Incentive Mechanism Design for Crowdsourcing: An All-Pay Auction Approach

TIE LUO, Institute for Infocomm Research, A*STAR
SAJAL K. DAS, Missouri University of Science and Technology
HWEE PINK TAN, Singapore Management University
LIRONG XIA, Rensselaer Polytechnic Institute

Crowdsourcing can be modeled as a principal-agent problem in which the principal (crowdsourcer) desires to solicit a maximal contribution from a group of agents (participants) while agents are only motivated to act according to their own respective advantages. To reconcile this tension, we propose an all-pay auction approach to incentivize agents to act in the principal's interest, i.e., maximizing profit, while allowing agents to reap strictly positive utility. Our rationale for advocating all-pay auctions is based on two merits that we identify, namely all-pay auctions (i) compress the common, two-stage "bid-contribute" crowdsourcing process into a single "bid-cum-contribute" stage, and (ii) eliminate the risk of task nonfulfillment. In our proposed approach, we enhance all-pay auctions with two additional features: an adaptive prize and a general crowdsourcing environment. The prize or reward adapts itself as per a function of the unknown winning agent's contribution, and the environment or setting generally accommodates incomplete and asymmetric information, risk-averse (and risk-neutral) agents, and a stochastic (and deterministic) population. We analytically derive this all-pay auction-based mechanism and extensively evaluate it in comparison to classic and optimized mechanisms. The results demonstrate that our proposed approach remarkably outperforms its counterparts in terms of the principal's profit, agent's utility, and social welfare.

CCS Concepts: • **Information systems** → **Incentive schemes**; • **Applied computing** → **Economics;** *Online auctions*; • **Theory of computation** → *Quality of equilibria*; • **Human-centered computing** → *Ubiquitous and mobile computing theory, concepts and paradigms;* Smartphones;

Additional Key Words and Phrases: Mobile crowd sensing, participatory sensing, incomplete information, risk aversion, Bayesian Nash equilibrium, shading effect

ACM Reference Format:
Tie Luo, Sajal K. Das, Hwee Pink Tan, and Lirong Xia. 2016. Incentive mechanism design for crowdsourcing: An all-pay auction approach. ACM Trans. Intell. Syst. Technol. 7, 3, Article 35 (February 2016), 26 pages.
DOI: http://dx.doi.org/10.1145/2837029

1. INTRODUCTION

Crowdsourcing as a new data-collection and problem-solving model has been widely used in place of the traditional outsourcing paradigm to tackle real-world challenges by leveraging the "power of crowds." In addition to first-generation crowdsourcing

This work was supported in part by A*STAR Singapore under SERC grant 1224104046, and in part by the U.S. National Science Foundation under grants CNS-1404677, IIS-1404673, CNS-1545037, and CNS-1545050.
Authors' addresses: T. Luo, Institute for Infocomm Research, A*STAR, Singapore 138632; email: luot@i2r.a-star.edu.sg; S. K. Das, Missouri University of Science and Technology, Rolla, MO 65409; email: sdas@mst.edu; H.-P. Tan, Singapore Management University, Singapore 178902; email: hptan@smu.edu.sg; L. Xia, Rensselaer Polytechnic Institute, Troy, NY 12180; email: xial@cs.rpi.edu.

DOI: http://dx.doi.org/10.1145/2837029

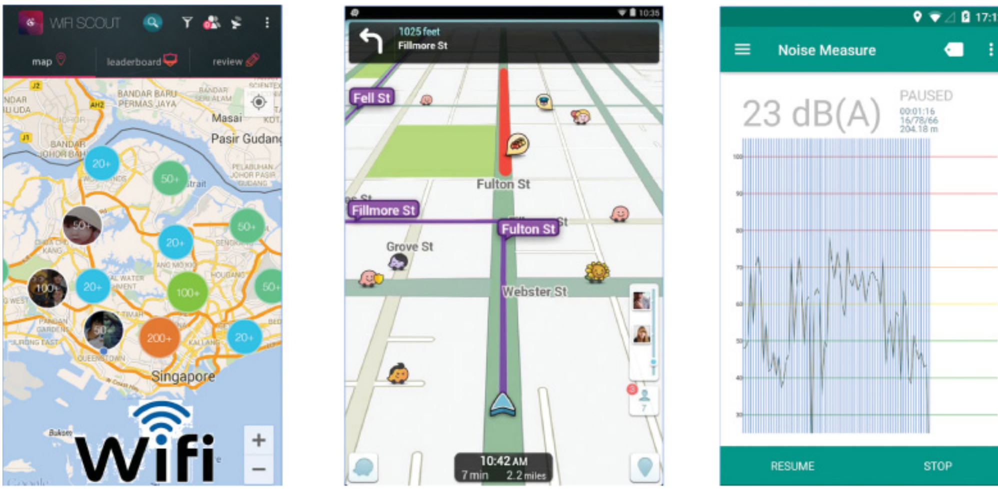

(a) Network quality monitoring (b) Transportation planning (c) Environmental noise measurement

Fig. 1. Example mobile crowdsourcing applications. Smartphone users contribute WiFi signal qualities [Wu and Luo 2014; WiFi-Scout 2014] (a), GPS readings [Waze 2009] (b), or noise measures [NoiseTube 2010] (c), through their phones to a cloud service for data analytics and visualization.

platforms such as Wikipedia, Amazon Mechanical Turk [2005], CrowdFlower [2007], and TaskRabbit [2008], the recent global proliferation of smartphones has spurred a new paradigm called *mobile crowdsensing* [Ganti et al. 2011; Zhang et al. 2014], otherwise known as participatory sensing [Guo et al. 2014, 2015], which ushered in a large number of projects, such as Waze [2009], NoiseTube [2010], ParkNet [Mathur et al. 2010], and WiFi-Scout [2014], some of which are illustrated in Figure 1. Although motivated by participatory sensing, the approach proposed in this article can be applied to crowdsourcing in general.

The viability of crowdsourcing critically depends on how well participants are motivated to contribute to such activities. There are a variety of possible motives, such as reputation [Zhang and van der Schaar 2012], social engagement [Kaufmann et al. 2011], and personal enjoyment [Han et al. 2011]; however, monetary reward is more generally applicable. This article pursues along this line and, in particular, uses *auctions* [Krishna 2009] as the framework to allocate such reward. Auction is an excellent tool for incentive mechanism design because it allows buyers and sellers to mutually agree on the price of some "commodity" (e.g., goods, time, or effort) that is otherwise hard to price by any third party. In fact, many prior studies [Lee and Hoh 2010; Yang et al. 2012; Koutsopoulos 2013; Feng et al. 2014; Zhao et al. 2014] were carried out in this regard as well. What is common in these prior studies is that they all adopt *winner-pay auctions*, where the bidders who bid (offer to contribute) higher than other bidders or a threshold win the auction (receive reward) and shall pay for their bids (by making corresponding contributions). Indeed, winner-pay auctions represent the majority of auctions, with first-price (Dutch) and second-price (English) auctions being their well-known classic examples.

In this article, by contrast, we take an *all-pay auction* approach to design incentive mechanisms for crowdsourcing. All-pay auctions are distinct from the winner-pay genre in that every bidder must pay for her bid regardless of whether she wins or loses the auction. This appears to be rather peculiar, but we identify two important merits that motivated us to explore this auction mechanism. First, all-pay auctions compress the common two-stage "bid-contribute" crowdsourcing process into a single "bid-cum-contribute" stage. In the two-stage pattern, which is commonly used in the literature (e.g., Lee and Hoh [2010], Yang et al. [2012], Koutsopoulos [2013], Feng et al. [2014], and Zhao et al. [2014]), all participants first need to enter a *bidding stage* to declare (i.e., bid) how much they would like to contribute or be paid. The highest or lowest bidders will be selected as winners and then enter the second stage, the *contribution stage*, where the actual crowdsourcing activity takes place (e.g., performing a sensing task). On the other hand, all-pay auctions do not require two stages but a single step: participants directly contribute to the crowdsourcing activity in which their respective contributions are their respective bids. Not only does this remarkably simplify the process of crowdsourcing, it also represents a very natural model for crowdsourcing because any contribution being crowdsourced is essentially human effort or some information, which is *irrevocable* once submitted (effort sunk or information disclosed).

The second advantage of all-pay auctions is that they eliminate the risk of *task nonfulfillment*. In winner-pay auction–based mechanisms such as Lee and Hoh [2010], Yang et al. [2012], Koutsopoulos [2013], Feng et al. [2014], and Zhao et al. [2014], winners are selected based on their *declared* bids (promised contributions or desired payments) to perform a task. However, the winners may not fulfill the task with the quality or quantity in accord with their bids, either intentionally or unintentionally. The consequence is typically that the crowdsourcer suffers from task delay, failure, or financial loss. In traditional auctions, such as those selling goods, one could stipulate regulations or rely on laws to enforce fulfilling the bids, but this is difficult to implement in crowdsourcing activities which are often ad hoc. On the other hand, all-pay auctions offer a good remedy by basing the winner selection on the *actual* rather than "promised" user contributions.

Therefore, in this article, we take all-pay auctions as our framework, and furthermore, we propose an approach with two enhancement features: (i) an *adaptive prize* that incentivizes participants (agents) to contribute insofar as the crowdsourcer (principal) reaps the maximum profit[1] and (ii) a *general crowdsourcing environment* that accommodates incomplete and asymmetric information, risk-averse agents, and stochastic population for the purpose of modeling more realistic scenarios.

Adaptive prize. We equip the principal with an adaptive prize as the reward, which scales as per a function of the unknown winner's contribution. This is in contrast to all conventional auctions (regardless of winner-pay or all-pay, including their variants) in which the item on sale or the reward is *fixed* ex-ante [Moldovanu and Sela 2001; Klemperer 2004; Krishna 2009; Archak and Sundararajan 2009; DiPalantino and Vojnovic 2009; Chawla et al. 2012].

Using an adaptive prize represents a new dimension of *optimal incentive mechanism design*. The optimality is traditionally defined as maximizing *revenue* [Myerson 1981] (e.g., the total contribution) for the principal, but in our case it is redefined as maximizing *profit* because the cost (prize) now is no longer a constant but a function. This also shows that our model subsumes the traditional goal as a special case.

In addition, we subject our mechanism to the constraint of strict individual rationality (SIR); in other words, the mechanism must be such that agents will receive strictly higher utility by participating than otherwise.

General crowdsourcing environment. This setting comprises three key elements.

(1) *Incomplete and asymmetric information*: As crowdsourcing typically involves a large and distributed group of people who are strangers to one another, we assume that each agent does not know any other agent's *type*, i.e., private information (e.g.,

[1]Henceforth, we use the terminology of *principal* and *agent* by framing crowdsourcing into a *principal-agent problem* on the basis that the crowdsourcer and the participants have asymmetric interests (maximizing their own, usually conflicting, utilities).

ability or cost), except her own (thus "incomplete"). On the other hand, the principal does not know any agent's type (thus "asymmetric"). The principal and agents only have distributional knowledge about agent types.

(2) *Risk-averse agents*: One common, yet often tacit, assumption in the mechanism design and game theory literature is that agents are *risk neutral*. Loosely speaking, it means that agents are indifferent between a sure-win \$50 reward and an uncertain \$100 on condition of tossing a fair coin. However, in reality, most people would prefer the guaranteed \$50 rather than the risky \$100, as they are *risk averse*, i.e., have stronger reluctance to lose than willingness to win. Therefore, we face our mechanism design problem with risk-averse agents, which also subsume risk-neutral agents as a special case.

