

Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices

Yang Liu, Yiling Chen
Harvard University, Cambridge MA, USA
{yangl,yiling}@seas.harvard.edu

Abstract

Peer prediction mechanisms are often adopted to elicit truthful contributions from crowd workers when no ground-truth verification is available. Recently, mechanisms of this type have been developed to incentivize effort exertion, in addition to truthful elicitation. In this paper, we study a sequential peer prediction problem where a data requester wants to dynamically determine the reward level to optimize the trade-off between the quality of information elicited from workers and the total expected payment. In this problem, workers have homogeneous expertise and heterogeneous cost for exerting effort, both unknown to the requester. We propose a sequential posted-price mechanism to dynamically learn the optimal reward level from workers' contributions and to incentivize effort exertion and truthful reporting. We show that (1) in our mechanism, workers exerting effort according to a non-degenerate threshold policy and then reporting truthfully is an equilibrium that returns highest utility for every worker, and (2) The regret of our learning mechanism w.r.t. offering the optimal reward (price) is upper bounded by $\tilde{O}(T^{3/4})$ where T is the learning horizon. We further show the power of our learning approach when the reports of workers do not necessarily follow the game-theoretic equilibrium.

Introduction

Crowdsourcing has arisen as a promising option to facilitate machine learning via eliciting useful information from human workers. For example, such a notion has been widely used for labeling training samples, e.g., Amazon Mechanical Turk. Despite its simplicity and popularity, one salient feature or challenge of crowdsourcing is the lack of evaluation for the collected answers, because ground-truth labels often are either unavailable or too costly to obtain. This problem is called *information elicitation without verification* (IEWV) (Waggoner and Chen 2014). A class of mechanisms, collectively called *peer prediction*, has been developed for the IEWV problem (Prelec 2004; Miller, Resnick, and Zeckhauser 2005; Jurca and Faltings 2007; 2009; Witkowski and Parkes 2012a; 2012b; Radanovic and Faltings 2013). In peer prediction, an agent is rewarded according to how his answer compares with those of his peers and the reward rules are designed so that everyone truthfully reporting their information is a game-theoretic equilibrium.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

More recent work (Witkowski et al. 2013; Dasgupta and Ghosh 2013; Shnayder et al. 2016) on peer prediction concerns *effort elicitation*, where the goal is not only to induce truthful report, but also to induce high quality answers by incentivizing agents to exert effort. In such work, the mechanism designer is assumed to know workers' expertise level and their cost for effort exertion and designs reward rules to induce optimal effort levels and truthful reporting at an equilibrium.

This paper also focuses on the effort elicitation of peer prediction. But different from prior work, our mechanism designer knows neither workers' expertise level nor their cost for effort exertion. We introduce a *sequential peer prediction* problem, where the mechanism proceeds in rounds and the mechanism designer wants to learn to set the optimal reward level (that balances the amount of effort elicited and the total payment) while observing the elicited answers in previous rounds. There are several challenges to this problem. First, effort exertion is not observable and no ground-truth answers are available for evaluating contributions. Hence, it is not immediately clear what information the mechanism designer can learn from the observed answers in a sequential mechanism. Second, forward-looking workers may have incentives to mislead the learning process, hoping for better future returns.

The main contributions of this paper are the following: (1) We propose a *sequential peer prediction* mechanism by combining ideas from peer prediction with multi-armed bandit learning (Lai and Robbins 1985; Auer, Cesa-Bianchi, and Fischer 2002). (2) In this mechanism, workers exerting effort according to a non-degenerate threshold policy and then reporting truthfully in each round is an equilibrium that returns highest utility for every worker. (3) We show that the regret of this mechanism w.r.t. offering the optimal reward is upper bounded by $\tilde{O}(T^{3/4})$ where T is the learning horizon. We also show that under a "mean-field" assumption, the sequential learning mechanism can be extended to a setting where workers may not be fully rational. (4) Our sequential peer prediction mechanism is *minimal* in that reported labels are the only information we need from workers.

In the rest of the paper, we first survey the most related work in Section . Section introduces our problem formulation. We then present a game-theoretic analysis of worker behavior in a one-stage static setting in Section . Based on

the equilibrium analysis of the one-stage setting, we propose and analyze a learning mechanism to learn the optimal bonus level using posted price in Section . We also discuss an extension of our learning mechanism to a setting where workers may not be fully rational. Section concludes this paper. All omitted details can be found in the full version of the paper (Liu and Chen 2016b).

Related work

Eliciting high-quality data from effort-sensitive workers hasn't been addressed within the literature of peer prediction until recently. Witkowski et al. [2013] and Dasgupta and Ghosh [2013] formally introduced costly effort into models of IEWV. The costs for effort exertion were assumed to be homogeneous and known and static, one-shot mechanisms were developed for effort elicitation and truthful reporting. Our setting allows participants to have heterogeneous cost of effort exertion drawn from a common unknown distribution and hence we consider a sequential setting that enables learning over time. Liu and Chen [2016a] is the closest to this work. It considered the same general setting and partially resolved the problem of learning the optimal reward level sequentially. There are two major differences however. First, the method developed in Liu and Chen [2016a] required workers to report their private cost in addition to their answer, which is arguably undesirable for practical applications. Our learning mechanism in contrast is "minimal" (Segal 2007; Witkowski and Parkes 2013) and only asks for answers (for tasks) from workers. Second, the mechanism of Liu and Chen [2016a] was built upon the output agreement mechanism as the single-round mechanism. Output agreement and hence the mechanism of Liu and Chen [2016a] suffer from potential, simple collusions of workers: colluding by reporting an uninformative signal will lead to a better equilibrium (higher utility) for workers. By building upon the mechanism of Dasgupta and Ghosh [2013], which as a one-shot mechanism is resistant to such simple collusion, we develop a collusion-resistant sequential learning mechanism.

