A Winners-Take-All Incentive Mechanism for Crowd-Powered Systems

Pengfei Jiang
School of Electrical, Computer and Energy Engineering
Arizona State University
pengfei.jiang@asu.edu

Weina Wang
School of Electrical, Computer and Energy Engineering
Arizona State University
and
Coordinated Science Laboratory
University of Illinois
weinaw@illinois.edu

Yao Zhou
School of Computing, Informatics, and Decision Systems Engineering
Arizona State University
yzhou174@asu.edu

Jingrui He
School of Computing, Informatics, and Decision Systems Engineering
Arizona State University
jingrui.he@asu.edu

Lei Ying
School of Electrical, Computer and Energy Engineering
Arizona State University
lei.ying.2@asu.edu

ABSTRACT

This paper studies incentive mechanisms for crowd-powered systems, including applications such as collection of personal data for big-data analytics and crowdsourcing. In big-data analytics using personal data, an individual may control the quality of reported data via a privacy-preserving mechanism that randomizes the answer. In crowdsourcing, the quality of the reported answer depends on the amount of effort spent by a worker or a team. In these applications, incentive mechanisms are critical for eliciting data/answers with target quality. This paper focuses the following two fundamental questions: *what is the minimum payment required to incentivize an individual to submit a data/answer with quality level ϵ?* and *what incentive mechanisms can achieve the minimum payment?*

Let ϵ_i denote the quality of the data/answer reported by individual i. In this paper, we first derive a lower bound on the minimum amount of payment required for guaranteeing quality level ϵ_i. Inspired by the lower bound, we propose an incentive mechanism, named Winners-Take-All (WINTALL). WINTALL first decides a winning answer based on the reported data, cost functions of individuals, and some prior distribution; and then pays to individuals whose reported data match the winning answer. Under some assumptions, we show that the expected payment of WINTALL matches the lower bound. In the application of private discrete distribution estimation, we show that WINTALL simply rewards individuals whose reported answers match the most popular answer from the reported ones (the prior distribution is not needed in this case).

NetEcon '18, June 18, 2018, Irvine, CA, USA

ISBN 978-1-4503-5916-0/18/06…$15.00

DOI: https://doi.org/10.1145/3230654.3230657

1. INTRODUCTION

A number of recent technology revolutions such as big-data analytics and crowdsourcing are powered by the *crowds*. The access to massive personal data and individual talent via online platforms enables new scientific discoveries, new personalized applications/services and new mechanisms of problem solving. The successes of these *crowd-powered* systems require the participation of, or access to, a massive population. In many systems, this massive participation is a result of a highly popular application or service, such as Gmail, iPhone or Amazon, which attracts millions of active users. In other systems such as Amazon Mechanical Turk, the massive participation is achieved using monetary incentives.

This paper focuses on the design of incentive mechanisms for crowd-powered systems which attract the crowd with monetary incentives instead of services. We consider a crowd-powered system where the platform elicits "answers" from a crowd. Depending on the applications, the answer may be a piece of personal information or a solution to a crowdsourcing task. The design of incentive mechanisms for these crowd-powered systems is interesting because it differs from other systems in the following aspects.

- The quality of the answer is "controlled" by an individual, and often "unverifiable" by the platform. For example, when collecting personal data, to protect individuals' privacy, the reported data may be randomized versions of the original data. In such a case, the platform does not have the access to the true data so will not be able to verify the truthfulness of the answers. Furthermore, when the randomization mechanism is controlled by individuals, the amount of randomization added to the data is also unverifiable to the platform. Therefore, the conventional wisdom of "pay according to the quality" is difficult

to implement.

- Many incentive mechanisms aim at "truthfulness" to make sure an individual has no incentive to alter her/his answer and reports the true data. This however is not necessary and sometimes should be avoided in crowd-powered systems. For example, platforms that collect personal information, such as Apple, increasingly emphasize privacy protection of their customers and prefer privacy-preserving data instead of raw personal information. Both Google and Apple have pioneered in using differential privacy in data collection [1, 5].

In light of the challenges above, this paper studies the design of incentive mechanisms aiming at obtaining answers with target quality with minimum payment, for which we need to understand *who should be paid* and *how much should be paid*? Both questions are highly nontrivial because the limited capability a platform has to assess the quality of reported answers. A flat-rate-payment mechanism, which gives a predetermined payment to each reported answer, is not cost-efficient. This is because in theory, rational individuals, who are interested in maximizing their payoffs (payoff=payment - cost) would spend zero effort and submit completely random answers under a flat-rate-payment mechanism.