However, a significant challenge is thus created: the most celebrated revenue equivalence theorem (RET)[2] [Myerson 1981], which is a powerful tool for analyzing many auction-related problems, breaks under the assumption of risk aversion. We tackle this challenge by employing *perturbation analysis* introduced by Fibich and Gavious [2003].

(3) *Stochastic population*: Another common assumption adopted by prior art is that the number of agents is known a priori. But in reality, this is often difficult to achieve due to the scale of crowdsourcing, and even if the number of *signed-up* agents can be retrieved from a database, it does not tell how many of them are actually contributing.

Therefore, we assume a stochastic population where the number of agents is uncertain. Notably, this endows our mechanism with another advantage: in other mechanisms of crowdsourcing, users have to declare their participation (like RSVP) in advance so that the number is known before crowdsourcing actually starts; but our assumption of a stochastic population gives the flexibility to permit *ad hoc user entry*, where an agent can enter a crowdsourcing activity to contribute at any time. Moreover, it subsumes deterministic population where the number is known and thus again lends us more generality.

1.1. Summary of Contributions

—We propose to use all-pay auctions as a framework to design incentive mechanisms for crowdsourcing by identifying two of its merits: (i) it compresses the two-stage bid-contribute crowdsourcing process into a single bid-cum-contribute stage, and (ii) it eliminates the risk of task nonfulfillment.

—We explore a new dimension of optimal incentive mechanism design by instrumenting an adaptive prize that scales as per a function of an unknown winner's contribution. We show that this new method results in superior performance.

—We cast our mechanism in a general crowdsourcing environment, which comprises incomplete and asymmetric information, risk averse agents, and stochastic population. This allows for a more realistic model and a wider range of applications.

—Despite the daunting all-pay nature (everyone has to pay regardless of winning or losing), our proposed mechanism satisfies SIR, which ensures that rational agents strictly have an incentive to participate. (In fact, our evaluations demonstrate that agent utilities are well bounded away from zero, which corresponds to nonparticipating.)

1.2. Article Organization

The rest of the article proceeds as follows. Section 2 reviews the literature, and Section 3 describes our model. Section 4 derives our optimal incentive mechanism, with its three

[2]RET states that any auction that satisfies a set of standard assumptions yields the same amount of expected revenue.

special cases presented in Section 5. Then, an alternative approach to optimal incentive mechanism design is explored in Section 6, which is then compared to our proposed approach as well as a standard counterpart mechanism in Section 7. Finally, Section 8 discusses some interesting findings, and Section 9 concludes this article.

2. RELATED WORK

The original idea appeared in Luo et al. [2014]. Since then, we have repositioned the problem in a broader scope, significantly augmented the technical content, and restructured and rewritten our original paper. Moreover, as this article reports a completely new and much more extensive set of performance evaluations, we also make the full package of our *Mathematica* source code available upon request.

In the following, we review both auction- and nonauction-related work, and for the former, we cover both standard and nonstandard settings.

2.1. Winner-Pay Auctions

Using a reverse auction to sell participatory sensing data was proposed in Lee and Hoh [2010]. It incentivizes participation with virtual credits using a heuristic to minimize cost and retain users. From a theoretical perspective, incentive mechanism design was investigated using auction theory in Yang et al. [2012], Koutsopoulos [2013], Feng et al. [2014], and Zhao et al. [2014]. Yang et al. [2012] studied a user-centric model and designed a truthful auction in which users bid their true costs for the crowdsourcer to select winners to perform a task set and pay them no lower than their bids. But as a trade-off, the payoff or profit of the crowdsourcer cannot be maximized due to NP-hardness. Koutsopoulos [2013] designed another truthful auction subject to a certain QoS requirement. But in addition to determining user payments like in Yang et al. [2012], the crowdsourcer also determines each user's participation level (e.g., data sampling rate) on users' behalf, which is more prone to task nonfulfillment. Feng et al. [2014] took location information into account when assigning mobile crowdsourcing tasks to users. Besides ensuring truthful cost bidding, similar to what was discussed earlier, they also proposed an approximate algorithm to find the near-optimal winning bids in polynomial time. Zhao et al. [2014] proposed two auction mechanisms to handle an online scenario where bidders only stay in the system for a (short) period of time and the crowdsourcer has to select a subset of bidders by a certain deadline to perform a task. All of the preceding auctions need an extra bidding stage.

2.2. All-Pay Auctions

In the relatively much smaller regime of all-pay auctions, Baye et al. [1996] were probably the first to analyze all-pay auctions with complete information, where all the bidders' types (valuations of the auctioned item) are common knowledge. On the other hand, all-pay auctions with incomplete information were studied by Moldovanu and Sela [2001] and Archak and Sundararajan [2009], but the aim was to examine whether allocating a single prize or dividing it into multiple prizes is optimal in terms of maximizing the total quality [Moldovanu and Sela 2001] or the highest k qualities [Archak and Sundararajan 2009]. Along a similar line, Chawla et al. [2012] showed that the highest bid is at least half the sum of all bids. Under the same (standard) model, Kaplan et al. [2002] and Cohen et al. [2008] explored optimal contest design with multiplicative or additive value functions when the revenue is defined as the total bids or the highest bid. DiPalantino and Vojnovic [2009] studied a crowdsourcing platform that offers multiple tasks each with a fixed reward. Using a standard all-pay auction model with a deterministic population of risk-neutral users, they found a logarithmic relationship between incentives and user participation levels.

All of these prior studies assume standard settings that are amicable to analysis but, on the other hand, leave the need for more realistic models.

2.3. Nonstandard Auction Settings

With nonstandard assumptions, analysis generally becomes more challenging or even intractable. For winner-pay auctions, McAfee and McMillan [1987] and Levin and Ozdenoren [2004] compared the revenue when the number of bidders is only known to the auctioneer but not to the bidders to the revenue when the number is known to all players. Both studies still assume risk-neutral bidders. Harstad et al. [1990] and Krishna [2009] characterized the equilibrium bidding strategies for an uncertain number of bidders, but a common and crucial vehicle used by both studies is RET because the bidders are assumed to be risk neutral. In the genre of all-pay auctions, Fibich et al. [2006] compared all-pay auctions with first-price auctions in terms of bids and revenue, with a known number of risk-averse players. The case of an unknown number of (risk-neutral) bidders was recently studied by Haviv and Milchtaich [2012], but to trade for tractability, it assumes identical bidders where all types are equal, which is clearly not the case in the context of crowdsourcing. In addition, all of the preceding assumes a constant auctioned item or reward.

2.4. Nonauction Incentives

Without using auctions, Luo and Tham took a resource-allocation approach to incentivize user participation by allocating service quota (where service is produced from user contributed data) to users based on their contribution levels and, particularly, their service demands [Luo and Tham 2012; Tham and Luo 2014]. Two incentive schemes, called *IDF* and *ITF*, were proposed to maximize fairness and social welfare, respectively. Restuccia and Das [2014] propose a trust-based incentive framework that takes data reliability into account to reward reliable users. The framework relies on the concept of mobile security agents and is shown to be resistant to a set of GPS-based attacks. The joint issue of incentive and trust is also addressed in Luo et al. [2014a], but it takes a socioeconomic approach to link participants, via a relationship called *endorsement*, into a *social network* overlaid by economic incentives. The endorsement relationship is based on three factors: social trust, financial return, and a new social notion called *nepotism*. Ghosh and McAfee [2011] studied the problem of incentivizing high-quality user-generated content (UGC), using a game-theoretical model with the solution concept of a free-entry Nash equilibrium. The authors designed an elimination mechanism that subjects contributions to ratings and eliminates contributions that do not receive good ratings. Xia et al. [2014] also studies a UGC problem but considers heterogeneous users in the sense that they can produce different best qualities (types).[3] The work focuses on proving the existence of pure Nash equilibria in various UGC mechanisms under different information settings. Quality of contribution is also a central concept in Tham and Luo [2013, 2015], but the authors expanded it to an aggregate notion referred to as quality of contributed service (QCS), which characterizes the overall service produced from individual user contributions. Under a market-based model, they show that QCS can converge to a good market equilibrium starting from purely random behaviors.

[3] A more general definition (and in standard auction terms) of heterogeneous users is that not only user types are different but also the beliefs (distributional knowledge; assuming incomplete information) about their respective types are different. This constitutes an *asymmetric auction*, and the corresponding mechanism design is addressed in Luo et al. [2014b, 2016].

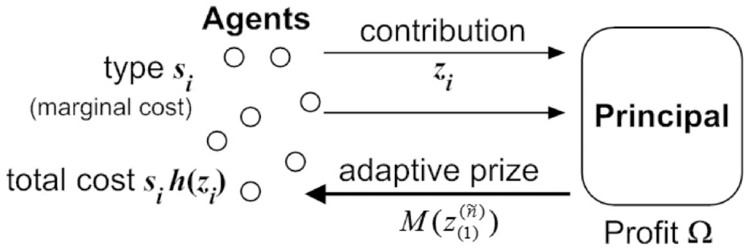

Fig. 2. Our all-pay auction–based crowdsourcing model. Each participant (agent) bids cum contributes z_i, incurring a cost $s_i h(z_i)$, and the agent who makes the highest contribution $z_{(1)}^{(\tilde{n})}$ will win a prize that is adaptive as per a function $M(\cdot)$. The crowdsourcer (principal) seeks to maximize profit, which is the total contribution $\sum_{i=1}^{\tilde{n}} z_i$ subtracted by the prize $M(z_{(1)}^{(\tilde{n})})$.

3. THE MODEL

A principal launches a crowdsourcing campaign to an uncertain population of agents to solicit their contributions to perform a task, such as collecting sensory data, producing creative work, or solving a problem. The campaign is flexible in the sense that agents can join and contribute at any time (due to the advantage of "ad hoc user entry"; cf. Section 1), and each agent can either submit a single contribution or multiple contributions (a specific campaign may choose, e.g., the best or the sum, as the valid submission).

The number of agents is unknown, denoted by $\tilde{n}$, which is a random variable following a probability mass function $\Pr(\tilde{n} = n) = p_n,\ n \geq 2$. An arbitrary agent i makes a (valid total) contribution of z_i, which can be a simple measure of quantity or quality alone, or a compound measure of both quantity and quality [Tham and Luo 2013, 2015]. The strategy of an agent i is to determine how much z_i to contribute.

To provide an incentive, the principal instruments a monetary prize to reward, at the end of the campaign or periodically, the agent who makes the highest contribution $\max_{i\in[1..\tilde{n}]} z_i$, which we denote by $z_{(1)}^{(\tilde{n})}$ following the notational convention from order statistics. The prize is *adaptive* as per a function $M(\cdot)$ of the (unknown) winning agent's contribution, i.e., $M(z_{(1)}^{(\tilde{n})})$. The function $M(\cdot)$ is known to all, whereas $z_{(1)}^{(\tilde{n})}$ is unknown.

Each agent is uniquely characterized by her type s_i, representing her marginal cost of participation; in other words, her total cost is $s_i h(z_i)$. Here, $h(\cdot)$ is a modulator function that allows for more generality, rather than being limited to the commonly used linear cost model $s_i \times z_i$. We assume that $h(0) = 0$ and $h(\cdot)$ is monotone increasing and continuously differentiable. Our model is depicted in Figure 2.