Generally speaking, when there is a lack of knowledge of agents, the design problem needs to incorporate learning from prior outputs from running the mechanism – see Chawla, Hartline, and Nikipelov [2014] for specific examples on learning with auction data. And this particular topic has also been studied within the domain of crowdsourcing. For example, Roth and Schoenebeck [2012] and Abernethy et al. [2015] consider strategic data acquisition for estimating the mean and for online learning respectively. Our problem differs from above in that both agents' action (effort exertion) and ground-truth outcomes are unavailable.

Problem Formulation

Formulation and settings

Suppose in our system we have one data requester (or a mechanism designer), and there are N candidate workers denoted by $\mathcal{C} = \{1, 2, \dots, N\}$, where $N \geq 4$. In all we have $N + 1$ interactive agents. The data requester has binary-answer tasks, with answer space $\{-1, +1\}$, that she'd like to

get labels for. The requester assigns tasks to workers.

Label generated by worker $i \in \mathcal{C}$ comes from a distribution that depends both on the ground-truth label and an effort variable e_i . Suppose there are two effort levels, High and Low, that a worker can potentially choose from: $e_i \in \{H, L\}$. We model the cost c for exerting High effort for each (worker, task) pair as drawn from a distribution with c.d.f. $F(c)$ on a bounded support $[0, c_{\max}]$; while exerting Low effort incurs no cost. We assume such costs are drawn in an i.i.d. fashion. Denote worker i 's probability of observing $s' \in \{-1, +1\}$ when the ground-truth label is $s \in \{-1, +1\}$ as $p_{i,e_i} = \Pr(s' = s | s, e_i)$, under effort level e_i . Note with above we have assumed that the labeling accuracy is symmetric, and is independent of the ground-truth label s . Further for simplicity of analysis, we will assume all workers have the same set of $p_{i,H}, p_{i,L}$, denoting as p_H, p_L . With higher effort, the expertise level is higher: $1 \geq p_H > p_L \geq 0.5$ – we also assume the labeling accuracy is no less than 0.5. The above are common knowledge among workers, while the mechanism designer doesn't know the form of $F(\cdot)$; neither does she know p_H, p_L . But we assume the mechanism designer knows the structural information, such as costs are i.i.d., workers are effort sensitive, and there are two effort levels etc.

The goal of the learner is to design a sequential peer prediction mechanism for effort elicitation via observing contributed data from workers, such that the mechanism will help the learner converge to making the optimal action (will be defined later).

Reward mechanism

Once assigned a task, worker i has a guaranteed base payment $b > 0$ for each task he completes. In addition to the base payment, the worker receives a reward $B_i(k)$ for task k that he has provided an answer for. The reward is determined using the mechanism proposed by Dasgupta and Ghosh [2013], where in this paper, we denote this specific peer prediction mechanism for effort elicitation as (DG13). In this mechanism, for each (worker, task) pair (i, k) , first a reference worker $j \neq i$ is selected randomly from the set of workers who are also assigned task k . Suppose any pair of such workers have been assigned d other distinct tasks $\{i_1, \dots, i_d\}$ and $\{j_1, \dots, j_d\}$ respectively. Then the mechanism pays $B'_i(k)$ to worker i on task k in the following way: the payment consists of two terms; one term that rewards agreement on task k , and another that penalizes on uninformative agreement on other tasks:

$$B'_i(k) = 1 \left(L_i(k) = L_j(k) \right) - L_i^d \cdot L_j^d - \bar{L}_i^d \cdot \bar{L}_i^d, \quad (1)$$

where we denote reports from worker i on task n as $L_i(n)$ and $L_i^d := \sum_{n=1}^d (1 + L_i(i_n)) / (2d)$, $\bar{L}_i^d = 1 - L_i^d$. Our bonus rule follows exactly the same idea except that we will multiply $B'_i(k)$ by a constant $B \in [0, \bar{B}]$ (which we can choose): $B_i(k) := B \cdot B'_i(k)$.

Task assignment: We'd like all workers to work on the same number of tasks, all tasks are assigned to the same

number of workers and any pair of workers have distinct tasks. In particular, each worker is assigned $M > 1$ tasks – denote the set of tasks assigned to worker i as $\mathcal{T}_i : |\mathcal{T}_i| = M$. This is to simplify the computation of workers’ utility and payment functions. Each task is assigned to at least two workers. For any pair of workers who has been assigned a common task, they have at least $1 \leq d < M$ distinct tasks. These are to ensure that the (DG13) mechanism is applicable. We also set the number of assignments for each task to be the same – denote this number as $1 \leq K < N$, so that when we evaluate the accuracy of aggregated labels later, all tasks receive the same level of accuracy. But note we do not assign all tasks to all workers, i.e., $K \neq N$. Suppose each assigned task k appears in $D \leq d$ tasks’ distinct set for each worker. The described assignment can be achieved by assigning N different tasks in each round. For more details on assignments please refer to our full version.

