:

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

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 write δ_i^* as shorthand for $\delta_i^*(v, s)$. Given our quasi-linear 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*.

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. Then in the next section we address 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 use of the following lemma, which demonstrates sufficient conditions under which *extreme-effort* policies are optimal.[5]

Lemma 1. *For arbitrary $i \in I$ with arbitrary skill s_i, for arbitrary skill levels of the other agents and value v for the principal, fix an arbitrary effort policy for agents other than i. Amongst effort levels for i in 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:* $\forall \beta \in [0, v]$, $\forall \epsilon \in [a, b)$,

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

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

Corollary 1. *For environments where the quality distribution functions satisfy the relationship of Eqs. (4) and (5) in Lemma 1 over the full range of effort levels $[0, 1]$, 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).

[5]All proofs are omitted due to space; we refer the interested reader to the full version of the paper.

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:

$$m^* = \arg\max_{m \in \{0,\ldots,n\}} \Big[m \int_0^v F^v_{\hat{s},1}(x)^{m-1} f^v_{\hat{s},1}(x) x\, dx - m \Big] \qquad (6)$$

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.

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, we take $F^v_{s_i,\delta_i}$ for each $i \in I$ to be the uniform distribution over quality levels between 0 and $\delta_i s_i v$, i.e.,

$$f^v_{s_i,\delta_i}(x) = \begin{cases} \frac{1}{\delta_i s_i v} & \text{if } x \in [0,\ \delta_i s_i v] \\ 0 & \text{otherwise} \end{cases} \qquad (7)$$

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

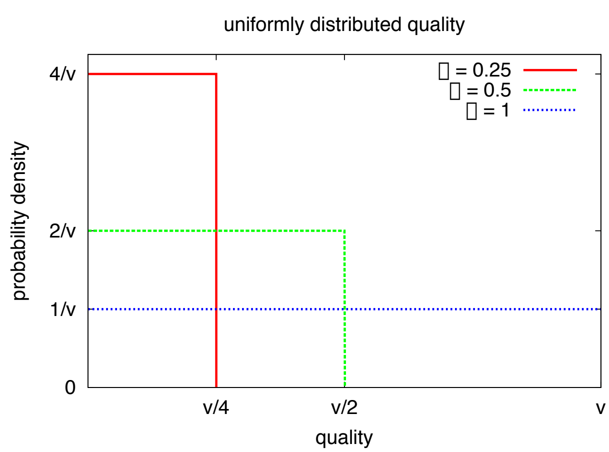

Figure 1: Probability density over quality levels for varying degrees of effort, for an agent with skill level 1 in the uniformly distributed quality case.

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 $\delta_i = 1$.*

Thus computing an efficient policy can be reduced to choosing the optimal set of agents to participate, who will participate with maximum effort. In the *constant skill* case where all agents have the same skill level (normalized to 1, below) this reduces to choosing the optimal *number* of agents to participate.

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:*

$$m^* = \begin{cases} \lfloor\sqrt{v}\rfloor - 1 & \text{if } \lfloor\sqrt{v}\rfloor^2 + \lfloor\sqrt{v}\rfloor > v \\ \lfloor\sqrt{v}\rfloor & \text{otherwise} \end{cases} \quad (8)$$

Figure 2 depicts how the optimal number of participating agents relates to the principal's value.

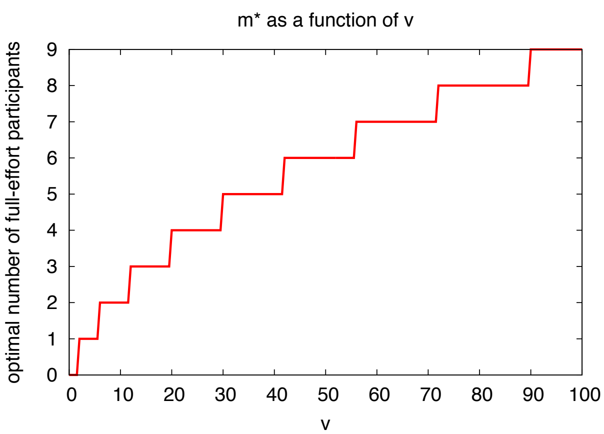

Figure 2: Optimal number of agents that should produce a good (with full effort) in the uniform quality case, as a function of the principal's value v.