Information is incomplete, i.e., each agent only knows her own type s_i but not any other's type. On the other hand, it is common knowledge that all agent types are independently drawn from $[\underline{s}, \bar{s}]$ according to a continuously differentiable c.d.f. $F(s)$. Such a setting is referred to as an independent-private-value (IPV) model [Krishna 2009] with a common Bayesian belief $F(s)$.[4]

Agents are risk averse, characterized by a von Neumann-Morgenstern (vNM) utility function $u(\cdot)$, i.e., an agent derives a utility of $u(x)$ from a net gain of x (can be negative, which indicates a net loss). For example, $u(x) = x$ is the most common utility formulation that assumes risk-neutral agents. The function $u(\cdot)$ is twice differentiable and satisfies $u(0) = 0$, $u' > 0$, $u'' \leq 0$ (i.e., diminishing marginal gain or increasing marginal loss).

[4]If, otherwise, the belief about each agent is different, i.e., there is $F_i(s_i)$ for each i and $F_i(s_i) \neq F_j(s_j)$, this constitutes an asymmetric model and a solution can be found in Luo et al. [2014b, 2016].

Now we explain how our model can be realized in practice. The contribution z_i can be measured by the agent's personal device (e.g., smartphone, PC, wearable widget) or by the principal using a computing facility, or evaluated by an expert committee (e.g., in the case of creative work). The belief or c.d.f. $F(s)$ can be constructed empirically from historical contribution records or, in the absence of historical data, by adopting the uniform distribution, as similarly discussed in Koutsopoulos [2013].

Utility formulation. Each agent's utility is defined as

$$\tilde{\pi}_i(s_i, \mathbf{z}) := \begin{cases} u\left(M(z_i) - s_i h(z_i)\right), & \text{if } z_i > z_j, \forall j \neq i, \\ u(-s_i h(z_i)), & \text{otherwise,} \end{cases} \tag{1}$$

where $\mathbf{z} := \{z_1, z_2, \ldots z_{\tilde{n}}\}$. Her expected utility is

$$\pi_i(s_i, \mathbf{z}) := \mathop{\mathbb{E}}_{z_{-i}}[\tilde{\pi}_i], \tag{2}$$

where $z_{-i} := \mathbf{z} \setminus \{z_i\}$. In (1), the event of forming a tie ($z_i = z_j$) happens with probability zero because (i) $F(s)$ is continuously differentiable, and hence the p.d.f. is atomless, and (ii) bidding strategy z is a strictly monotone function of type s in the Bayesian game induced by such an IPV model [Krishna 2009].

The principal's utility or profit Ω is defined as the total contribution solicited from all agents less the adaptive prize, i.e.,

$$\Omega(\tilde{n}, \mathbf{z}) := \sum_{i=1}^{\tilde{n}} z_i - M\big(z_{(1)}^{(\tilde{n})}\big). \tag{3}$$

Problem statement. Our objective is to design an incentive mechanism that, at equilibrium (if exists), (a) maximizes the expected profit of the principal, i.e., it solves (note that throughout, we use "*" to indicate variables at equilibrium)

$$\mathop{\arg\max}_{M(\cdot)} \Omega^*,$$

where

$$\Omega^* := \mathop{\mathbb{E}}_{\tilde{n}}[\Omega(\tilde{n}, \mathbf{z}^*)], \tag{4}$$

(b) satisfies SIR for each agent, i.e., the expected utility of each agent is strictly positive if and only if she contributes;[5] formally,

$$\pi_i^* := \pi_i(s_i, \mathbf{z}^*) > 0 \quad \text{if and only if } z_i^* > 0, \tag{5}$$

where, as a common practice, the outside option (not participating) is assumed to receive zero utility. Compared to the canonical individual rationality (IR) which is defined as $\pi_i^* \geq 0$, SIR implies a stronger motivation to participants.

Remark. Incentive compatibility or truthfulness is often a mechanism design objective such as in Yang et al. [2012], Koutsopoulos [2013], Feng et al. [2014], and Zhao et al. [2014]. However, this is technically irrelevant to—or inherently resolved by—our all-pay auction model. This is because our model does not require agents to report their types but rather bases winner selection on agents' actual (and observable) contributions that already endogenize their respective (true) types and cannot be lied about. In fact, distinct from most prior work, this article represents an example of *indirect mechanism* design (as opposed to direct-revelation mechanisms).

[5]We use "expected" utility because of the incomplete-information setting. In standard game-theoretical terms, it corresponds to the "interim" stage (which is the most natural case) as opposed to the ex-ante and ex-post stages. Hence, our definition of SIR (5) can also be more precisely referred to as *interim* SIR.

4. OPTIMAL INCENTIVE MECHANISM DESIGN

We derive the optimal mechanism in two steps. In Section 4.1, we characterize the Bayesian Nash equilibrium (BNE) by assuming that the adaptive prize is *given*. In Section 4.2, we optimize the BNE by solving for the profit-maximizing prize function $M(\cdot)$, where we employ perturbation analysis.

4.1. Equilibrium Contribution Strategy

Our incomplete-information setting induces a Bayesian game and thereby requires an extended notion of Nash equilibrium as follows.

Definition 1 (*Bayesian Nash Equilibrium [Harsanyi 1968]*). A (pure strategy) Bayesian Nash equilibrium is a strategy profile $\mathbf{z}^* := (z_1^*, z_2^*, ...)$ that satisfies

$$\pi_i(s_i, \mathbf{z}^*; F_{-i}(s)) \geq \pi_i(s_i, z_i, z_{-i}^*; F_{-i}(s)), \quad \forall i, \forall z_i,$$

where $F_{-i}(s)$ is i's belief (a joint probability distribution) of all types other than s_i.

In other words, each agent plays a strategy z_i^* that maximizes her expected utility π_i given her belief about other agents' types and that other agents also play their respective equilibrium strategies z_{-i}^*.

LEMMA 1 (EXISTENCE AND MONOTONICITY). *Our model admits a pure strategy Bayesian Nash equilibrium in which each agent's strategy is a monotone decreasing function of her type.*

PROOF. Due to space constraints, most proofs in this article are contained in the Appendix, which is available in the ACM Digital Library (http://dl.acm.org). □

In our case, the BNE is symmetric, i.e., all agents adopt the same strategy (function) because of the homogeneous belief $F(\cdot)$ (see Luo et al. [2014b, 2016] for the heterogeneous case). To solve it, we first need to express the expected utility of each agent in terms of her winning probability. However, the stochastic population size presents a challenge. One possible solution is to use *contingent bidding* proposed in Harstad et al. [1990], where each bidder submits a *list* of bids in the form of "I bid z_1 if there are n_1 bidders, z_2 if $n_2, \ldots$" Then, after collecting all of the lists, the auctioneer knows the exact number of bidders, say n_k, and thus can use the corresponding z_k's to determine the winner. However, this contingent bidding method does not apply to our all-pay auction case, because each bid represents *irrevocable* contribution of effort or information. Therefore, we use a conditional winning probability:

$$\Pr\left(z_i^*(s_i) > z_j^*(s_j) | \tilde{n} = n\right) = (1 - F(s_i))^{n-1}, n = 2, 3, \ldots, \forall i \neq j,$$

where $1 - F(s_i)$ is due to the decreasing monotonicity of the equilibrium (Lemma 1). By taking the expected value, we obtain the winning probability for an agent of arbitrary type s, as

$$P(s) = \sum_{n=2}^{\infty} p_n (1 - F(s))^{n-1}. \tag{6}$$

This may lure one to hypothesize that the equilibrium contribution strategy could simply be a weighted sum $z^*(s) = \sum_n p_n z_{dp}^*(s|n)$, where $z_{dp}^*(s|n)$ is the strategy under a deterministic population size n. This is not true, because z^* is not an affine optimizer of the utility π^* (optimizer in the sense that z^* is the best-response strategy to maximize π^*), which is analogous to Jensen's inequality. In fact, such a weighted sum is not the case even when bidders are risk neutral [Krishna 2009].

Given $P(s)$, we can now express the expected utility of an agent of type s, based on (1), as

$$\begin{aligned} \pi(s, \mathbf{z}^*) &= u(M(z^*) - h(z^*)s)P(s) + u(-h(z^*)s)(1 - P(s)) \\ &\equiv P(s)[u(\alpha^*) - u(-\beta^*)] + u(-\beta^*), \qquad (7) \\ \text{where} \qquad \alpha^* &:= M(z^*) - h(z^*)s, \qquad \beta^* := h(z^*)s. \end{aligned}$$

LEMMA 2 (EQUILIBRIUM STRATEGY). *In an all-pay auction with incomplete information and a stochastic population of risk-averse agents, given an adaptive prize $M(z)$, the equilibrium strategy $z^*(s)$ is determined by*

$$\int_s^{\bar{s}} [P(s_1)(u'(\alpha^*) - u'(-\beta^*)) + u'(-\beta^*)]h(z^*(s_1))\,\mathrm{d}s_1 = P(s)[u(\alpha^*) - u(-\beta^*)] + u(-\beta^*). \quad (8)$$

Remark. In general, there is no closed-form expression of $z^*(s)$ when agents are risk averse or the prize is not constant. However, when agents are risk neutral *and* the prize is constant, we can obtain an explicit solution using Lemma 2, as will be demonstrated later by Corollary 2 (26).

4.2. Optimal Mechanism

We first derive the profit at equilibrium for a given prize and then maximize the profit by finding the optimal adaptive prize function using perturbation analysis.

LEMMA 3 (PROFIT AT EQUILIBRIUM). *The profit of the principal at equilibrium is given by*

$$\Omega^* = \sum_n np_n \int_{\underline{s}}^{\bar{s}} [z^*(s) - M(z^*)(1 - F)^{n-1}]\,\mathrm{d}F. \quad (9)$$

Next, to maximize the profit Ω^*, we need an explicit expression of $M(z^*)$. However, directly solving it from (8), where $M(z^*)$ is embedded in α^*, will be of no avail because of the reason remarked below Lemma 2. To tackle this challenge, we employ the *perturbation analysis* introduced by Fibich and Gavious [2003].

Perturbation analysis. This method applies to the case of *weak risk aversion*. Intuitively speaking, although a weakly risk-averse agent still prefers \$50 for certain to \$100 with half chance, she will, however, take that "risky" \$100 rather than be guaranteed but a mere \$20 or \$30. Weak risk aversion is most common in reality and is not a restrictive assumption.

Formally, we say an agent is *weakly risk averse* if, for any net payment x (possibly negative) she receives, her vNM utility function $u(x)$ satisfies $u(x) \approx x$. In general, the utility function can be written (e.g., using Taylor series expansion) as

$$u(x) = x + \epsilon u_1(x) + O(\epsilon^2), \quad 0 < \epsilon \ll 1,$$

where ϵ is referred to as the risk aversion parameter and $\epsilon \ll 1$ indicates weak risk aversion. The function u_1 satisfies $u_1(0) = 0, u_1'' \le 0$ and $u_1' > -\frac{1}{\epsilon}$ (u_1' can be either positive or negative) for $u' > 0$ and $u'' \le 0$. For example, $u(x) = [1 - e^{-\epsilon x}]/\epsilon$ is a constant absolute risk averse (CARA) utility function, and $u(x) = x^{1-\epsilon}$ is a constant relative risk averse (CRRA) utility function.[6]

[6]CARA means that the *Arrow-Pratt measure* of absolute risk aversion, which is defined as $-u''(x)/u'(x)$, is a constant. CRRA means that the *Arrow-Pratt-De Finetti measure* of relative risk aversion, which is defined as $-xu''(x)/u'(x)$, is a constant.