Worker model

After receiving each task k , worker i first realizes the cost $c_i(k)$ for exerting High effort. Then worker i decides his effort level $e_i(k) \in \{H, L\}$ and observes a signal $L_i(k)$ (label of the task). Worker i can decide either to truthfully report his observation $r_i(k) = 1$ (denote by $r_i(k)$ the decision variable on reporting) or to revert the answer $r_i(k) = 0$:

$$L_i^r(k) = \begin{cases} L_i(k), & \text{if } r_i(k) = 1 \\ -L_i(k), & \text{if } r_i(k) = 0 \end{cases}$$

Workers are utility maximizers. Denote the utility function at each time (or step) for each worker as u_i , which is assumed to have the following form (payment – cost):

$$u_i = Mb + \sum_{k \in \mathcal{T}_i} B_i(k) - \sum_{k \in \mathcal{T}_i} c_i(k), \quad \forall i.$$

Data requester model

After collecting labels for each task, the data requester will aggregate labels via majority voting. Denote workers who labeled task k as $w_k(1), \dots, w_k(K)$. Then the aggregate label for k is given by

$$L^A(k) = 1 \left(\sum_{n=1}^K L_{w_k(n)}^r(k) / K > 0 \right) \cdot 2 - 1.$$

The data requester’s objective is to find a bonus level B (as in $B_i(k)$) that balances the accuracy of labels collected from workers, and the total payment. Denote requester’s objective function at each step as $\mathcal{U}(B)$ (assigning N tasks):

$$\mathcal{U}(B) := \sum_{k=1}^N \left[\Pr[L^A(k) = L(k)] - \eta \sum_{n=1}^K \mathbb{E}[B_{w_k(n)}(k)] \right],$$

where $L(k)$ denotes the true label of task k , and $\eta > 0$ is a weighting constant balancing the two terms in the objective.

Since we have assumed that all tasks have been assigned the same number of workers, and workers are homogeneous in their labeling accuracy and cost (i.i.d.), we know all tasks enjoy the same probability of having a correct label (a-priori). We denote this probability as $\mathcal{P}(B) := \Pr[L^A(k) =$

$L(k)], \forall k$. Further as workers do not receive payment when a task is not assigned to him, $\mathcal{U}(B)$ can be simplified (and normalized¹) to the following form:

$$\mathcal{U}(B) = \mathcal{P}(B) - \frac{\eta}{N} \sum_{i \in \mathcal{C}} \sum_{k=1}^N \mathbb{E}[B_i(k)], \quad (2)$$

Suppose there exists a maximizer $B^* = \operatorname{argmax}_B \mathcal{U}(B)$.

Sequential learning setting

Suppose our sequential learning algorithm goes for T stages. At each stage $t = 1, \dots, T$, learner assigns a certain number of tasks $M_i(t)$ to a set of selected workers $i \in S(t)$ ². The learner offers a bonus bundle $B_{i,t}$ to each worker $i \in S(t)$ (the bonus constant in reward mechanism). The regret of offering $\{B_{i,t}\}_{i,t}$ w.r.t. B^* is defined as follows:

$$R(T) = T \cdot \mathcal{U}(B^*) - \sum_{t=1}^T \sum_{i \in S(t)} \frac{M_i(t) \cdot \mathbb{E}[\mathcal{U}(B_{i,t})]}{\sum_{j \in S(t)} M_j(t)}. \quad (3)$$

Note we normalize $\mathbb{E}[\mathcal{U}(B_{i,t})]$ using the number of assignments – intuitively the more the requester assigned with a wrong price, the more regret will be incurred. The goal of the data requester is to design an algorithm such that $R(T) = o(T)$. We can also define $R(T)$ as being unnormalized, which will add a constant (bounded number of assignments at each step) in front of our results later.

One stage game-theoretic analysis

From the data requester’s perspective, we need to first understand workers’ actions towards effort exertion and reporting under different bonus levels, in order to figure out the optimal B^* . We start with the case that the data requester knows the cost distribution, and we characterize the equilibria for effort exertion and reporting, i.e. $(\mathbf{e}_i, \mathbf{r}_i)$, on workers’ side. Note $\mathbf{e}_i, \mathbf{r}_i$ are both vectors defined over all tasks – this is a simplification of notation as workers do not receive all tasks. We are safe as if task k is not assigned to i , worker i does not make decisions on $(e_i(k), r_i(k))$. We define Bayesian Nash Equilibrium (BNE) in our context as follows:

Definition 1. We say $\{(\mathbf{e}_i^*, \mathbf{r}_i^*)\}_{i \in \mathcal{C}}$ is a BNE if $\forall j, (\mathbf{e}_j, \mathbf{r}_j)$:

$$\mathbb{E}[u_j | \{(\mathbf{e}_i^*, \mathbf{r}_i^*)\}_{i \in \mathcal{C}}] \geq \mathbb{E}[u_j | \{(\mathbf{e}_i^*, \mathbf{r}_i^*)\}_{i \neq j}, (\mathbf{e}_j, \mathbf{r}_j)].$$

In this paper, we restrict our attention to symmetric BNE. For each assigned task, we have a Bayesian game among workers in \mathcal{C} : a worker’s decision on effort exertion is a function of $\mathbf{c}_i, \mathbf{e}_i(\mathbf{c}_i) : [0, c_{\max}]^M \rightarrow \{H, L\}^M$, which specifies the effort levels for worker i when his realized cost is \mathbf{c}_i and $\mathbf{r}_i(\mathbf{e}_i) : \{0, 1\}^M \rightarrow \{0, 1\}^M$ gives the reporting strategy for the chosen effort level. We focus on threshold policies: that is, there is a threshold c^* such that $e_i(k) = H$ for all $c_i(k) \leq c^*$ and $e_i(k) = L$ otherwise. In fact, players must play a threshold strategy for effort exertion at any symmetric BNE: workers’ outputs do not depend on $c_i(k)$ and worker i ’s chance of getting a bonus will not change

¹which does not affect optimizing the utility function.