This paper considers a model where a platform is interested in getting answers from N individuals for some question. Each individual has a private answer to the question, denoted by S_i. The distribution of S_i is parameterized by Θ. An individual reports X_i, where the conditional distribution of X_i given S_i is controlled by user i via ϵ_i. In the application of personal data collection, S_i is the private information and X_i is the reported data with privacy level ϵ_i. In the application of crowdsourcing, S_i is the best answer individual i can provide when the individual spends unlimited effort to solve the task; and X_i is an answer reported when ϵ_i amount of effort is used. Given this model, we first derive a lower bound on the minimum payment required given a target quality level, and then propose an incentive mechanism that matches the lower bound under some assumptions. The main results are summarized below.

1.1 Main Results

- We formulate the design of minimum payment incentive mechanisms as an optimization problem in Section 2 assuming a cost-aware platform who is interested in minimizing total payment when eliciting answers with required quality, and strategic individuals who are interested in maximizing their payoffs. In Section 3, we derive a lower bound on the minimum payment. The lower bound is derived by introducing a genie-aided mechanism, where a genie knows Θ, and pays an individual based on only her reported answer and Θ. So under the genie-aided mechanism, the minimum payment problem for the crowd becomes a minimum payment problem for an individual. We further show that any incentive mechanism that does not have access to Θ can be mimicked by the genie-aided mechanism with the same expected payment. Therefore, the minimum payment under the genie-aided mechanism is a lower bound on the original problem.

- Inspired by the lower bound, we propose WINTALL, a winners-take-all incentive mechanism in Section 3. WINTALL first estimates θ from the reported answers, denoted by $\tilde{\theta}$, from which WINTALL decides the winning answer for each individual. Under some assumptions, we prove that the expected payment under WINTALL matches the lower bound given $\tilde{\theta} = \theta$, i.e. the payment is close to the lower bound if the platform can accurately estimate θ from the reported answers with a high probability.

  To implement WINTALL, the platform needs to know the conditional distribution of X_i given S_i and the distribution of S_i given $\tilde{\theta}$, which could be difficult to obtain in practice. However, it turns out that both are not necessarily needed in some applications. In an example of discrete distribution estimation, we show that under the k-ary randomized response mechanism [6], the winning answer simply is the most popular answer among the reported answers, and the amount of payment is

$$\frac{\partial g(\epsilon)}{\partial \epsilon} \frac{(M+\epsilon)\epsilon}{M\frac{\sum_{i=1}^{N} 1_{x_i=m^*}}{N} - 1}, \tag{1}$$

  where M is the size of the alphabet (sample space), m^* is the winning answer, ϵ is the target quality level, and $g(\epsilon)$ is the cost of reporting an answer with privacy level ϵ.

1.2 Related work

Incentive mechanisms for crowd-powered systems have gained a lot of interest in recent years. A popular approach used in crowdsourcing is the peer prediction mechanisms [4,9–13,17, 18, 22], under which each individual is paired with another randomly selected individual and is paid based on how well her reported data predicts the data from her paired individual. [18] proposed an output agreement mechanism where a positive payment is made if two answers agree. [10] introduced the "Bayesian Truth Serum" (BTS) mechanism, which requires a data subject to provide her own answers as well as her belief of others' answers. A high score is given to an answer when the actual frequency is larger than the prediction. The mechanism has been further extended in various different settings [11–13, 22]. [4, 17] introduced strong truthfulness mechanisms for binary and non-binary signals in the presence of multiple questions. [14–16] developed incentive mechanisms for improving quality of labelling in crowdsourcing. The mechanisms incentivize workers to self-correct their answers in a second stage after comparing their answers with a reference answer from other workers. The goal of these mechanisms is to obtain "truthful" answers, while under our model, the platform is interested in eliciting answers with target quality instead of truthful answers.

[7] studies the problem of maximizing accuracy of crowdsouced data given a fixed budget, under which crowdsoucing tasks are assigned adaptively based on the answers collected. The paper introduced an adaptive mechanism combined with inference scheme. [21] introduced a Bayesian inference model for crowdsourcing which integrates data collection and learning. The focus of [7, 21] is on task assignment instead of incentive mechanisms. [8] developed a learning algorithm to identify low- and high-quality labelers and further used this information to improve the labelling quality in crowdsourcing. In [3], the authors considered a model under which the data collector can buy data with different variance levels with different prices. The focus is on incentivizing data providers to report their true cost functions, and it assumes the variance levels tagged are the true variance levels of the data, which is fundamental different from

our model where the variance level is unverifiable to the data collector.