Now, revisiting Lemma 2, we write the following using the perturbation method:

$$\begin{cases} u(\alpha^*) & = \alpha^* + \epsilon u_1(\alpha^*) + O(\epsilon^2) = \alpha^* + \epsilon u_1(\alpha^*_{rn}) + O(\epsilon^2), \\ u(-\beta^*) & = -\beta^* + \epsilon u_1(-\beta^*) + O(\epsilon^2) = -\beta^* + \epsilon u_1(-\beta^*_{rn}) + O(\epsilon^2), \\ u'(\alpha^*) & = 1 + \epsilon u_1'(\alpha^*_{rn}) + O(\epsilon^2), \\ u'(-\beta^*) & = 1 + \epsilon u_1'(-\beta^*_{rn}) + O(\epsilon^2), \end{cases} \tag{10}$$

where the subscript rn indicates the risk-neutral case, i.e.,

$$\alpha^*_{rn} = M_{rn}(z^*_{rn}) - h(z^*_{rn})s, \qquad \beta^*_{rn} = h(z^*_{rn})s. \tag{11}$$

THEOREM 1 (OPTIMAL PRIZE AND MAXIMUM PROFIT). *In an all-pay auction with incomplete information and a stochastic population of weakly risk-averse agents, the optimal adaptive prize that maximizes profit for the principal is given by*

$$\mathring{M}(z) = \frac{1}{P(\mathring{s}(z))}\left[\mathring{s}(z)h(z) - \mathring{A}(\mathring{s}(z)) - \int_{\mathring{z}(\bar{s})}^{z} \mathring{B}(\mathring{s}(z_1))h(z_1)\,\mathrm{d}\mathring{s}(z_1)\right], \tag{12}$$

where $\mathring{s}(z)$ is the inverse function of $\mathring{z}(s)$, which is the equilibrium strategy induced by $\mathring{M}(z)$ and is given by

$$\mathring{z}(s) = (h')^{-1}\left(\frac{aF'(s)}{G'(s)s + G(s)\mathring{B}(s)}\right), \tag{13}$$

where $(h')^{-1}(\cdot)$ is the inverse function of $h'(\cdot)$, $h'' > 0$, and $a = \sum_n np_n$. The resultant maximum profit is given by

$$\mathring{\Omega} = \int_{\underline{s}}^{\bar{s}} [a\mathring{z}F' - h(\mathring{z})(sG' + \mathring{B}(s)G) + \mathring{A}(s)G']\,\mathrm{d}s. \tag{14}$$

In the preceding,

$$\begin{aligned} G(s) &= \int_{\underline{s}}^{s} F'(s_1)\frac{\sum_n np_n(1-F(s_1))^{n-1}}{\sum_n p_n(1-F(s_1))^{n-1}}\,\mathrm{d}s_1, \\ \mathring{A}(s) &= \epsilon P(s)[u_1(\mathring{\alpha}_{rn}) - u_1(-\mathring{\beta}_{rn})] + \epsilon u_1(-\mathring{\beta}_{rn}), \\ \mathring{B}(s) &= \epsilon P(s)[u_1'(\mathring{\alpha}_{rn}) - u_1'(-\mathring{\beta}_{rn})] + \epsilon u_1'(-\mathring{\beta}_{rn}) + 1. \end{aligned} \tag{15}$$

Remark on notation. Henceforth, we use the overhead circle "∘" to indicate the *optimal* (i.e., profit-maximizing) equilibrium quantities (prize, contribution strategy, profit, agent utility), whereas the superscript "*", as is commonly used, indicates general equilibrium quantities.

Remark on complexity. Despite looking complex, Theorem 1 captures a very general crowdsourcing environment. Furthermore, it can be easily simplified to special cases, which are presented in Section 5, and intuitive insights can also be drawn, which are presented in Sections 7 and 8.

Let us recall the second objective in our problem statement.

THEOREM 2 (STRICT INDIVIDUAL RATIONALITY (SIR)). *The mechanism specified by Theorem 1 satisfies strict individual rationality for both risk-neutral and weakly risk-averse agents. In other words, any such agent expects strictly positive utility at equilibrium.*

5. SPECIAL CASES

In this section, we show how our main result (Theorem 1), which accommodates a general setting, can be easily simplified to three special and representative cases:

risk-neutral agents (Section 5.1), deterministic population (Section 5.2), and the intersection of them (Section 5.3). These results will also be used by our evaluation in Section 7.

5.1. Risk-Neutral Agents

When agents are risk neutral, $\epsilon = 0$ and it immediately follows from (15) that

$$\mathring{A}(s) = 0, \ \ \mathring{B}(s) = 1. \tag{16}$$

COROLLARY 1 (RN). *In an all-pay auction with incomplete information and a stochastic population of risk-neutral agents, the optimal adaptive prize that maximizes profit for the principal is given by*

$$\mathring{M}_{rn}(z) \ = \ \frac{\mathring{s}_{rn}(z)h(z) - \int_{\mathring{z}_{rn}(\bar{s})}^{z} h(z_1)\,\mathrm{d}\mathring{s}_{rn}(z_1)}{P(\mathring{s}_{rn}(z))}, \tag{17}$$

where $\mathring{s}_{rn}(z)$ is the inverse function of $\mathring{z}_{rn}(s)$, which is the equilibrium strategy induced by $\mathring{M}_{rn}(z)$ and is given by

$$\mathring{z}_{rn}(s) = (h')^{-1}\left(\frac{aF'(s)}{G'(s)s + G(s)}\right), \tag{18}$$

where $(h')^{-1}(\cdot)$ is the inverse function of $h'(\cdot)$, $h'' > 0$, and $a = \sum_n np_n$. The resultant maximum profit is given by

$$\mathring{\Omega}_{rn} = \int_{\underline{s}}^{\bar{s}} [a\mathring{z}_{rn}F' - h(\mathring{z}_{rn})(G'(s)s + G(s))]\,\mathrm{d}s. \tag{19}$$

PROOF. Substituting (16) into Theorem 1 yields the result. □

5.2. Deterministic Population

When the number of agents is known as n, we have the following simplified expressions:

$$\begin{aligned} a &= n, & P(s) &= (1 - F(s))^{n-1}, \\ G(s) &= nF(s), & G'(s) &= nF'(s). \end{aligned} \tag{20}$$

Thus, we obtain the result for this DP case similarly as Section 5.1. We omit the details for brevity.

5.3. Risk Neutral and Deterministic Population

This is the most common setting in the literature.

COROLLARY 2 (RN-DP). *In an all-pay auction with incomplete information and a deterministic population of n risk-neutral agents, the optimal adaptive prize that maximizes profit for the principal is given by*

$$\mathring{M}_{rn,dp}(z) = \frac{\mathring{s}_{rn,dp}(z)h(z) - \int_{\mathring{z}_{rn,dp}(\bar{s})}^{z} h(z_1)\,\mathrm{d}\mathring{s}_{rn,dp}(z_1)}{\left(1 - F(\mathring{s}_{rn,dp}(z)\right)^{n-1}}, \tag{21}$$

where $\mathring{s}_{rn,dp}(z)$ is the inverse function of $\mathring{z}_{rn,dp}(s)$, which is the equilibrium strategy induced by $\mathring{M}_{rn,dp}(z)$ and is given by

$$\mathring{z}_{rn,dp}(s) = (h')^{-1}\left(\frac{F'}{sF' + F}\right), \tag{22}$$

where $(h')^{-1}(\cdot)$ *is the inverse function of* $h'(\cdot)$*, and* $h'' > 0$*. The resultant maximum profit is given by*

$$\mathring{\Omega}_{rn,dp} = n \int_{\underline{s}}^{\bar{s}} \left[\mathring{z}_{rn,dp} - h(\mathring{z}_{rn,dp}) \left(s + \frac{F}{F'} \right) \right] \mathrm{d}F. \tag{23}$$

Proof. Substituting (20) into Corollary 1 proves the result. □

6. ALTERNATIVE OPTIMAL INCENTIVE MECHANISM DESIGN

Our proposed approach represents a new way of optimal incentive mechanism design. Alternatively, we could explore another, and perhaps more intuitive, way by optimizing the (constant) auction prize value by exploiting all information available to the principal.

This section pursues this idea under two cases: (i) the *general case*, a stochastic population of risk-averse agents, which corresponds to our environment, and (ii) the *special case*, a deterministic population of risk-neutral agents.

6.1. General Case

To state the problem, a principal needs to determine, over all possible prizes M_0 to offer, an optimal one to maximize its profit. The profit can be written using (A.3) or (9) as

$$\Omega^*_{con} = \sum_n np_n \int_{\underline{s}}^{\bar{s}} z^*_{con}\, \mathrm{d}F - M_0, \tag{24}$$

where the equilibrium strategy z^*_{con} is given next.

Proposition 1 (CON: Equilibrium Strategy). *In an all-pay auction with incomplete information and a stochastic population of weakly risk-averse agents, given a constant prize* M_0*, the equilibrium strategy* z^*_{con} *is given by*

$$\frac{\mathrm{d}h(z^*_{con})}{\mathrm{d}s} + \frac{1 - B^*(s)}{s} h(z^*_{con}) = M_0 \frac{P'(s)}{s} + \frac{A^{*\prime}(s)}{s}, \tag{25}$$

where $A^*(s)$ *and* $B^*(s)$ *are defined in* (A.6).

Proof. Substituting M_0 for $M(z^*)$ in (A.5) (cf. proof of Theorem 1) and differentiating (A.5) with respect to s proves the result. □

In (25), $A^*(s)$ and $B^*(s)$ depend on $z^*_{rn,con,}$ which is given next.

Proposition 2 (RN-CON: Equilibrium Strategy). *In an all-pay auction with incomplete information and a stochastic population of risk-neutral agents, given a constant prize* M_0*, the equilibrium strategy is given by*

$$z^*_{rn,con} = h^{-1} \left(-M_0 \int_{s}^{\bar{s}} \frac{P'(s_1)}{s_1} \mathrm{d}s_1 \right). \tag{26}$$

Optimization procedure. Given the parameters for a specific application, first obtain $z^*_{rn,con}$ using Proposition 2 and then solve for z^*_{con} using Proposition 1. Next, substitute z^*_{con} together with other given parameters into (24), which explicitly spells out the optimization problem $\max_{M_0} \Omega^*_{con}$, and solve it (subject to $M_0 \geq 0$) using traditional optimization techniques such as Kuhn-Tucker conditions.

6.2. Special Case

In this case, the problem remains the same, and the profit at equilibrium follows from (3) to be

$$\Omega_0^* = n \int_{\underline{s}}^{\bar{s}} z_0^* \, \mathrm{d}F - M_0, \tag{27}$$

where z_0^* is given next.

Corollary 3 (RN-DP-CON: Equilibrium Strategy). *In an all-pay auction with incomplete information and a deterministic population of n risk-neutral agents, given a constant prize M_0, the equilibrium strategy is given by*

$$z_0^* = h^{-1}\left((n-1)M_0 \int_{s}^{\bar{s}} (1-F)^{n-2} \frac{F'}{s_1} \, \mathrm{d}s_1\right). \tag{28}$$

Proof. Substituting $P(s) = (1-F)^{n-1}$ into (26) proves the result. □

The result (28) coincides with the one in Moldovanu and Sela [2001], but we use a different method.

Optimization procedure. Given the parameters (in particular, the c.d.f. $F(\cdot)$ and the modulator function $h(\cdot)$) for a specific application, first obtain an explicit expression of z_0^* using (28). Then plug it into (27), and we can express the optimization problem $\max_{M_0} \Omega_0^*$, which can be solved similarly as in the general case shown earlier.