²For details please refer to our algorithm.

when he has a cost $c'_i(k) < c_i(k)$; so a worker will choose to exert effort, if it is a better move for an even higher cost. We will use $(c^*, r_i(k))$ to denote this threshold (c^*) strategy for workers. Denote $r_i(\cdot) \equiv 1$ the reporting strategy that $r_i(H) = r_i(L) = 1$, i.e. reporting truthfully regardless of the choice of effort.

Theorem 2. *When $p_L > 0.5$ and $F(c)$ is concave, there exists a unique threshold $c^*(B) > 0$ for $B > 0$ such that $(c^*(B), 1)$ is a symmetric BNE for all workers on all tasks.*

Other equilibria: The above threshold policy is unique only in non-degenerate effort exertion ($c^* > 0$). There exist other equilibria. We summarize them here:

- *Un-informative equilibrium:* Colluding by always reporting the same answer to all tasks is an equilibrium. Similarly as mentioned in (Dasgupta and Ghosh 2013), when colluding (pure or mixed strategies) the bonus index defined in Eqn. (1) reduces to 0 for each worker, which leads to a worse equilibrium.
- *Low effort:* When $p_L = 1/2$, $c^* = 0$, i.e., no effort exertion (followed by either truthful or untruthful reporting) is also an equilibrium: when no one else is exerting effort, each worker's answer will be compared to a random guess. So there would be no incentive for effort exertion.
- *Permutation:* Exerting the same amount of effort and then reverting the reports ($r_i \equiv 0$) is also an equilibrium.

But we would like to note that though there may exist multiple equilibria, all others lead to strictly less utility for each worker at equilibrium compared to the threshold equilibrium with $c^* > 0$ followed by truthful reporting, except for the permutation equilibria, which gives the same expected utility to workers.

Solve for optimal B^ :* After characterizing the equilibria c^* on effort exertion as a function of B , we can compute $\mathcal{P}(B)$ and $\mathbb{E}[B_i(k)]$ for each reward level B . Then solving for the optimal reward level becomes a programming problem in B , which can be solved efficiently when certain properties, e.g. convexity, can be established for $\mathcal{U}(\cdot)$.

Sequential Peer Prediction

In this section we propose an adaptive learning mechanism to learn to converge to the optimal or nearly optimal reward level. As mentioned earlier, a recent work (Liu and Chen 2016a) attempted to resolve this challenge. But besides the output labels, workers are also required to report the private costs, in which sense the proposed learning mechanism is not "minimal". We try to remove this requirement by learning only through the label information reported by the workers. In this section, we assume the requirements for Theorem 2 hold, and workers will follow an equilibrium that returns the highest utility.

Challenges

In designing the mechanism, we face two main challenges. The first challenge is on the learning part. In order to select the best B^* , we need to compute $\mathcal{U}(B), \forall B$, which can be

computed as a function of B and $\bar{p}(B)$, the probability of labeling accurately when B is offered and the threshold policy $c^*(B)$ is adopted by workers:

$$\bar{p}(B) := F(c^*(B))p_H + [1 - F(c^*(B))]p_L. \quad (4)$$

The dependency on B is straight-forward. For $\bar{p}(B)$, e.g. when using Chernoff bound for approximating $\mathcal{P}(B)$:

$$\begin{aligned} \mathcal{P}(B) &= \Pr \left[\frac{\sum_i 1(\text{worker } i \text{ is correct})}{M} \geq 0.5 \right] \\ &\geq 1 - \exp(-2(\bar{p}(B) - 0.5)^2 M), \end{aligned}$$

it is clear $\mathcal{P}(B)$ is a function of $\bar{p}(B)$. In fact both $\mathbb{E}[B_i(k)]$ and $\mathcal{P}(B)$ are functions of $\bar{p}(B)$, so is $\mathcal{U}(\cdot)$. For details, please see Appendix of (Liu and Chen 2016b). The question pings down to learn $\bar{p}(B)$. Since we do not have the ground-truth labels, we have no way to directly evaluate $\bar{p}(B)$ via checking workers' answers. Also since we do not elicit reports on private costs, we are un-able to estimate the amount of induced effort for each reward level.

The second challenge we have is that when workers are returning and participating in a/an sequential/adaptive learning mechanism, they have incentives to mislead the learning process by deviating from the one-shot BNE strategy for a task, so to create untruthful samples (and then collected by learner), which will lead the learner into believing that inducing certain amount of effort requires a much higher reward level. The cost-reporting mechanism described in (Liu and Chen 2016a) proposes a method to deter such a deviation by eliminating workers who over-reported from receiving potentially higher bonus. We will describe a two-fold cross validation approach to decouple such incentives, which aims to remove the requirement of reporting additional information.