[2] investigated a similar problem under a model, where the reported data is the true answer plus an additive noise with mean zero and variance as a function of effort level ϵ. This paper does not assume an additive noise model. In fact, noises introduced by many popular privacy protection mechanisms based on differential privacy are not additive. Another closely related line of work is [19, 20], which studies the market value of private data by casting the problem as eliciting private data from privacy-sensitive individuals. Both papers consider binary data, which is a special case of the model studied in this paper.

2. MODEL

Let i be the index of individuals and $S_i \in \mathcal{S}$ be the private data of individual i, where $\mathcal{S}$ is a finite set. We assume $S_i \in \mathcal{S}$ to be a random variable whose distribution $\mathbb{P}_{S_i}(s;\Theta)$ is parameterized by Θ, where Θ can be a random variable, a random vector or even a random matrix. We further assume that given $\Theta = \theta$, $\{S_i\}$ are independent. We note that if $\{S_i\}$ is an infinite sequence of exchangeable random variables, then there exists a latent random variable Θ such that $\{S_i\}$ are independent conditioned on $\Theta = \theta$ according to de Finetti's theorem. When individuals are chosen uniformly at random from a large population like in most online or offline surveys, $\{S_i\}$ are exchangeable random variables, and satisfy the assumption of this paper.

Let $X_i \in \mathcal{X}$ denote the data that individual i reports to the platform, where $\mathcal{X}$ is a finite set and may be different from $\mathcal{S}$. Furthermore, denote by $\sigma_i(\epsilon) : \mathcal{S} \to \mathcal{X}$ a data-reporting mechanism that generates reported data X_i according to the private data with quality level ϵ.

For applications which collect personal data, ϵ can be viewed as the privacy budget (a.k.a privacy loss). For example, Google RAPPOR and Apple iPhone have implemented privacy-preserving mechanisms based on differential privacy [1, 5] where ϵ is the privacy budget defined in differential privacy. In crowdsourcing case, the quality of an answer is determined by a worker's effort, so the quality level ϵ can be viewed as the amount of time (or effort) used by the worker.

We assume that the quality level ϵ uniquely determines the distribution

$$\mathbb{P}_{X_i|S_i}(x|s;\epsilon). \tag{2}$$

We assume ϵ_i is controlled by individual i and $\mathbb{P}_{X_i|S_i}(x|s;\epsilon)$ is differentiable with respect to ϵ. Note that this assumption means an individual controls $\mathbb{P}_{X_i|S_i}$ via the effort level instead of controlling it directly. For personal data collection, this assumption means the privacy preserving mechanism is fixed but an individual can select the privacy level. In crowdsourcing, it means a worker controls the quality of an answer via the amount of the effort used to finish the task. We further assume that individual i chooses the effort level before seeing her private data S_i.

We remark that it would be more general if the model allows individual i to directly choose X_i based on S_i, instead of controlling the answer via ϵ_i. However, we believe our model is more suitable for applications such as Google RAPPOR or Apple iPhone, where we can easily envision that Google or Apple in the future may let users determine the level of privacy-loss they prefer, but it is difficult to image that an individual would customize every single bit of her personal data reported to Google or Apple.

2.1 Cost-Aware Platform

We assume the platform is cost-aware and is interested in collecting data with target quality level $\boldsymbol{\epsilon}$ with minimum payment. The platform therefore uses incentive mechanism R (also called a payment mechanism) such that $R_i(\mathbf{X})$ is the payment to individual i when the reported data is $\mathbf{X}$ which is a vector such that the ith entry is the reported data of individual i, i.e. X_i. The goal of the platform is to minimize the total payment $\sum_i R_i(\mathbf{X})$ under the constraint $\epsilon_i \geq \epsilon_i^{(t)}$, where $\epsilon_i^{(t)}$ is the target quality level chosen by the platform and ϵ_i is the actual quality level of the data from individual i. In other words, the platform aims at solving the following problem:

$$\min_R \mathbb{E}\left[\sum_i R_i(\mathbf{X})\right] \tag{3}$$

$$\text{subject to:} \quad \boldsymbol{\epsilon} \geq \boldsymbol{\epsilon}^{(t)} \tag{4}$$

$$R_i(\mathbf{X}) \geq 0, \quad \forall \mathbf{X},\ \forall i, \tag{5}$$

where $\boldsymbol{\epsilon}$ is a vector such that the ith entry is ϵ_i. We impose constraint $R_i(\mathbf{X}) \geq 0$ for all $\mathbf{X}$ and all i so that negative payment (i.e. penalty) is not allowed, which is common in practice.