7. EVALUATION

Not only does this section evaluate the performance of our mechanism (as specified by Theorem 1) in comparison to two other incentive mechanisms, but equally important, it also demonstrates how to apply our analytical results, step by step, to obtain concrete mechanisms for various cases. The three incentive mechanisms are:

—Standard Prize: An all-pay auction with a normalized prize, which is the standard version and can be found in any auction textbook such as Krishna [2009] and Klemperer [2004].
—Optimal Prize: An all-pay auction with an optimized (constant) prize as solved in Section 6.
—Adaptive Prize: This is our proposed mechanism, presented in Section 4 and 5.

For each comparison, we consider both the general case and the special case specified by Section 6, and evaluate three performance metrics for each case:

—Maximum Profit (Ω): The total agent contribution minus the reward expense.
—Agent Utility (π): The expected utility of an agent of arbitrary type, then averaged over all of the types, i.e., $\mathbb{E}_s[\pi(s)]$, where π is expressed in (7).
—Social Welfare (W): The aggregate utility of all players, namely the principal and agents.

Moreover, we also examine the prizes allocated by the three mechanisms in the special case to shed light on how the respective performances are achieved.

7.1. Simulation Setup

For a more concrete understanding, let us revisit the crowdsourcing applications in Figure 1, where smartphone users contribute WiFi signal qualities, GPS readings, or noise measurements to a crowdsourcer. A software agent runs on each participant's phone to compute the best contribution strategy and ensure the execution of the strategy.[7] At the end of the crowdsourcing campaign or periodically (e.g., every month), the crowdsourcer selects the top contributor as the winner and rewards her with a prize that is adapted to her contribution level.

The simulation setup is as follows. Each agent is characterized by her marginal cost s of making a contribution, which is private information known to herself only. The common belief is that all the marginal costs are independently and uniformly drawn from $[1, 2]$, i.e., $F(s) = s - 1$. An agent i who makes a total contribution of z_i will incur a total cost of $s_i h(z_i)$, where $h(z) = z^w$, $w > 1$. This corresponds to a superlinear increase of cumulative cost when contribution accrues, which is common in practice. For example, if z measures the value of information (VoI) and $h(z)$ is the time or effort spent on producing the VoI, then as long as the production process exhibits diminishing marginal returns (which is a prevalent phenomenon in practice), producing more VoI will consume increasingly more time or effort. For another example, suppose that the production process is linear instead, i.e., producing z consumes a proportional amount of time or effort. But as the time/effort expense becomes higher, participating in the crowdsourcing activity interferes increasingly more with the agent's regular life and work, or she simply becomes increasingly more impatient, then the same cost behavior arises as captured by $h(z)$. For numeric computation, we take the quadratic form, i.e., $w = 2$.

Agents' risk aversion profile is characterized by a vNM function $u(x) = x - \epsilon x^2$, where $0 \le \epsilon \ll 1$ ($\epsilon = 0$ indicating risk neutral). The stochastic population has an uncertain size that follows a probability mass function $p_n = \frac{1}{N}$, $n = 2, 3, \ldots, N + 1$. The upper bound, $N + 1$, could be the total number of registered users, which can be retrieved from the campaign database.

7.2. Evaluation for General Case

7.2.1. Adaptive Prize. Theorem 1 applies to this case.

Step 1. Compute $G(s)$ and $G'(s)$.

$$
\begin{aligned}
G(s) &= \int_1^s \frac{\sum_{n=2}^{N+1} n(2-s_1)^{n-1}}{\sum_{n=2}^{N+1}(2-s_1)^{n-1}}\,\mathrm{d}s_1 = \int_1^s \frac{-\left[\sum_{n=2}^{N+1}(2-s_1)^n\right]'_{s_1}}{\sum_{n=2}^{N+1}(2-s_1)^{n-1}}\,\mathrm{d}s_1 \\
&= \int_{2-s}^1 \frac{\left[\frac{t^2(1-t^N)}{1-t}\right]'_t}{\frac{t(1-t^N)}{1-t}}\,\mathrm{d}t \quad (\text{denote } t := 2 - s_1) \\
&= \int_{2-s}^1 \frac{\frac{\left[2t-(N+2)t^{N+1}\right](1-t)+t^2-t^{N+2}}{(1-t)^2}}{\frac{t(1-t^N)}{1-t}}\,\mathrm{d}t = \int_{2-s}^1 \left[(N+1) - \frac{N}{1-t^N} + \frac{1}{1-t}\right]\mathrm{d}t.
\end{aligned}
$$

This integral exists despite that the second and third terms in the integrand appear to be individually unbounded at $t = 1$. To see this, check the limit

$$
\lim_{t\to 1}\left(-\frac{N}{1-t^N} + \frac{1}{1-t}\right) = \lim_{t\to 1}\frac{1-t^N-N(1-t)}{(1-t)(1-t^N)} \xlongequal{\text{L'Hôpital's rule}} \lim_{t\to 1}\frac{N-Nt^{N-1}}{2Nt^N - Nt^{N-1} - 1} = 0.
$$

[7]To ensure strategy execution, the software either takes over the task from the user (e.g., by periodically turning on sensors and sending sensing data to a server) or prompts the user in such a way that she stops contributing when z_i^* is reached (or simply enforces it by turning off a component).

Thus, we can proceed to calculate, e.g., when $N = 2$ ($n = 2, 3$; $p_n = 1/2$),

$$\begin{aligned} G(s) &= 3(s-1) - \ln(1+t)\big|_{2-s}^{1} = 3s + \ln\frac{3-s}{2} - 3, \\ G'(s) &= 3 - \frac{1}{3-s}. \end{aligned}$$

Step 2. Compute $\mathring{A}(s)$ and $\mathring{B}(s)$.

First we solve for $\mathring{M}_{rn}(s)$ and $\mathring{z}_{rn}(s)$ using Corollary 1 (RN). So we rewrite (17) in terms of s (in place of z as in $z = \mathring{z}_{rn}(s)$):

$$\mathring{M}_{rn}(s) = \left[s\mathring{z}_{rn}^2(s) + \mathring{\pi}_{rn}(s)\right]/P(s), \tag{29}$$

$$\text{where} \qquad \mathring{\pi}_{rn}(s) := \int_s^2 \mathring{z}_{rn}^2(s_1)\,\mathrm{d}s_1. \tag{30}$$

Then, using (6),

$$\begin{aligned} P(s) &= \frac{1}{N}\sum_{n=2}^{N+1}(2-s)^{n-1} = \frac{(2-s)[1-(2-s)^N]}{(s-1)N} \\ &= (2-s)(3-s)/2 \qquad (\text{when } N = 2). \end{aligned} \tag{31}$$

For $\mathring{z}_{rn}(s)$, because $a = \sum_{n=2}^{N+1} np_n = \frac{N+3}{2}$ and $(h')^{-1}(x) = x/2$, it follows from (18) that

$$\begin{aligned} \mathring{z}_{rn}(s) &= \frac{N+3}{4}\Big/[G'(s)s + G(s)] \\ &= \frac{5/4}{6s - \frac{3}{3-s} + \ln\frac{3-s}{2} - 2} \qquad (\text{when } N = 2). \end{aligned}$$

Thus, and using $u_1(x) = -x^2$, we can spell out (15) as

$$\begin{aligned} \mathring{A}(s) &= \epsilon P(s)[2\mathring{M}_{rn}(s)h(\mathring{z}_{rn})s - \mathring{M}_{rn}^2(s)] - \epsilon h^2(\mathring{z}_{rn})s^2 \\ &= -\tfrac{\epsilon}{P(s)}\left[\mathring{\pi}_{rn}^2 - 2(1-P)\mathring{\pi}_{rn}\mathring{z}_{rn}^2 s + (1-P)\mathring{z}_{rn}^4 s^2\right], \end{aligned} \tag{32}$$

$$\begin{aligned} \mathring{B}(s) &= -2\epsilon P(s)\mathring{M}_{rn}(s) + 2\epsilon h(\mathring{z}_{rn})s + 1 \\ &= 1 - 2\epsilon\mathring{\pi}_{rn}. \end{aligned} \tag{33}$$

Step 3. Compute $\mathring{\Omega}$.

Now we apply Theorem 1 to obtain, first, the equilibrium strategy according to (13):

$$\mathring{z}(s) = \frac{(N+3)/4}{G'(s)s + G(s)(1 - 2\epsilon\mathring{\pi}_{rn})}, \tag{34}$$

and next, the maximum profit using (14):

$$\begin{aligned} \mathring{\Omega} &= \int_1^2 \left[\frac{(N+3)\mathring{z}}{2} - \mathring{z}^2\left[sG' + (1-2\epsilon\mathring{\pi}_{rn})G\right] + \mathring{A}(s)G'\right]\mathrm{d}s \\ &= \int_1^2 \left[\frac{(N+3)^2/16}{G'(s)s + G(s)(1-2\epsilon\mathring{\pi}_{rn})} + \mathring{A}(s)G'\right]\mathrm{d}s, \end{aligned}$$

where $\mathring{\pi}_{rn}(s)$ can be computed according to (30) numerically (using, e.g., `NIntegrate` of *Mathematica*). Then, we apply the following approximation to compute $\mathring{\Omega}$ for different values of ϵ:

$$\mathring{\Omega} \approx \sum_{k=0}^{m-1} I(1 + k\Delta s) \cdot \Delta s,$$

where $I(s)$ denotes the integrand of $\mathring{\Omega}$, and $m = 1/\Delta s$, in which $\Delta s \ll 1$. In our numerical computation, we chose $m = 200$ because we found that when $\epsilon = 0.1$, e.g., an increase of m from 50 to 100 improves precision by 1% and from 100 to 200 improves by 0.7%.

Step 4. Compute agent utility and social welfare.

The expected utility of an agent of arbitrary type s can be written using (12) as

$$\mathring{\pi}(s) \;=\; P(s)\mathring{M}(s) - h(\mathring{z})s = \int_s^2 \mathring{B}(s_1)h(\mathring{z}(s_1))\,\mathrm{d}s_1 - \mathring{A}(s), \tag{35}$$

where $\mathring{A}$, $\mathring{B}$, $\mathring{z}$ have all been obtained in the preceding. Its mean value is given by

$$\mathop{\mathbb{E}}_s[\mathring{\pi}(s)] = \int_1^2 \mathring{\pi}(s)\,\mathrm{d}F,$$

which we compute using `NIntegrate`. The social welfare is then obtained as

$$W = \mathring{\Omega} + \frac{N+3}{2}\mathop{\mathbb{E}}_s[\mathring{\pi}(s)].$$

7.2.2. Optimal Prize. We follow the optimization procedure outlined in Section 6.1.

Step 1. Obtain $z^*_{rn,con}$ using Proposition 2.

First, calculate $P'(s)$ using (31):

$$\begin{aligned} P'(s) \;&=\; \frac{(2-s)^N[N(s-1)+1]-1}{N(s-1)^2} \\ &=\; (2s-5)/2 \qquad (\text{when } N=2). \end{aligned}$$

Then, we obtain $z^*_{rn,con}$ using Proposition 2:

$$\begin{aligned} z^*_{rn,con} \;&=\; \sqrt{-M_0 \int_s^2 \tfrac{P'(s_1)}{s_1}\,\mathrm{d}s_1} \\ &=\; \sqrt{\left[\tfrac{5}{2}\ln\tfrac{2}{s} - (2-s)\right] M_0} \qquad (\text{when } N=2). \end{aligned}$$

Step 2. Obtain z^*_{con} using Proposition 1.