Learning w/o ground-truth The following observation inspires our method for learning without ground-truth. For each bonus level B , we can estimate $\bar{p}(B)$ (at equilibrium) through the following experiments: the probability of observing a pair of matching answers for any pair of workers i, j (denoted by $p_m(B)$ for each bonus level B) on equilibrium can be written as follows:

$$p_m(B) = \underbrace{\bar{p}^2(B)}_{\text{match on correct label}} + \underbrace{(1 - \bar{p}(B))^2}_{\text{match on wrong label}}. \quad (5)$$

The above matching formula forms a quadratic equation of $\bar{p}(B)$. From Eqn. (4) we know $\bar{p}(B) \geq 0.5, \forall B$, when $p_H, p_L \geq 0.5$. Then the only solution to the matching Eqn. (5) that is larger than 0.5 is

$$\bar{p}(B) = 1/2 + \sqrt{2p_m(B) - 1}/2.$$

Above solution is well defined, as from Eqn. (5) we can also deduce that $p_m(B) \geq 1/2$. Therefore though we cannot evaluate each worker's labeling accuracy directly, we can make such an inference using the matching probability, which is indeed observable.

Decoupling incentives via cross validation To solve the incentive issue, we propose the following cross validation approach (illustrated in Fig. 1). First the entire crowd \mathcal{C} is separated into two groups G_1, G_{-1} uniformly random, but with equal size (when N is even) or their sizes differ by at most 1 (when N is odd). Suppose we have at least $N \geq 4$. Denote worker i 's group ID as $g(i) \in \{-1, 1\}$. Then we have $|G_1|, |G_{-1}| \geq 2$. For our learning algorithm, only the data/samples collected from group $-g(i)$ will be used to reward any worker i in group $g(i)$. Secondly when selecting reference worker for comparing answers for mechanism (DG13), we select from the same group $g(i)$.




Figure 1: Illustration of our mechanism.

Mechanism

We would like to treat each bonus level as an ‘‘arm’’ (as in standard MAB context) to explore with. Since we have a continuous space of bonus level B , we separate the support of bonus level B ($[0, \bar{B}]$) into finite intervals. Then we treat each bonus interval as an arm. Our goal is to select the best one of them, and bound the performance in such a selection.

We set up $\lceil T^z \rceil$ arms as follows: chooses a $0 < z < 1$, separate $[0, \bar{B}]$ into $N_a = \lceil T^z \rceil$ uniform intervals:

$$[0, \bar{B}/N_a], \dots, [(k-1)\bar{B}/N_a, k\bar{B}/N_a], \dots, [(N_a-1)\bar{B}/N_a, \bar{B}]$$

For each interval we take its right end point as the bonus level to offer: $B_k = k\bar{B}/N_a$. Denote by $\tilde{p}_{m,t}^g(B_k)$ the estimated matching probability for agents in group g under bonus level B_k , and $\tilde{p}_t^g(B_k)$ s the estimated $\bar{p}(B_k)$ for group g , at stage t ; and we use $\tilde{U}_{\tilde{p}}(B)$ to denote the estimated utility function when using a noisy $\bar{p}(B)$ (\tilde{p}), instead of the true ones. We present Mechanism 1.³

Note since we assign the same number of tasks to each labeler at all stages, we have the regret defined in Eqn. (3) become equivalent with the following form:

$$R(T) = \sum_{t=1}^T \left| \sum_{g \in \{1, -1\}} \omega_g(t) \cdot \mathbb{E}[\mathcal{U}(B_{g,t}^*)] - \mathcal{U}(B^*) \right|,$$

where when t is in exploration stages $\omega_g(t) \equiv 1/2$, otherwise $\omega_g(t) := |G_g|/N$.

Equilibrium analysis: workers' side

Denote a worker's action profile at step t as $\mathbf{a}_i(t) := (e_i(t), r_i(t))$. We adopt BNE as our solution concept:

³We assume we know p_L or a non-trivial lower bound on $p_L > 0.5$.

Mechanism 1 (SPP_PostPrice)

Initialization: $t = 1$. $D(t) := t^\theta \log t$, $0 \leq \theta \leq 1$. Explore each bonus level B_k once and update $\tilde{p}_{m,t}^g(B_k)$ (for details please refer to the exploration phases), and set the number of explorations as $n_i(t) = 1$.

for $t = 1$ to T **do**

Set $\mathcal{E}(t) := \{i : n_i(t) < D(t)\}$.

if $\mathcal{E}(t) \neq \emptyset$ **then**

Exploration: randomly pick $B_k, k \in \mathcal{E}(t)$ to offer.
Follow subroutine (Explore_Crowd).

else

Exploitation: $S(t) = \mathcal{C}$. Offer $B_{g,t}^*$ to $i \in G_g$:

$$B_{g,t}^* = \operatorname{argmax}_{B_k, k=1,2,\dots,\lceil T^z \rceil} \tilde{U}_{\tilde{p}_t^g(B_k)}(B_k),$$

Follow (DG13) for workers in each group G_g .

end if

end for

Mechanism 2 (Explore_Crowd)