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 $\mu = \delta_i s_i$ and scale parameter $\sigma = v/8$.

For arbitrary v, β, and δ_i 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 value, given any β, of effort for arbitrarily fine discretizations of the effort space $\delta_i \in [0, 1]$. Then, we 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 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 indicate that for all values of v above a very low threshold (2.8), when skill equals 1 there is no possibility that anything other than extreme-effort could be optimal. Checking for lower skill levels, we notice that the small range of values for which extreme-effort is not demonstrably optimal shrinks further.

We emphasize that we are only checking certain *sufficient conditions*: extreme-effort may well be optimal even when the conditions fail. Then, given that an extreme-effort policy is efficient we can easily compute the optimal number m^* of full-effort participants according to Eq. (6). Figure 3 depicts efficiency as a function of number of participants, for various v.

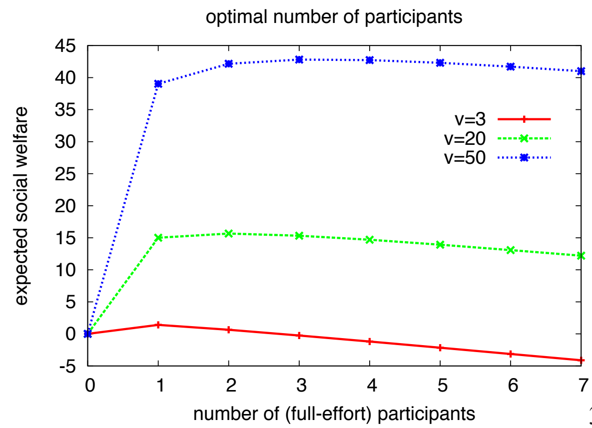

3.1in

Figure 3: Expected social welfare as a function of the number of agents participating with full-effort (where others exert no effort), for various instantiations of the principal's value v in the normally distributed quality case. As v rises the marginal value of extra quality increases, and accordingly the efficient level of participation rises.

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: our goal will be to realign every player's incentives with the maximization of *social* welfare, so that each player maximizes his own utility exactly when the sum of all players' utilities is maximized. We seek an equilibrium solution 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 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 if skill is variable and private knowledge.[6] 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 other constraints: the mechanism should be weakly *budget-balanced*, in expectation never paying out more than it takes in, since otherwise an external subsidy would be required for its implementation. The mechanism should also be *individually rational*, meaning no agent should expect to be worse off in equilibrium from truthfully participating.

In this short version of the paper for the incentives analysis we present results only for a context of *constant skill*, where two agents exerting the same effort produce quality according to the same distribution. In the full paper we develop a theory of incentives for the general private skill model. 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, and also individually rational and budget balanced. It defines payments that, in some cases, depend on a priori *expected* quality for a given effort level.

We use 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, and $Q_i(v)$ to denote an actual quality level *realized* by i (Q_i will be 0 by default if $\delta_i = 0$). 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 proposed below can be slightly (and harmlessly) modified to have the principal report the quality level of each produced good. $Q(v)$ 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.

Mechanism 1. *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:*

$$\sum_{i \in I} \delta_i^*, \tag{9}$$

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

$$\delta_h^* + Q_h(v) - Q^{(2)}(v) - \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*) - \mathcal{Q}_{-h}^{(1)}(v, \delta_{-h}^*)], \tag{10}$$

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

$$\delta_i^* - \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*) - \mathcal{Q}_{-i}^{(1)}(v, \delta_{-i}^*)] \tag{11}$$

[6] 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) obtains a "bonus" payment $(Q_h(v) - Q^{(2)}(v))$.