2.2 Strategic Individuals

We assume individuals are rational and strategic. Each individual is associated with a cost function $g_i(\epsilon)$ which is the cost incurred to individual i when the quality level is ϵ. For collecting personal data, ϵ is the cost of privacy loss; and in crowdsourcing, the cost could be the monetary loss from other sources when worker i works on finishing the task. We assume $g_i(\cdot)$ is an increasing function.

We assume individual i has the following information:

- cost function $g_i(\cdot)$,
- a (possibly biased) belief of parameter Θ, denoted by $\tilde{\Theta}_i$ and with distribution $f_{\tilde{\Theta}_i}(\theta)$,
- the payment mechanism based on reported data $R(\mathbf{X})$, which is announced by the platform,
- distribution $\mathbb{P}_{S_i}(s;\theta)$ for any given θ,
- distribution $\mathbb{P}_{S_j}(s;\theta)$ for any given θ for all $j \neq i$, and
- quality-level used by other individuals $\boldsymbol{\epsilon}_{-i}$.

Let $R_i(X_i, \mathbf{X}_{-i})$ denote the payment received by individual i given reported data $\mathbf{X}$, which is simply a different notation for $R_i(\mathbf{X})$. The payment individual i expects to receive with quality level ϵ, based on her belief $\tilde{\Theta}_i$, is

$$h_i(\epsilon) = \mathbb{E}_{\tilde{\Theta}_i, \epsilon, \boldsymbol{\epsilon}_{-i}}\left[R_i(X_i, \mathbf{X}_{-i})\right].$$

We assume individuals are strategic and are interested in maximizing the expected payoff, i.e. finding a quality level ϵ_i^* such that

$$\epsilon_i^* \in \arg\max_{\epsilon}\left(h_i(\epsilon) - g_i(\epsilon)\right). \tag{6}$$

Individuals are also rational so that they will not participate if

$$\max_{\epsilon}\left(h_i(\epsilon) - g_i(\epsilon)\right) < 0.$$

Remarks on the assumptions:

- We assume each individual has personalized $\tilde{\Theta}_i$ to model her bias (or lack of information). An individual is strategic but may not have the perfect unbiased information about Θ.

- The assumption that individual i has the information about other individuals $\boldsymbol{\epsilon}_{-i}$ and $\mathbb{P}_{S_j}(s;\theta)$ is mainly for the analysis purpose so that the expected payment to individual i, i.e.

$$\mathbb{E}_{\tilde{\theta}_i,\boldsymbol{\epsilon},\boldsymbol{\epsilon}_{-i}}\left[R_i(X_i,\mathbf{X}_{-i})\right]$$

is well defined. As we will see in the analysis, under the proposed payment mechanism and some minor assumptions, individual i does not need these two pieces of information for calculating ϵ_i^*. In other words, under the proposed incentive mechanism, individual i makes the same decision with or without knowing $\boldsymbol{\epsilon}_{-i}$ and $\mathbb{P}_{S_j}(s;\theta)$. Therefore, our main results hold even individuals do not have other individuals' information. We will see that for some application, the proposed incentive mechanism can be remarkably simple and practical with minimal information needed by both the platform and individuals. The arguably most critical assumption is that the cost function $g_i(\epsilon)$ is known to the platform. Learning and estimating the costs functions are beyond the scope of this paper.

2.3 Minimum Payment Incentive Mechanism

Summarizing the discussions in the previous two subsections, the design of a minimum cost incentive mechanism is to solve the following problem:

$$\min_R \mathbb{E}\left[\sum_i R_i(\mathbf{X})\right] \tag{7}$$

$$\text{subject to: } R_i(\mathbf{X}) \geq 0, \quad \forall \mathbf{X},\ \forall i \tag{8}$$

$$\arg\max_{\epsilon}\left(h_i(\epsilon) - g_i(\epsilon)\right) \geq \epsilon_i^{(t)} \quad \forall \tilde{\Theta}_i, \forall i. \tag{9}$$

We next comment on constraint (9), which is called **Bias-Free** condition in this paper.