At exploration phases,

- 1: Randomly select two workers $S_g(t) = \{i^g(t), j^g(t)\}$ from each group G_g ; $S(t) := \cup_g S_g(t)$.
- 2: Assign 1 common and d distinct tasks to each $S_g(t)$.
- 3: Denote by $\mathcal{E}_k(t)$ the set of time steps the algorithm enters exploration and offered price B_k by time t . Estimate the probability of matching $\tilde{p}_{m,t}^g(B_k)$ for each crowd $G_g, g \in \{-1, 1\}$ by averaging:

$$\tilde{p}_{m,t}^g(B_k) = \frac{1}{|\mathcal{E}_k(t)|} \sum_{t' \in \mathcal{E}_k(t)} 1(L_{i-g(t')}(t') = L_{j-g(t')}(t')),$$

and reset $\tilde{p}_{m,t}^g(B_k)$ to $\max\{\tilde{p}_{m,t}^g(B_k), 1/2\}$. $L_i(t')$ denotes the label for the common task at time t' .

- 4: Compute $\tilde{p}_t^g(B_k)$ (estimate for $\bar{p}(B_k)$ at time t):

$$\tilde{p}_t^g(B_k) = 1/2 + \sqrt{2\tilde{p}_{m,t}^g(B_k) - 1/2}.$$

Reset $\tilde{p}_t^g(B_k) := \max\{\tilde{p}_t^g(B_k), p_L\}$.

Definition 3. A set of reporting strategy $\{\tilde{\mathbf{a}}_i := \{\tilde{\mathbf{a}}_i(t)\}_{t=1,\dots,T}\}_{i \in \mathcal{C}}$ is a BNE if for any $i, \forall \tilde{\mathbf{a}}_i' \neq \tilde{\mathbf{a}}_i$ we have

$$\begin{aligned} & \sum_t \mathbb{E}[\max_{e_i(t), r_i(t)} u_i(\tilde{\mathbf{a}}_i, \tilde{\mathbf{a}}_{-i}) | \tilde{\mathbf{a}}_i(1:t-1), \tilde{\mathbf{a}}_{-i}(1:t-1)] \\ & \geq \sum_t \mathbb{E}[\max_{e_i(t), r_i(t)} u_i(\tilde{\mathbf{a}}_i', \tilde{\mathbf{a}}_{-i}) | \tilde{\mathbf{a}}_i'(1:t-1), \tilde{\mathbf{a}}_{-i}(1:t-1)]. \end{aligned}$$

We first characterize the equilibrium strategy for workers' effort exertion and reporting with (SPP_PostPrice).

Lemma 4. At a symmetric BNE, the strategy for workers can be decoupled into a composition of myopic equilibrium strategies, that is $e_i(t) = c^*(B^{*,g(i)}(t))$, combined with $r_i(t) \equiv 1, \forall i, t$.

Proof. W.l.o.g., consider worker i in G_1 . We are going to reason that deviating from one step BNE strategy for effort

exertion is non-profitable for worker i , when other players are following the equilibria. Since the one stage equilibrium strategy maximizes the utility at current stage, and it does not affect the past utilities that have been already collected, the potential gain by deviating comes from the future gains in utilities in: (1) the offered bonus level (2) matching probability from other peers. For the bonus level offered to worker i , it will only be computed using observed data from workers in G_{-1} at exploration phases. Note for our online learning process, the exploration phases only depend on the pre-defined parameter $D(t)$, which does not depend on worker i 's data (deterministic exploration). Similarly for all other workers $j \in G_1$ (reference workers), their future utility gain is not affected by worker i 's data. Therefore an unilateral deviation from worker i will not affect the matching probability from other peers. So no deviation is profitable. \square

Again consider colluding workers. Potentially when offered a certain bonus level, workers can collude by not exerting effort regardless of their cost realization, so to mislead the learner into believing that in order to incentivize certain amount of effort, a much higher bonus level is needed.⁴ There potentially exist infinitely many colluding strategies for workers to game against a sequential learning algorithm. We focus on the following easy-to-coordinate (for workers), yet powerful strategy (Collude_Learn):

Definition 5 (Collude_Learn). *Workers collude by agreeing to exert effort when the offered bonus level $B_{g(i),t}^*$ satisfies $B_{g(i),t}^* \geq B(c_{i,t}(k)) + \nabla B$, where $c_{i,t}(k)$ is the cost for workers i to exert effort for task k at time t . $\nabla B > 0$ is a collusion constant.*

In doing so workers mislead the learner into believing that a higher bonus (differs by ∇B) is needed to induce certain effort. The next lemma establishes the collusion-proofness:

Lemma 6. *Colluding in (SPP_PostPrice) via (Collude_Learn) is not an equilibrium.*

In fact the reasoning for removing *all* symmetric colluding equilibria is similar – regardless of how others collude on effort exertion, when a worker's realized cost is small enough, he will deviate.

Performance of (SPP_PostPrice)

We impose the following assumptions⁵:

Assumption 7. $U(\cdot)$ is Lipschitz in both $\bar{p}(B)$ and B :

$$|\tilde{U}_{\bar{p}}(\tilde{B}) - U(B)| \leq L_1 |\tilde{p}(\tilde{B}) - \bar{p}(\tilde{B})| + L_2 |\tilde{B} - B|,$$

$L_1, L_2 > 0$ are the Lipschitz constants.