First, to solve for $A^*(s)$, we take (32) as a shortcut:

$$A^*(s) \;=\; \epsilon P(s)(2M_0 {z^*_{rn,con}}^2 s - M_0^2) - \epsilon {z^*_{rn,con}}^4 s^2.$$

Hence, when $N = 2$,

$$\begin{aligned} A^*(s) \;&=\; \epsilon M_0^2 \left[(2-s)(3-s)\left(s\left[\frac{5}{2}\ln\frac{2}{s} - (2-s)\right] - \frac{1}{2}\right) - s^2\left[\frac{5}{2}\ln\frac{2}{s} - (2-s)\right]^2\right], \\ A^{*\prime}(s) \;&=\; -\epsilon M_0^2\left[2sq^2 - pq + \left(s - \frac{5}{2}\right)(1-p)\right], \end{aligned}$$

where $q := \frac{5}{2}\ln\frac{2}{s} - (2-s)$ and $p := (2-s)(3-s)$.

To solve for $B^*(s)$, we take (33) as a shortcut:

$$B^*(s) \;=\; -2\epsilon P(s)M_0 + 2\epsilon {z^*_{rn,con}}^2 s + 1.$$

Hence, when $N = 2$, $B^*(s) = 2\epsilon M_0(sq - \frac{p}{2}) + 1$.

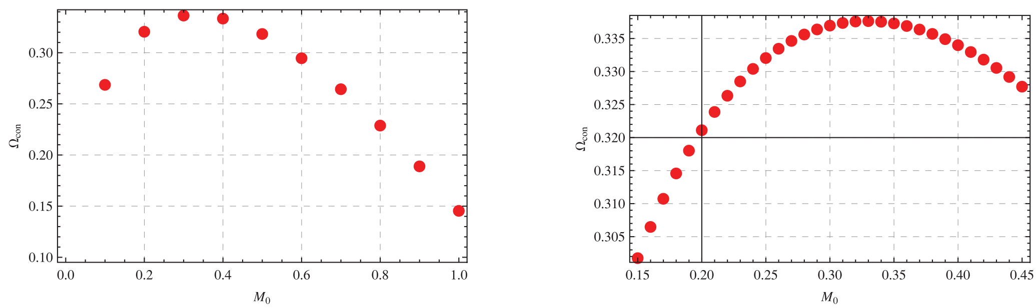

(a) Coarse-grained search (step size 0.1). (b) Fine-grained search (step size 0.01).

Fig. 3. Finding the maximum profit for Optimal prize: the general case (Section 7.2.2, Step 3). $\epsilon = 0.1$, $N = 2$.

Therefore, Proposition 1 gives, when $N = 2$,

$$H' - 2\epsilon M_0 \left(q - \frac{p}{2s}\right) H = Q(s), \tag{36}$$

where $H(s) := {z^*_{con}}^2(s)$, $Q(s) := (1 - \frac{5}{2s})M_0 - \epsilon M_0^2[2q^2 - \frac{pq}{s} + (1 - \frac{5}{2s})(1 - p)]$, and the boundary condition is $H(2) = 0$. Note that (36) is a first-order linear differential equation.

Step 3. Solve $\max_{M_0} \Omega^*_{con}$.

Rewrite (24) with $p_n = 1/N$, as

$$\Omega^*_{con} = \frac{N+3}{2} \int_1^2 z^*_{con}\, \mathrm{d}F - M_0. \tag{37}$$

The typical way to solve $\max_{M_0} \Omega^*_{con}$, if using *Mathematica*, is to first solve (36) using `DSolve`, next plug z^*_{con} $(=\sqrt{H})$ into (37), and then apply `Maximize` or `NMaximize`. In our case, however, `DSolve` does not yield a result when the boundary condition is specified. Therefore, we use numerical search as follows. We assign a sequence of values to M_0 and apply `NDSolve` to (36) to obtain a corresponding sequence of z^*_{con}. With this sequence of z^*_{con}, we compute a corresponding sequence of Ω^*_{con} using (37), and the maximum Ω^*_{con} is thus found. Of course, the range of the sequence must be chosen such that it includes the maximum point. Furthermore, for a fast and accurate search, we first perform a coarse-grained search to locate a narrower range of the maximizer $\mathring{M}_0$, and then, within that range, we perform a fine-grained search to pinpoint the $\mathring{M}_0$. An illustration is given by Figure 3, where the maximum $\mathring{\Omega}_{con} = 0.337747$ is found at $\mathring{M}_0 = 0.33$.

Step 4. Compute agent utility and social welfare.

The expected utility of an agent of arbitrary type s is

$$\mathring{\pi}_{con}(s) = P(s)\mathring{M}_0 - h(\mathring{z}_{con})s,$$

whose mean value is $\mathbb{E}_s[\mathring{\pi}_{con}(s)] = \int_1^2 \mathring{\pi}_{con}(s)\, \mathrm{d}F$. The social welfare is then obtained as $\mathring{W}_{con} = \mathring{\Omega}_{con} + \frac{N+3}{2}\, \mathbb{E}_s[\mathring{\pi}_{con}(s)]$.

7.2.3. Standard Prize. In this case, the crowdsourcing campaign assumes a normalized prize $M_0 = 1$. Substitute it into (36) to compute z^*_{std} (in place of z^*_{con}) and (37) to compute Ω^*_{std} (in place of Ω^*_{con}).

The procedures of computing agent utility and social welfare are similar to the preceding and hence are omitted.

7.2.4. Results. We compare the three metrics, namely maximum profit, agent utility, and social welfare, for Adaptive Prize, Optimal Prize, and Standard Prize in Figure 4.

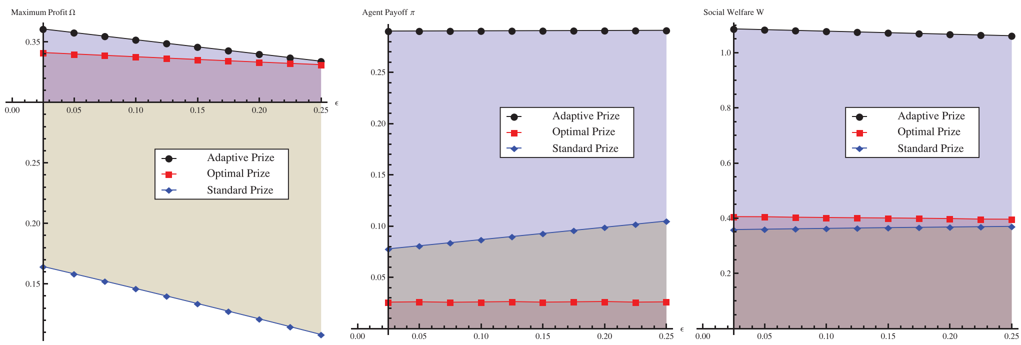

(a) Principal's profit. (b) Agent's utility. (c) Social welfare.

Fig. 4. Comparison among three all-pay auction based incentive mechanisms: general case.

For profit as shown in Figure 4(a), Adaptive Prize constantly outperforms both Optimal Prize and Standard Prize over the entire range of risk aversion parameter ϵ. Note that ϵ cannot be too large, as otherwise it violates the assumption of weakly risk aversion. We also see that when ϵ increases (agents become more averse to risk), the profits of all three incentive mechanisms decrease. This conforms to our intuition, as risk-averse agents are more reluctant to lose than risk-neutral agents, and therefore they become more conservative in spending effort. Particularly, Standard Prize, which is without optimization, decreases the fastest among the three mechanisms.

For agent utility as shown in Figure 4(b), Adaptive Prize is the clear winner (i.e., about 11.2 times of Optimal Prize and 3.35 times of Standard Prize at $\epsilon = 0.1$). Optimal Prize turns out to be lower than Standard Prize, because its optimization aims at maximizing the principal's profit rather than agent utility, and this result reveals that it has to slightly sacrifice agent utility (while still ensuring IR) to achieve the aim. Another observation is that Standard Prize can be taken advantage of by agents (and not the principal) who can gain slightly higher expected utility by being more and more risk averse. This is due to their conservation of effort and leads to the lower profit for the principal (cf. Figure 4(a)). On the other hand, Adaptive Prize and Optimal Prize are much more resistant to risk aversion.

Last, on social welfare as shown in Figure 4(c), Adaptive Prize again wins over the other mechanisms remarkably. We also see that Optimal Prize only outperforms Standard Prize by a small margin. This implies that our exploration on the new dimension of optimal incentive mechanism design—instrumenting an adaptive prize as a function—is worthwhile as compared to the traditional optimization approach, which is to optimize the constant prize value.

7.3. Evaluation for Special Case

In the special case considered in this section, agents are risk neutral and the population size is known as n.

7.3.1. Adaptive Prize. The equilibrium strategy follows from Corollary 2 (22) as $\mathring{z}_{rn,dp} = \frac{1}{4s-2}$, or inversely, $\mathring{s}_{rn,dp} = \frac{1}{4z} + \frac{1}{2}$. The optimal adaptive prize then follows from (21) as

$$\mathring{M}_{rn,dp}(z) = \frac{\frac{z}{4} + \frac{z^2}{2} - \int_{\frac{1}{6}}^{z}\left(-\frac{1}{4}\right) dz_1}{\left(\frac{3}{2} - \frac{1}{4z}\right)^{n-1}} = \frac{(4z)^n(z+1)}{8(6z-1)^{n-1}} - \frac{1}{24}\left(\frac{4z}{6z-1}\right)^{n-1},$$

or equivalently, in terms of s, as

$$\mathring{M}_{rn,dp}(s) = \frac{4s-1}{8(2-s)^{n-1}(2s-1)^2} - \frac{1}{24(2-s)^{n-1}}. \tag{38}$$

Thus, the maximum profit is obtained from (23) as

$$\mathring{\Omega}_{rn,dp} = n\int_1^2 \frac{\mathrm{d}s}{4(2s-1)} = \frac{\ln 3}{8}n. \tag{39}$$

The expected utility of an agent of arbitrary type s is

$$\mathring{\pi}_{rn,dp}(s) = (1-F)^{n-1}\mathring{M}_{rn,dp}(s) - h(\mathring{z}_{rn,dp})s = \frac{1}{8(2s-1)} - \frac{1}{24},$$

whose mean value is

$$\mathop{\mathbb{E}}_s[\mathring{\pi}_{rn,dp}(s)] = \int_1^2 \mathring{\pi}_{rn,dp}(s)\,\mathrm{d}F = \frac{\ln 3}{16} - \frac{1}{24}. \tag{40}$$

The social welfare is thus

$$\mathring{W}_{rn,dp} = \mathring{\Omega}_{rn,dp} + n\mathop{\mathbb{E}}_s[\mathring{\pi}_{rn,dp}(s)] = \frac{3\ln 3}{16}n - \frac{1}{24}. \tag{41}$$

7.3.2. Optimal Prize. We follow the optimization procedure outlined in Section 6.2.
Step 1. Compute equilibrium strategy.
First, Corollary 3 gives the equilibrium strategy

$$z_0^*(s) = \sqrt{(n-1)M_0\int_s^2 \frac{(2-s_1)^{n-2}}{s_1}\,\mathrm{d}s_1} = \sqrt{\frac{(n-1)2^n}{8}\left(\ln\frac{4}{s^2} - 2\mathcal{H} + s(n-2)\mathcal{G}\right)M_0},$$

where $\mathcal{H}$ stands for $\mathcal{H}(n-2)$, which gives the $(n-2)^{\text{th}}$ harmonic number,[8], and $\mathcal{G}$ stands for $\mathcal{G}(1,1,3-n;2,2;\frac{s}{2})$, which is a generalized hypergeometric function.[9] For example, when $n = 2, 3, 4$,

$$z_0^*(s) = \begin{cases} \sqrt{\frac{M_0}{2}\ln\frac{4}{s^2}}, & n=2 \\ \sqrt{2\left(s-2+\ln\frac{4}{s^2}\right)M_0}, & n=3 \\ \sqrt{6\left(2s(1-\frac{s}{8})-3+\ln\frac{4}{s^2}\right)M_0}, & n=4. \end{cases}$$

Step 2. Profit maximization.
The profit at equilibrium can be obtained using (27):

$$\begin{aligned} \Omega_0^* &= n\int_1^2 z_0^*\,\mathrm{d}F - M_0 & (42)\\ &= \begin{cases} 2\sqrt{M_0}\left(\sqrt{\pi}\,\mathrm{erf}(\sqrt{\ln 2}) - \sqrt{\ln 2}\right) - M_0, & n=2 \\ 3\int_1^2\sqrt{2M_0\left(s-2+\ln\frac{4}{s^2}\right)}\,\mathrm{d}s - M_0, & n=3 \\ 4\int_1^2\sqrt{6M_0\left(s\left(2-\frac{s}{4}\right)-3+\ln\frac{4}{s^2}\right)}\,\mathrm{d}s - M_0, & n=4. \end{cases} \end{aligned}$$

from which we obtain the maximum profit $\mathring{\Omega}_0$ and the optimal prize (maximizer) $\mathring{M}_0$, using `NMaximize` with *Mathematica*.