- **Bias-Free:** We require condition (9) holds for all $\tilde{\Theta}_i$ because the bias $\tilde{\Theta}_i$ in general is unknown (or just partial known) to the platform. This condition guarantees that individual i chooses quality level at least $\epsilon_i^{(t)}$ regardless of her bias, which we feel is important in practice where individuals often have only limited and heterogeneous knowledge about the underlying parameter Θ.

3. A WINNERS-TAKE-ALL INCENTIVE MECHANISM

Before we present WINTALL, we first derive a lower bound on the payment to individual i with quality level ϵ. We define

$$\mathbb{P}_{X_i}(x;(\theta,\epsilon)) = \sum_{s\in\mathcal{S}} \mathbb{P}_{S_i}(s;\theta)\mathbb{P}_{X_i|S_i}(x|s;\epsilon),$$

which is the probability that individual i reports x when the underlying parameter is θ and the quality level of individual i is ϵ.

THEOREM 1. *Given any nonnegative and bias-free payment mechanism R, if ϵ is the quality-level of individual i at a Nash equilibrium under payment mechanism R, then the expected payment to individual i when $\Theta = \theta$ is lower bounded by*

$$V_i^l(\epsilon,\theta) = \frac{\partial g_i(\epsilon)}{\partial \epsilon} A_i(\epsilon,\theta), \tag{10}$$

where

$$A_i(\epsilon,\theta) = \min_{x\in\mathcal{X}}\left\{\frac{\mathbb{P}_{X_i}(x;(\theta,\epsilon))}{\frac{\partial \mathbb{P}_{X_i}(x;(\theta,\epsilon))}{\partial\epsilon}} : \frac{\partial \mathbb{P}_{X_i}(x;(\theta,\epsilon))}{\partial\epsilon} > 0\right\}. \tag{11}$$

□

The proof of this theorem can be found in our technical report.

Given the lower bound, the question now is whether the lower bound can be achieved? To answer this question, we note that by defining

$$x^*_{i,\theta,\epsilon_i} \in \arg\min_{x\in\mathcal{X}}\left\{\frac{\mathbb{P}_{X_i}(x;(\theta,\epsilon_i))}{\frac{\partial \mathbb{P}_{X_i}(x;(\theta,\epsilon_i))}{\partial\epsilon}} : \frac{\partial \mathbb{P}_{X_i}(x;(\theta,\epsilon_i))}{\partial\epsilon} > 0\right\}, \tag{12}$$

the lower bound given $\Theta = \theta$ is

$$\frac{\partial g_i(\epsilon_i)}{\partial\epsilon}\frac{1}{\frac{\partial \mathbb{P}_{X_i}(x^*_{i,\theta,\epsilon_i};(\theta,\epsilon_i))}{\partial\epsilon}}\mathbb{P}_{X_i}(x^*_{i,\theta,\epsilon_i};(\theta,\epsilon_i)), \tag{13}$$

which suggests that the lower bound can be achieved by paying individual i only when she reports $x^*_{i,\theta,\epsilon_i}$ with a payment of $\frac{\partial g_i(\epsilon_i)}{\partial\epsilon}\frac{1}{\frac{\partial \mathbb{P}_{X_i}(x^*_{i,\theta,\epsilon_i};(\theta,\epsilon_i))}{\partial\epsilon}}$. Therefore, we propose the following incentive mechanism.

A Winners-Take-All Incentive Mechanism (WINTALL)

(1) The platform announces target quality level $\boldsymbol{\epsilon}^{(t)}$.

(2) Each individual reports her data (which can also be an decision of not participating).

(3) For non-participating individual, the payment is zero.

(4) Given the collected data $\mathbf{X}$, the platform estimates θ, denoted by $\tilde{\theta}$.

(5) For each participating individual i, the platform pays according to the reported X_i, the estimation $\tilde{\theta}$, and the target quality level $\epsilon_i^{(t)}$. Specifically, if the reported data is $x^*_{i,\tilde{\theta},\epsilon_i^{(t)}}$, individual i receives a payment of

$$W_{i,\epsilon_i^{(t)},\tilde{\theta}} = \frac{\partial g_i(\epsilon_i^{(t)})}{\partial\epsilon}\frac{1}{\frac{\partial \mathbb{P}_{X_i}(x^*_{i,\tilde{\theta},\epsilon_i^{(t)}};(\tilde{\theta},\epsilon_i^{(t)}))}{\partial\epsilon}},$$