We have the following theorem:

Theorem 8. *The regret of (SPP_PostPrice) is:*

$$R(T) \leq O(\lceil T^{\theta+z} \log T \rceil) + \frac{L_1 C(\theta)}{\sqrt{2p_L - 1}} T^{1-\theta/2} + L_2 C(z) \bar{B} T^{1-z} + \text{const.}$$

⁴This is similar to the colluding strategy that contributes uninformative signals we studied in Section .

⁵Please refer to our full version for justification.

$0 < \theta, z < 1$ are tunable parameters. $C(\theta), C(z) > 0$ are constants. The optimal regret is $R(T) \leq \tilde{O}(T^{3/4})$ when setting $z = 1/4, \theta = 1/2$.

Proof. (Sketch) First notice by triangular inequality we know $R(T) \leq \sum_{g \in \{1, -1\}} \omega_g R_g(T)$, i.e., the total regret is upper bounded by the weighted sum of each group's regret. Since the two learning processes for the two groups G_1, G_{-1} parallel each other, we omit the g super- or subscript. We analyze for the regret incurred at the exploration and exploitation stages respectively. For exploitation regret, we first characterize the estimation error for estimating each $U(B_k)$. First by mean value theorem we can show:

$$|\bar{p}(B_k) - \tilde{p}_t(B_k)| \leq \frac{1}{2\sqrt{2p_L - 1}} |p_m(B_k) - \tilde{p}_{m,t}(B_k)|.$$

At time t , an exploitation phase, there are $D(t)$ number of samples guaranteed; so the estimation $\tilde{p}_{m,t}(B_k)$ satisfies: $\Pr[|p_m(B_k) - \tilde{p}_{m,t}(B_k)| \geq \frac{1}{t^{\theta/2}}] \leq \frac{2}{t^2}$. Then w.h.p., we have established that $|\bar{p}(B_k) - \tilde{p}_t(B_k)| \leq \frac{t^{-\theta/2}}{2\sqrt{2p_L - 1}}$. Further by Lipschitz condition we know $|\tilde{U}(B_k) - U(B_k)| \leq \frac{L_1 \cdot t^{-\theta/2}}{2\sqrt{2p_L - 1}}$. For any B that falls in the same interval as B_k we know: $|\mathcal{U}(B) - U(B_k)| \leq L_2 |B - B_k| \leq L_2 \bar{B} T^{-z}$.

Denote by $B_t^* = \operatorname{argmax}_{B_k, k=1,2,\dots,\lceil T^z \rceil} \tilde{U}_t(B_k)$, i.e., B_t^* is the estimated optimal bonus level at time t – at any time t , by searching through all arm and we can find the one maximizes the utility function from the empirically estimated function the learner currently has. Combine above arguments, we can prove $\mathcal{U}(B^*) - \mathcal{U}(B_t^*) \leq \frac{L_1 \cdot t^{-\theta/2}}{2\sqrt{2p_L - 1}} + L_2 \bar{B} T^{-z}$. Then the exploration regret can be bounded as

$$\begin{aligned} & \sum_{t=1}^T \left(\frac{L_1 \cdot t^{-\theta/2}}{2\sqrt{2p_L - 1}} + L_2 \bar{B} T^{-z} \right) + O\left(\sum_{t=1}^T \frac{2}{t^2}\right) \\ & \leq \frac{L_1 \cdot C(\theta)}{2\sqrt{2p_L - 1}} T^{1-\theta/2} + L_2 \cdot C(z) \bar{B} T^{1-z} + \text{const.} \end{aligned}$$

where we have used the fact that for any $0 < \alpha < 1$, there exists a constant $C(\alpha)$ such that $\sum_{t=1}^T \frac{1}{t^\alpha} \leq C(\alpha) T^{1-\alpha}$. The total number of explorations are $(T^\theta \log T)$ number of explorations needed for each of the $\lceil T^z \rceil$ arms: $\lceil T^z \rceil \cdot T^\theta \log T = \lceil T^{\theta+z} \log T \rceil$. Sum over we finish the proof. \square

Beyond game theoretical model

So far we have modeled the workers as being fully rational, and the reports as coming from game theoretical responses. Consider the case workers' responses do not necessarily follow a game (or arguably no one is fully rational). Instead we assume each worker has a labeling accuracy $p_i(B)$ for different B , where $p_i(B)$ can come from different models, being game theory driven, behavioral model driven or decision theory driven, and can be different for different workers.

Challenges and a mean filed approach: With this model, we can again write $\mathcal{U}(B)$ as a function of $\{p_i(B)\}_i$ and B . In order to select B^* , again we need to learn

$p_i(B)$. We re-adopt the bandit model we described earlier, and estimate $p_i(B)$ s via observing the matching probability between worker i and a randomly selected reference worker j : For each i, B we define $\bar{p}_{-i}(B) := \sum_{j \in G_{g(i)} \setminus i} p_j(B) / |G_{g(i)} \setminus i|$, and we have

$$p_m^i(B) = p_i(B)\bar{p}_{-i}(B) + (1 - p_i(B))(1 - \bar{p}_{-i}(B)), \quad (6)$$

where $p_m^i(B)$ is the probability of observing a matching for worker i , when a random reference worker is drawn (uniformly) from his group. The above forms a system of quadratic equations in $\{p_i(B)\}_i$ when $\{p_m^i(B)\}_i$ s are known. We then need to solve a perturbed quadratic equations for $\{p_i(B)\}_i$, with $\{p_m^i(B)\}_i$ s being estimated via observations (Step 3 of (EXplore_Crowd)). The following challenges arise for analysis: (1) it is hard to tell whether the solution for above quadratic equations is unique or not. (2) Solving a set of perturbed (error in estimating $\{p_m^i(B)\}_i$ s) quadratic equations for each B incurs heavy computations.⁶

Instead, by observing the availability of relatively large and diverse population of crowd workers, we make the following mean filed assumption:

Assumption 9. For any worker i , $\bar{p}_{-i}(B) \equiv \bar{p}_{g(i)}(B)$.