[8] $\mathcal{H}(0,1,2,3,4,5,\ldots) = 1, \frac{3}{2}, \frac{11}{6}, \frac{25}{12}, \frac{137}{60}, \frac{49}{20}, \ldots$
[9] $\mathcal{G}(1,1,\{1,0,-1,-2,\ldots\};2,2;1) = \frac{\pi^2}{6}, 1, \frac{3}{4}, \frac{11}{18}, \ldots$

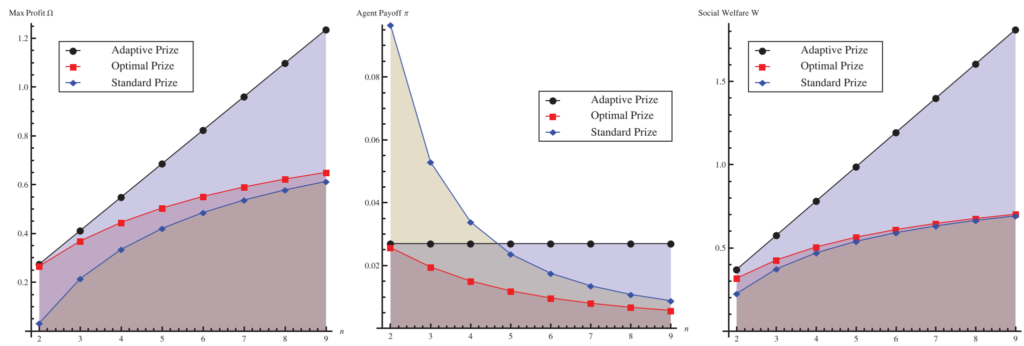

(a) Principal's profit. (b) Agent's utility. (c) Social welfare.

Fig. 5. Comparison among three all-pay auction–based incentive mechanisms: special case.

Step 3. Agent utility and social welfare.
The expected utility of an agent of arbitrary type s is

$$\mathring{\pi}_0(s) = (1-F)^{n-1}\mathring{M}_0 - h(\mathring{z}_0)s = (2-s)^{n-1}\mathring{M}_0 - \mathring{z}_0^2 s \tag{43}$$
$$= \begin{cases} 0.266(2-s) - 0.133s\ln\frac{4}{s^2}, & n=2 \\ 0.368(2-s)^2 - 0.738s(s-2+\ln\frac{4}{s^2}), & n=3 \\ 0.445(2-s)^3 - 2.673s\left[s(2-\frac{s}{4}) - 3 + \ln\frac{4}{s^2}\right], & n=4. \end{cases}$$

Its mean value, and the social welfare, can then be determined similarly as in Section 7.3.1.

7.3.3. Standard Prize. In this case, we plug $M_0 = 1$ into (42) and (43) to obtain the profit and agent utility, respectively. The mean of agent utility and the social welfare are then obtained similarly as in Section 7.3.1.

7.3.4. Results. The results are presented in Figure 5. First, on profit (Figure 5(a)), Adaptive Prize is again the clear winner. Notably, its profit increases *linearly* as the number of agents increases, whereas the profits of Optimal Prize and Standard Prize exhibit diminishing marginal return and tend to saturate. This implies a highly desirable *scalability*, which is important to crowdsourcing systems, and gives a competitive edge to our approach.

Second, see agent utility in Figure 5(b). The foremost observation is that the agent utilities of both Standard Prize and Optimal Prize decrease as n increases, but that of Adaptive Prize remains constant. In addition, Adaptive Prize maintains the highest agent utility among the three mechanisms except at $n = 2, 3, 4$. Therefore, since n is much larger in typical crowdsourcing scenarios, this set of results show that our approach provides a *much stronger incentive* to the agents. Moreover, since the agent utilities in all three mechanisms are positive, it indicates that a properly designed all-pay auction can indeed incentivize agents to participate.

Last, for social welfare as shown in Figure 5(c), we see a similar yet more prominent trend as compared to Figure 5(a). The similarity owes to the superposition of profit and n times agent utility (the multiplying factor n counteracts the decreasing agent utility of Standard Prize and Optimal Prize shown in Figure 5(b)). The prominence can be illustrated by a comparison. For instance, when $n = 9$, the profit of Adaptive Prize is about 2.0 and 1.9 times of Standard Prize and Optimal Prize, respectively, whereas its social welfare is about 2.6 times of the other two. We can also see that this difference will become larger and larger when n increases. This set of results implies that the community as a whole can derive substantial benefit from our proposed mechanism.

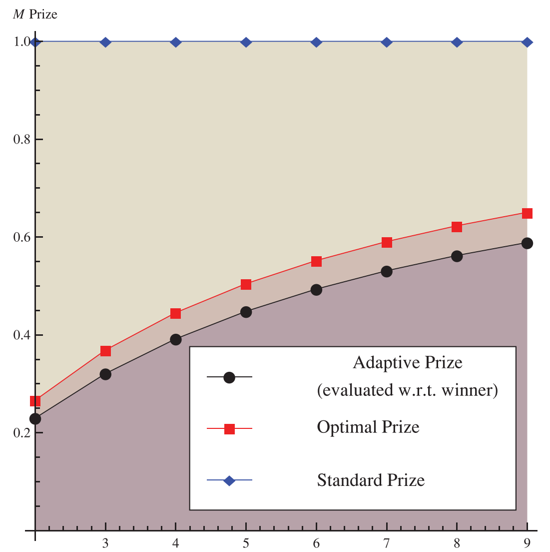

Fig. 6. Prize comparison among different mechanisms.

Investigating prizes. Naturally, one would wonder if Adaptive Prize has provisioned a much higher prize to achieve the superior performance. To this end, we evaluate its prize function (38) with respect to the winner's type, which is $\underline{s} + \frac{\bar{s}-\underline{s}}{n+1}$ in expectation due to the uniform distribution of agent type. Then we plot the function value with respect to n in Figure 6, together with (i) Optimal Prize, whose optimum $\mathring{M}_0$ is obtained in Step 2 of Section 7.3.2, and (ii) Standard Prize, whose prize remains one.

Counterintuitively, we see that Adaptive Prize gives out the *lowest* prize among all mechanisms. The insight to draw from this observation is that it is not necessary to provide a higher prize to incentivize agents so as to be more profitable; what really matters is how the whole mechanism is designed such that agents are incentivized to an "optimal" extent, whereas the principal can live on a budget that is as small as possible. This rationale is also reflected by Optimal Prize: although it has to offer slightly higher reward than Adaptive Prize, it is still much more budget friendly than Standard Prize.

Nonetheless, it may still be a little perplexing as to why a lower prize can actually achieve a higher profit. Thus, diving deeper, we examine how the prize function of Adaptive Prize varies with agent type and contribution, as plotted in Figure 7, where we fix $n = 6$ for illustration. Figure 7(a) reveals that when agent type is higher, the prize increases exponentially. This implies that Adaptive Prize "dares" to provide very high prize to incentive those agents who are very reluctant to participate because of their high contribution costs, whereas the other two mechanisms only provide constant prizes that do not appeal to those "weak" agents. On the other hand, Figure 7(b) surprisingly shows that making higher contribution actually wins a lower prize. In fact, this unveils a sensible design principle: higher contributions are made by lower-cost agents, who are already more willing to participate than others and do not need a high prize so as to be incentivized; in other words, by contributing more, their benefit from increasing the winning odds *outweighs* their cost increase. Therefore, this insight would allow a crowdsourcer to *downsize* its offered prize insofar as this outweighing leverage is not negated, which results in substantial cost saving and eventually profit gain.

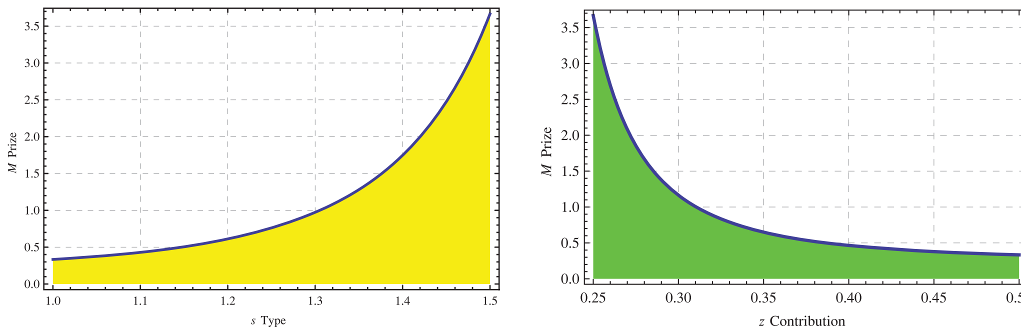

(a) Adaptive prize versus winning agent's type.

(b) Adaptive prize versus winning agent's contribution.

Fig. 7. Investigating adaptive prize with respect to the winning agent; $n = 6$. Note that $s = 1$ corresponds to $z = 0.5$ and $s = 1.5$ corresponds to $z = 0.25$.

8. DISCUSSION

8.1. Population Size Indifference and Shading Effect

It is interesting to note that in the case of RN-DP (the most common case in the literature), Adaptive Prize induces an agent contribution strategy (22) that is independent of the population size n. This is in contrast to all well-known standard auctions, regardless of all-pay or winner-pay, in which the agent strategy always depends on n (e.g., see Krishna [2009] and (28) herein). Indeed, this (coupling with n) conforms to our intuition, because when n increases, an auction becomes more competitive since one has to outbid all others to win, and hence an agent should accordingly adjust her strategy.

So why are agents *oblivious* to this number now (with our approach)?

To answer this question, let us first explain how n affects agent strategy in standard auctions. The "bridge" between n and strategy is the *winning probability*. For example, in IPV model–based all-pay auctions, the winning probability is $(1 - F(s))^{n-1}$, where agent type s denotes cost. Therefore, the chance of winning drops if n increases, and thus any rational agent will shade her bid *downward* to minimize the "more likely" loss. We refer to this as a shading effect.