where

$$x^*_{i,\tilde{\theta},\epsilon_i^{(t)}} \in \arg\min_{x\in\mathcal{X}}\left\{\frac{\mathbb{P}_{X_i}(x;(\tilde{\theta},\epsilon_i^{(t)}))}{\frac{\partial \mathbb{P}_{X_i}(x;(\tilde{\theta},\epsilon_i^{(t)}))}{\partial\epsilon}} : \frac{\partial \mathbb{P}_{X_i}(x;(\tilde{\theta},\epsilon_i^{(t)}))}{\partial\epsilon} > 0\right\}$$

otherwise, no payment is made to individual i. □

Now to understand whether the proposed incentive mechanism actually achieves the lower bound, we first prove the following theorem, which shows that $\boldsymbol{\epsilon}^{(t)}$ is a bias-free Nash equilibrium under WINTALL. Note that under WINTALL,

given the estimation $\tilde{\theta}$, the payment to individual i is independent of other individuals' reports. Furthermore individual i needs to decide on the quality level ϵ_i before they receive their private data. Therefore, we assume that individual i is confident about her belief and uses its belief $\tilde{\Theta}_i$ when choosing the quality level ϵ_i.

THEOREM 2. *If for any i,*

$$W_{i,\epsilon_i^{(t)},\theta}\mathbb{P}_{X_i}\left(x^*_{i,\epsilon_i^{(t)},\theta};(\theta,\epsilon)\right) - g_i(\epsilon)$$

is strictly concave in ϵ for given $\boldsymbol{\epsilon}^{(t)}$ and any θ, then the target quality level $\boldsymbol{\epsilon}^{(t)}$ is a bias-free Nash equilibrium under WINTALL. With quality level $\epsilon_i^{(t)}$, the expected payment individual i receives is

$$W_{i,\epsilon_i^{(t)},\tilde{\theta}}\mathbb{P}_{X_i}\left(x^*_{i,\tilde{\theta},\epsilon_i^{(t)}};(\tilde{\theta},\epsilon_i^{(t)})\right), \tag{14}$$

which equals to $V_i^l\left(\epsilon_i^{(t)},\theta\right)$ when $\tilde{\theta}=\theta$ (i.e. the platform can accurately estimate Θ from the collected data). □

The proof of this theorem can be found in our technical report.

"Winners-Take-All": Suppose $\{S_i\}$ are identically distributed and the target quality level is the same for all individuals. In this case, $x^*_{i,\tilde{\theta},\epsilon_i^{(t)}}$ is independent of i and can be written as $x^*_{\tilde{\theta},\epsilon}$. Therefore, only individuals who report $x^*_{\tilde{\theta},\epsilon}$ will be paid. In other words, under WINTALL, after collecting all data, the platform determines a "winning" report $x^*_{\tilde{\theta},\epsilon}$ and all payments go to the "winners".

Remark: The theorem above shows that the expected payment under WINTALL matches the lower bound when the underlying parameter θ can be estimated accurately. In practice, we expect that the platform can estimate θ accurately after it collects data from many individuals. So we can expect WINTALL achieves the minimum (or near minimum) payment in large-scale crowd-powered systems.

4. APPLICATIONS

An important application of the proposed model and incentive mechanism is private discrete distribution estimation, widely used in usage statistics breakdowns and count-based machine learning models. Let θ denote an M-dimensional probability distribution and define S_i to be a random variable such that $\Pr(S_i = m) = \theta_m$. Here S_i can represent whether user i has visited website m, or her opinion about a certain subject. The platform collects data from N individuals randomly selected from the crowd to estimate θ. As we mentioned earlier, when the individuals are selected uniformly at random, $\{S_i\}$ are exchangeable. A pictorial illustration is shown in Figure 1.

We assume individuals use the following privacy-preserving mechanism

$$\mathbb{P}_{X|S}(k|m;\epsilon) = \begin{cases} \frac{\epsilon+1}{\epsilon+M}, & k=m \\ \frac{1}{\epsilon+M}, & k\neq m \end{cases}.$$

This is the k-ary randomized response mechanism proposed in [6] for discrete distribution estimation which guarantees differential privacy budget

$$\epsilon^{(d)} = \log(\epsilon+1)$$

and is proved to be optimal in the low-privacy regime [6] .