Theorem 2. *Mechanism 1 is efficient, incentive compatible, individually rational, and budget balanced in expectation for constant skill settings.*

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 $\mathbb{E}[\mathcal{Q}^{(1)}(8, (1, 1, 0))] = \frac{16}{3}$ and $\mathbb{E}[\mathcal{Q}^{(1)}(8, (1, 0))] = 4$, agent 1 is paid:

$$1 - \left(\frac{16}{3} - 4\right) = -\frac{1}{3}, \tag{12}$$

i.e., he is charged $1/3$. Agent 2, the maximum quality-producing agent, is paid:

$$1 + (5 - 3) - \left(\frac{16}{3} - 4\right) = \frac{5}{3} \tag{13}$$

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.

One thing about the mechanism that could arguably be called a flaw is the fact that agents do not have a *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 in some cases we can go further. Note that the principal *does* obtain positive surplus from the mechanism; in fact he is the only player (including the mechanism designer) that does so in expectation. We can seek to distribute this more broadly. If the mechanism designer *knows*—independent of his announcement—that the principal's value is greater than or equal to some value $\underline{v}$, then let:

$$G = \mathbb{E}[\mathcal{Q}^{(1)}(\underline{v}, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) \tag{14}$$

We can amend Mechanism 1 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 the principal's value, G constitutes a lower bound (guarantee) on the expected surplus the principal obtains in equilibrium under Mechanism 1, and so individual rationality will still hold in the amended mechanism. $\forall v$ in the space of possible values, the principal's expected utility will equal:

$$\mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*(v))] - \sum_{i \in I} \delta_i^*(v) - G \tag{15}$$

$$\geq \mathbb{E}[\mathcal{Q}^{(1)}(v, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) - G \tag{16}$$

$$\geq \mathbb{E}[\mathcal{Q}^{(1)}(\underline{v}, \delta^*(\underline{v}))] - \sum_{i \in I} \delta_i^*(\underline{v}) - G = 0, \tag{17}$$

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 $v \neq \underline{v}$; if $v = \underline{v}$ his expected utility would be 0).

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.

Conclusion

In this paper we addressed *socially* optimal crowdsourcing, seeking to maximize the aggregate efficiency to all stakeholders in the system. As most prior work seeks to maximize the utility of the principal alone, we believe this holds the potential to bring significant added value to crowdsourcing marketplaces. Beyond their theoretical interest, we hope our findings make a practical contribution to the design of procurement marketplaces. Our results inform a designer of a crowdsourcing mechanism how to compute the optimal number of participants (workers) given the principal's value and the distribution over quality, and also how to award the payments or prizes. An interesting facet of the efficient payment scheme we propose is that if the optimal number of participants is k, then k payments are awarded, one to each. This is in contrast with the winner-take-all schemes currently characteristic of paid crowdsourcing, where the participant who submits the highest quality good is the sole recipient of a lump sum prize, the magnitude of which is independent of the value the principal obtains from the good.

References

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

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

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

Cavallo, R. 2006. 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)*, 882–889.

Chawla, S.; Hartline, J.; and Sivan, B. 2011. Optimal crowdsourcing contests. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA-12)*, 856–868.

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

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

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

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

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

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

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

Jain, S., and Parkes, D. C. 2009. The role of game theory in human computation systems. In *KDD-HCOMP*, 58–61.

Jain, S.; Chen, Y.; and Parkes, D. C. 2009. Designing incentives for online question and answer forums. In *Proc. of the 10th ACM Conference on Electronic Commerce (EC-09), Stanford, CA*, 129–138.

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

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

Minor, D. 2011. Increasing efforts through incentives in contests. Manuscript.

Moldovanu, B., and Sela, A. 2001. The optimal allocation of prizes in contests. *American Economic Review* 91:542–558.

Moldovanu, B., and Sela, A. 2006. Contest architecture. *Journal of Economic Theory* 126(1):70–96.

Moulin, H. 1986. *Game Theory for the Social Sciences*. New York University Press.

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

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

Tullock, G. 1980. Efficient rent-seeking. In Buchanan, J. M.; Tollison, R. D.; and Tullock, G., eds., *Towards a Theory of the Rent-Seeking Society*. Texas A&M University Press. 97–112.

Weber, R. 1985. Auctions and competitive bidding. In Young, H., ed., *Fair Allocation*. American Mathematical Society. 143–170.