That is one particular worker's expertise level does not affect the crowd's mean. This is not a entirely unreasonable assumption to make, as the candidate pool of crowd workers is generally large. With above $p_m^i(B)$ then becomes

$$p_m^i(B) = p_i(B)\bar{p}_{g(i)}(B) + (1 - p_i(B))(1 - \bar{p}_{g(i)}(B)).$$

Averaging over $i \in G_{g(i)}$ we have:

$\sum_{i \in G_{g(i)}} p_m^i(B) / |G_{g(i)}| = \bar{p}_{g(i)}^2(B) + (1 - \bar{p}_{g(i)}(B))^2$, which is very similar to the matching equation we derived earlier on. Again we can solve for $\bar{p}_{g(i)}(B)$ as a function of $\sum_{i \in G_{g(i)}} p_m^i(B) / |G_{g(i)}|$. Plugging $\bar{p}_{g(i)}(B)$ back into Eqn. (6), we obtain an estimate of $p^i(B)$ as follows:

$$p_i(B) = (p_m^i(B) + \bar{p}_{g(i)}(B) - 1) / (2\bar{p}_{g(i)}(B) - 1).$$

Similar regret can be obtained – the difference only lies in estimating $p_i(B_k)$ s. Details can be found in (Liu and Chen 2016b).

Conclusion

We studied the sequential peer prediction mechanism for eliciting effort using posted price. We improve over status quo towards making the peer prediction mechanism for effort elicitation more practical: (1) we propose a posted-price and “minimal” sequential peer prediction mechanism with bounded regret. The mechanism does not require workers to report additional information, except their answers for assigned tasks. Further we show our learning results can generalize to the case when workers may not necessarily be fully rational, under a mean-filed assumption. (2) Workers exerting effort according to an informative threshold strategy and reporting truthfully is an equilibria that returns highest utility.

⁶We do not claim this is impossible to do. Rather, analyzing the output from such a system of perturbed quadratic equations merits a further study.

Acknowledgement: We acknowledge the support of NSF grant CCF-1301976.

References

- Abernethy, J.; Chen, Y.; Ho, C.-J.; and Waggoner, B. 2015. Actively Purchasing Data for Learning. In *ACM EC 2015*.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47:235–256.
- Chawla, S.; Hartline, J.; and Nekipelov, D. 2014. Mechanism Design for Data Science. In *ACM EC 2014*, 711–712.
- Dasgupta, A., and Ghosh, A. 2013. Crowdsourced Judgment Elicitation with Endogenous Proficiency. In *WWW 2013*, 319–330.
- Jurca, R., and Faltings, B. 2007. Collusion-resistant, Incentive-compatible Feedback Payments. In *ACM EC 2007*, 200–209.
- Jurca, R., and Faltings, B. 2009. Mechanisms for Making Crowds Truthful. *JAIR* 34(1):209.
- Lai, T. L., and Robbins, H. 1985. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics* 6:4–22.
- Liu, Y., and Chen, Y. 2016a. Learning to Incentivize: Eliciting Effort via Output Agreement. In *IJCAI 2016*.
- Liu, Y., and Chen, Y. 2016b. Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices. *AAAI 2017, full version*, <http://arxiv.org/abs/1611.09219>.
- Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9):1359–1373.
- Prelec, D. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306(5695):462–466.
- Radanovic, G., and Faltings, B. 2013. A Robust Bayesian Truth Serum for Non-Binary Signals. In *AAAI 2013*.
- Roth, A., and Schoenebeck, G. 2012. Conducting Truthful Surveys, Cheaply. In *ACM EC 2012*, 826–843.
- Segal, I. 2007. The Communication Requirements of Social Choice Rules and Supporting Budget Sets. *Journal of Economic Theory* 136(1):341–378.
- Shnayder, V.; Agarwal, A.; Frongillo, R.; and Parkes, D. C. 2016. Informed Truthfulness in Multi-Task Peer Prediction. *ACM EC 2016*.
- Waggoner, B., and Chen, Y. 2014. Output Agreement Mechanisms and Common Knowledge. In *HCOMP 2014*.
- Witkowski, J., and Parkes, D. 2012a. A Robust Bayesian Truth Serum for Small Populations. In *AAAI 2012*.
- Witkowski, J., and Parkes, D. 2012b. Peer Prediction without a Common Prior. In *ACM EC 2012*, 964–981.
- Witkowski, J., and Parkes, D. C. 2013. Learning the Prior in Minimal Peer Prediction. In *the 3rd Workshop on SCUGC*.
- Witkowski, J.; Bachrach, Y.; Key, P.; and Parkes, D. C. 2013. Dwelling on the Negative: Incentivizing Effort in Peer Prediction. In *HCOMP 2013*.