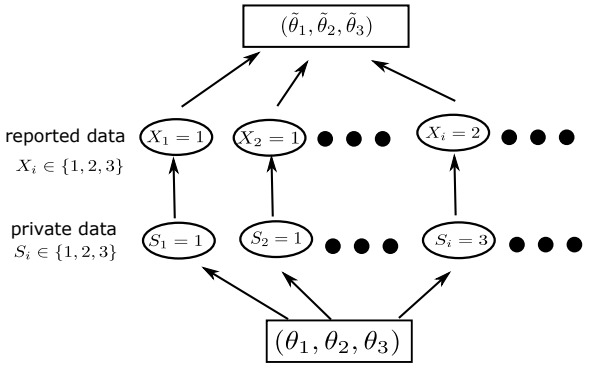

Figure 1: Illustration of Discrete Distribution Estimation

Given the k-ary randomized response mechanism, we have

$$\begin{aligned}\mathbb{P}_X(m;(\epsilon;\theta)) &= \sum_{k=1}^{M}\mathbb{P}_{X|S}(m|k;(\epsilon,\theta))\mathbb{P}_S(k;(\epsilon,\theta)) \\ &= \theta_m\frac{\epsilon+1}{\epsilon+M} + (1-\theta_m)\frac{1}{\epsilon+M} \\ &= \frac{\theta_m\epsilon+1}{\epsilon+M},\end{aligned}$$

and

$$\frac{\partial}{\partial\epsilon}\mathbb{P}_X(m;(\epsilon;\theta)) = -\frac{\theta_m\epsilon+1}{(\epsilon+M)^2} + \frac{\theta_m}{\epsilon+M} = \frac{\theta_m M-1}{(\epsilon+M)^2}.$$

Therefore, we have

$$\begin{aligned}\frac{\mathbb{P}_X(m;(\epsilon;\theta))}{\frac{\partial}{\partial\epsilon}\mathbb{P}_X(m;(\epsilon;\theta))} &= \frac{\epsilon\theta_m+1}{\epsilon+M}\frac{(\epsilon+M)^2}{M\theta_m-1} \\ &= (\epsilon+M)\frac{\epsilon}{M}\left(1+\frac{\frac{1}{\epsilon}+\frac{1}{M}}{\theta_m-\frac{1}{M}}\right),\end{aligned}$$

which is a decreasing function of θ_m. From that, we conclude that

$$m^* = \arg\min_m \frac{\mathbb{P}_X(m;(\epsilon;\theta))}{\frac{\partial}{\partial\epsilon}\mathbb{P}_X(m;(\epsilon;\theta))} = \arg\max_m \theta_m.$$

Note that unless $\theta_m = \frac{1}{M}$ for all m, i.e. uniform distribution, we have $\theta_{m^*} > \frac{1}{M}$, which implies that

$$\frac{\partial}{\partial\epsilon}\mathbb{P}_X(m^*;(\epsilon;\theta)) = \frac{\theta_{m^*}M-1}{(\epsilon+M)^2} > 0,$$

and $\mathbb{P}_X(m^*;(\epsilon;\theta))$ is strictly concave in ϵ because $\frac{\partial}{\partial\epsilon}\mathbb{P}_X(m^*;(\epsilon;\theta))$ is a decreasing function in ϵ. We also note that for any $\epsilon>0$,

$$m^* = \arg\max_m \theta_m = \arg\max_m \mathbb{P}_X(m;(\theta,\epsilon)).$$

In other words, the most popular answers in the private data and in the reported data are the same.

WINTALL with target quality level ϵ in this case is as follows.

WINTALL for Private Discrete Distribution Estimation

- After collecting data from N individuals, denoted by $\{x_i\}_{i=1,\cdots,N}$, the platform identifies the most popular answer m^* :

$$m^* \in \arg\max_m \frac{\sum_{i=1}^{N} 1_{x_i=m}}{N}.$$

Ties are broken uniformly at random.

- The platform pays each user who reports m^* an amount of

$$\frac{\partial g(\epsilon)}{\partial \epsilon} \frac{(M+\epsilon)\epsilon}{M\frac{\sum_{i=1}^{N} 1_{x_i=m^*}}{N} - 1}. \quad (15)$$

□

Remark: We note that the most popular answer in the private data is consistent with that in the reported data, which creates the incentive for an individual to report an answer close to her private data because the individual expects her private answer to be the dominating one. Note that in this example, the only prior information the platform needs is $\frac{\partial g(\epsilon)}{\partial \epsilon}$.

5. CONCLUSIONS

This paper studied incentive mechanisms for crowd-powered systems. We first derived a lower bound on the minimum payment required for guaranteeing quality level, and then proposed WINTALL — a novel incentive mechanism. The expected payment of WINTALL matches the lower bound when the underlying parameter θ can be estimated by the platform accurately. We present its application to private discrete distribution estimation, where WINTALL rewards individuals whose reported answers match the most popular one.