In our proposed approach, the principal is offered another degree of freedom to *functionize* the prize, which is then subjected to optimization. As a result, the winning probability—which contains the shading factor n—becomes part of the prize function (see (21) in the denominator). Thus, population size is no longer a concern of agents but is taken care of by the principal, thereby overcoming the undesired shading effect.

8.2. Limitations and Applicable Scenarios

The main challenge of applying all-pay auctions in general to practice is probably not technological but psychological. This is because the all-pay nature requires everyone to pay for her bid (i.e., to actually contribute) without any guaranteed reward. Although each person can anticipate a positive "expected" utility (by virtue of SIR), it still sounds lacking a sense of security or reassurance. Therefore, how to actually run such mechanisms in practice may need extra assistance, where behavioral economics and marketing strategies would play interesting roles.

On the other hand, a very recent study presents a user-friendly alternative called *Tullock contests* [Luo et al. 2015], which is complementary to all-pay auctions. The most distinctive feature of Tullock contests is that "everyone has a chance to win" no matter how "weak" she is.

Regarding the general criticism on game theory that human rationality is often limited or bounded in reality, we have overcome this limitation by equipping each user with a software agent to compute and ensure execution of the optimal contribution strategy on each user's behalf rather than assuming that users will do all of the calculations themselves. This has been described in Section 7.

Other than these, our proposed approach can be generally applied to a broad range of crowdsourcing systems. Such systems could be traditional (problem-solving) crowdsourcing Web sites such as Amazon Mechanical Turk [2005], CrowdFlower [2007], and TaskRabbit [2008], or new generation (data-centric) crowdsensing mobile platforms such as Waze [2009] and NoiseTube [2010]. The application domain would generally cover transportation, environmental, healthcare, communications and networking, and so forth.

9. CONCLUSION

This article advocates using all-pay auctions to design incentive mechanisms for crowdsourcing, based on their two merits that we identify: simplicity and risk elimination. In this spirit, we propose a particular all-pay auction approach that features an adaptive prize and a general crowdsourcing environment. We show that this approach generates significantly higher (and scalable) profit for the crowdsourcer and offers stronger incentive to participants, and it provides much better social welfare to the whole community.

Our approach of using an adaptive prize or reward also represents a new dimension of optimal incentive mechanism design. The accompanying general crowdsourcing environment provides a more realistic model to permit a wider range of applications. Finally, the theoretical results for the special cases (as outlined in Section 5) would allow for a convenient use by other future research.

APPENDIX

The Appendix is available in the ACM Digital Library (http://dl.acm.org).

REFERENCES

Amazon Mechanical Turk. 2005. Mechanical Turk Home Page. Retrieved June 5, 2015, from http://www.mturk.com.

Nikolay Archak and Arun Sundararajan. 2009. Optimal design of crowdsourcing contests. In *Proceedings of the 30th International Conference on Information Systems*.

Michael R. Baye, Dan Kovenock, and Casper G. de Vries. 1996. The all-pay auction with complete information. *Economic Theory* 8, 2, 291–305.

Shuchi Chawla, Jason D. Hartline, and B. Sivan. 2012. Optimal crowdsourcing contests. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*.

Chen Cohen, Todd Kaplan, and Aner Sela. 2008. Optimal rewards in contests. *RAND Journal of Economics* 39, 2, 434–451.

CrowdFlower. 2007. CrowdFlower Home Page. Retrieved June 5, 2015, from http://www.crowdflower.com.

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

Zhenni Feng, Yanmin Zhu, Qian Zhang, Lionel Ni, and Athanasios Vasilakos. 2014. TRAC: Truthful auction for location-aware collaborative sensing in mobile crowdsourcing. In *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14)*. 1231–1239.

Gadi Fibich and Arieh Gavious. 2003. Asymmetric first-price auctions—a perturbation approach. *Mathematics of Operations Research* 28, 4, 836–852.

Gadi Fibich, Arieh Gavious, and Aner Sela. 2006. All-pay auctions with risk-averse players. *International Journal of Game Theory* 34, 583–599.

Raghu K. Ganti, Fan Ye, and Hui Lei. 2011. Mobile crowdsensing: Current state and future challenges. *IEEE Communications Magazine* 49, 11, 32–39.

Arpita Ghosh and Preston McAfee. 2011. Incentivizing high-quality user-generated content. In *Proceedings of the 20th International Conference on World Wide Web (WWW'11)*. 137–146.

Bin Guo, Alvin Chin, Zhiwen Yu, Runhe Huang, and Daqing Zhang. 2015. An introduction to the special issue on participatory sensing and crowd intelligence. *ACM Transactions on Intelligent Systems and Technology* 6, 3, 36:1–36:4.

Bin Guo, Zhiwen Yu, Xingshe Zhou, and Daqing Zhang. 2014. From participatory sensing to mobile crowd sensing. In *Proceedings of the PERCOM Workshops*. 593–598.

Kyungsik Han, Eric A. Graham, Dylan Vassallo, and Deborah Estrin. 2011. Enhancing motivation in a mobile participatory sensing project through gaming. In *Proceedings of the 2011 IEEE 3rd International Conference on Social Computing (SocialCom'11)*. 1443–1448.

John C. Harsanyi. 1967/1968. Games with incomplete information played by Bayesian players, part I-III. *Management Science* 14, 3,5,7, 159–183,320–334,486–502.

Ronald M. Harstad, John H. Kagel, and Dan Levin. 1990. Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters* 33, 35–40.

Moshe Haviv and Igal Milchtaich. 2012. Auctions with a random number of identical bidders. *Economics Letters* 114, 2, 143–146.

Todd Kaplan, Israel Luski, Aner Sela, and David Wettstein. 2002. All-pay auctions with variable rewards. *Journal of Industrial Economics* 50, 4, 417–430.

Nicolas Kaufmann, Thimo Schulze, and Daniel Veit. 2011. More than fun and money. Worker motivation in crowdsourcing—a study on Mechanical Turk. In *Proceedings of the 17th Americas Conference on Information Systems (AMCIS'11)*.

Paul Klemperer. 2004. *Auctions: Theory and Practice*. Princeton University Press.

Iordanis Koutsopoulos. 2013. Optimal incentive-driven design of participatory sensing systems. In *Proceedings of IEEE INFOCOM*.

Vijay Krishna. 2009. *Auction Theory* (2nd ed.). Academic Press, New York, NY.

Juong-Sik Lee and Baik Hoh. 2010. Sell your experiences: A market mechanism based incentive for participatory sensing. In *Proceedings of the 2010 IEEE International Conference on Pervasive Computing and Communications (PerCom'10)*. 60–68.

Dan Levin and Emre Ozdenoren. 2004. Auctions with uncertain numbers of bidders. *Journal of Economic Theory* 118, 229–251.

Tie Luo, Salil S. Kanhere, Sajal K. Das, and Hwee-Pink Tan. 2014b. Optimal prizes for all-pay contests in heterogeneous crowdsourcing. In *Proceedings of the IEEE International Conference on Mobile Ad-Hoc and Sensor Systems (MASS'14)*. 136–144.

Tie Luo, Salil S. Kanhere, Sajal K. Das, and Hwee-Pink Tan. 2016. Incentive mechanism design for heterogeneous crowdsourcing using all-pay contests. *IEEE Transactions on Mobile Computing*. DOI:http://dx.doi.org/10.1109/TMC.2015.2485978

Tie Luo, Salil S. Kanhere, and Hwee-Pink Tan. 2014a. SEW-ing a simple endorsement Web to incentivize trustworthy participatory sensing. In *Proceedings of the IEEE International Conference on Sensing, Communication, and Networking (SECON'14)*. 636–644.

Tie Luo, Salil S. Kanhere, Hwee-Pink Tan, Fan Wu, and Hongyi Wu. 2015. Crowdsourcing with Tullock contests: A new perspective. In *Proceedings of the 34th Annual IEEE International Conference on Computer Communications (INFOCOM'15)*. 2515–2523.

Tie Luo, Hwee-Pink Tan, and Lirong Xia. 2014. Profit-maximizing incentive for participatory sensing. In *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14)*. 127–135.

Tie Luo and Chen-Khong Tham. 2012. Fairness and social welfare in incentivizing participatory sensing. In *Proceedings of the 2012 9th Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks (SECON'12)*. 425–433.

Suhas Mathur, Tong Jin, Nikhil Kasturirangan, Janani Chandrasekaran, Wenzhi Xue, Marco Gruteser, and Wade Trappe. 2010. ParkNet: Drive-by sensing of road-side parking statistics. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services (MobiSys'10)*. 123–136.

Preston McAfee and John McMillan. 1987. Auctions with a stochastic number of bidders. *Journal of Economic Theory* 43, 1–19.

Benny Moldovanu and Aner Sela. 2001. The optimal allocation of prizes in contests. *American Economic Review* 91, 3, 542–558.

Roger Myerson. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1, 58–73.

NoiseTube. 2010. NoiseTube Home Page. Retrieved June 5, 2015, from http://www.noisetube.net.

Francesco Restuccia and Sajal K. Das. 2014. FIDES: A trust-based framework for secure user incentivization in participatory sensing. In *Proceedings of the 2014 IEEE 15th International Symposium on a World of Wireless, Mobile, and Multimedia Networks (WoWMoM'14).* 1–10.

Task Rabbit. 2008. TaskRabbit Home Page. Retrieved January 15, 2015, from https://www.taskrabbit.com.

Chen-Khong Tham and Tie Luo. 2013. Quality of contributed service and market equilibrium for participatory sensing. In *Proceedings of the IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS'13)*. 133–140.

Chen Khong Tham and Tie Luo. 2014. Fairness and social welfare in service allocation schemes for participatory sensing. *Elsevier Computer Networks* 73, 58–71.

Chen Khong Tham and Tie Luo. 2015. Quality of contributed service and market equilibrium for participatory sensing. *IEEE Transactions on Mobile Computing* 14, 4, 829–842.

Waze. 2009. Waze Home Page. Retrieved June 5, 2015, from http://www.waze.com.

WiFi-Scout. 2014. WiFi-Scout Home Page. Retrieved June 5, 2015, from http://wifi-scout.sns-i2r.org.

Fang-Jing Wu and Tie Luo. 2014. WiFiScout: A crowdsensing wifi advisory system with gamification-based incentive. In *Proceedings of the IEEE International Conference on Mobile Ad-Hoc and Sensor Systems (MASS'14)*. 533–534.

Yingce Xia, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2014. Incentivizing high-quality content from heterogeneous users: On the existence of Nash equilibrium. arXiv:1404.5155 [cs.GT].

Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. 2012. Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking (MobiCom'12)*. 173–184.

Daqing Zhang, Leye Wang, Haoyi Xiong, and Bin Guo. 2014. 4W1H in mobile crowd sensing. *IEEE Communications Magazine* 52, 8, 42–48.

Yu Zhang and Mihaela van der Schaar. 2012. Reputation-based incentive protocols in crowdsourcing applications. In *Proceedings of the 31st Annual International Conference on Computer Communications (INFOCOM'12)*.

Dong Zhao, Xiang-Yang Li, and Huadong Ma. 2014. How to crowdsource tasks truthfully without sacrificing utility: Online incentive mechanisms with budget constraint. In *Proceedings of the 33rd Annual International Conference on Computer Communications (INFOCOM'14)*.

Received June 2015; revised August 2015; accepted October 2015