Acknowledgement

This work was supported in part by the NSF under Grants ECCS-1547294, CNS-1739344, IIS-1552654, and IIS-1813464.

6. REFERENCES

[1] https://images.apple.com/privacy/docs/Differential_Privacy_Overview.pdf.

[2] Cai, Y., Daskalakis, C., and Papadimitriou, C. Optimum statistical estimation with strategic data sources. In *Proc. Conf. Learning Theory (COLT)* (2015), pp. 280–296.

[3] Cummings, R., Ligett, K., Roth, A., Wu, Z. S., and Ziani, J. Accuracy for sale: Aggregating data with a variance constraint. In *Proc. Conf. Innovations in Theoretical Computer Science* (2015), pp. 317–324.

[4] Dasgupta, A., and Ghosh, A. Crowdsourced judgement elicitation with endogenous proficiency. In *Proc. Int. Conf. World Wide Web (WWW)* (2013), pp. 319–330.

[5] Erlingsson, Ú., Pihur, V., and Korolova, A. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proc. ACM Conf. Computer and Communications Security (CCS)* (Scottsdale, AZ, 2014), pp. 1054–1067.

[6] Kairouz, P., Bonawitz, K., and Ramage, D. Discrete distribution estimation under local privacy. In *Int. Conf. Machine Learning (ICML)* (2016), pp. 2436–2444.

[7] Khetan, A., and Oh, S. Achieving budget-optimality with adaptive schemes in crowdsourcing. In *Advances Neural Information Processing Systems (NIPS)* (2016), pp. 4844–4852.

[8] Liu, Y., and Liu, M. An online learning approach to improving the quality of crowd-sourcing. In *Proc. Ann. ACM SIGMETRICS Conf.* (New York, NY, USA, 2015), pp. 217–230.

[9] Miller, N., Resnick, P., and Zeckhauser, R. Eliciting informative feedback: The peer-prediction method. In *Computing with Social Trust*, Human–Computer Interaction Series. Springer London, 2009, pp. 185–212.

[10] Prelec, D. A bayesian truth serum for subjective data. *Science 306*, 5695 (2004), 462–466.

[11] Radanovic, G., and Faltings, B. A robust bayesian truth serum for non-binary signals. In *AAAI Conf. Artificial Intelligence* (2013), pp. 833–839.

[12] Radanovic, G., and Faltings, B. Incentives for truthful information elicitation of continuous signals. In *AAAI Conf. Artificial Intelligence* (2014), pp. 770–776.

[13] Radanovic, G., and Faltings, B. Incentive schemes for participatory sensing. In *Proc. Int. Conf. Autonomous Agents and Multiagent Systems* (2015), International Foundation for Autonomous Agents and Multiagent Systems, pp. 1081–1089.

[14] Shah, N., and Zhou, D. No oops, you won't do it again: Mechanisms for self-correction in crowdsourcing. In *International Conference on Machine Learning* (2016), pp. 1–10.

[15] Shah, N., Zhou, D., and Peres, Y. Approval voting and incentives in crowdsourcing. In *International Conference on Machine Learning* (2015), pp. 10–19.

[16] Shah, N. B., and Zhou, D. Double or nothing: Multiplicative incentive mechanisms for crowdsourcing. In *Advances in neural information processing systems* (2015), pp. 1–9.

[17] Shnayder, V., Agarwal, A., Frongillo, R., and Parkes, D. C. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation* (2016), pp. 179–196.

[18] Von Ahn, L., and Dabbish, L. Labeling images with a computer game. In *Proc. the SIGCHI conf. Human factors in computing systems* (2004), pp. 319–326.

[19] Wang, W., Ying, L., and Zhang, J. A game-theoretic approach to quality control for collecting privacy-preserving data. In *Proc. Annu. Allerton Conf. Communication, Control and Computing* (Monticello, IL, Sept. 2015).

[20] Wang, W., Ying, L., and Zhang, J. The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits. In *Proc. Ann. ACM SIGMETRICS Conf.* (Antibes Juan-les-Pins, France, June 2016).

[21] Wauthier, F. L., and Jordan, M. I. Bayesian bias mitigation for crowdsourcing. In *Advances Neural Information Processing Systems (NIPS)* (2011), pp. 1800–1808.

[22] Witkowski, J., and Parkes, D. C. A robust bayesian truth serum for small populations. In *AAAI Conf. Artificial Intelligence* (2012), vol. 12, pp. 1492–1498.