A Quality Assuring Mechanism for Crowdsourcing with Strategic Experts

Satyanath Bhat[1], Swaprava Nath[1], Onno Zoeter[2], Sujit Gujar[3], Y. Narahari[1], and Chris Dance[2]

[1] Indian Institute of Science, Bangalore, India
[2] Xerox Research Centre Europe, Meylan, France
[3] Xerox Research Centre India, Bangalore, India

Abstract. We consider a class of crowdsourcing problems where the owner of a task benefits from the high quality opinions provided by experts. Executing the task at an assured quality level in a cost effective manner in such situations becomes a mechanism design problem when the individual qualities are private information of the experts. The considered class of task execution problems falls into the category of interdependent values, where one cannot simultaneously achieve truthfulness and efficiency in the unrestricted setting due to an impossibility result [7]. We propose a mechanism QUEST, that leverages the structure of our special class of problems and guarantees allocative efficiency, ex-post incentive compatibility, ex-post individual rationality, and strict budget balance, which even the mechanism given by Mezzetti [13] cannot achieve in this setting. The ex-post individual rationality comes under a tight sufficiency condition, and we show that the above four properties are not simultaneously satisfiable if the sufficient condition is violated. To the best of our knowledge, this is the first attempt in developing a quality assuring crowdsourcing mechanism in an interdependent value setting with quality levels held private by strategic agents.

1 Introduction

This paper is motivated by several real-world situations where an organization or a task owner would benefit from the opinions of several experts in making crucial binary decisions. Crowdsourcing often facilitates such opinion aggregation from a set of skilled workforce or experts. We motivate our problem by describing a few such situations.

Suppose a company wants to launch a high risk new product and would like to take a go/no-go decision by collecting the opinions of domain experts. Using a crowdsourcing platform, the company may select a certain number of experts to get a binary (yes/no) opinion from the experts. The experts selected will delve into the details of the product, may glean many relevant documents, use their domain knowledge, etc., to come up with a decision. There is a cost involved here in procuring the services of each expert, which is often a common knowledge. Each expert selected for providing opinion has a certain level of proficiency or *quality* in arriving at a decision. This quality is private to the expert; however, an expert with strategic intent may not reveal this quality truthfully in the hope of gaining higher rewards. The company arrives at its final decision by aggregating the opinions collected using some algorithm and the value that accrues to the company will depend on the true qualities of the experts selected. This value is observable once the company implements the decision and the experts could be rewarded based on the actual outcome.

As our second example, let us consider a healthcare provider who is trying to decide whether or not to carry out a high risk but critical surgery on a patient depending on her medical reports. In addition to an in-house doctor, the administration of the hospital would like to take expert opinions from elsewhere. With the advent of Internet, it is easy to outsource the reports to remote experts anywhere in the world and gather their opinions. Whether or not the patient recovers is observable, and the experts can be rewarded accordingly. The practice of outsourcing is quite intense in the area of healthcare, since there are very few occupational health professionals such as doctors and nurses covering different specialties at the hospitals [15]. Also, some U.S. hospitals are outsourcing the tasks of reading and analyzing scan reports to companies in Bangalore, India [1]. Gupta et al. [3] give a detailed description of how the healthcare industry uses the outsourcing tool.

Our third example concerns a common situation facing a funding agency such as the World Bank while taking a decision on whether or not to fund a major project proposal. Typically the funding agency would select a team of experts to study the project proposal in detail and come up with their individual opinions (yes or no decisions). The funding agency will aggregate the opinion of all the experts commissioned for this purpose and take a final overall decision. Here again, the actual outcome is observable.

Aggregating the opinions of the experts and awarding monetary compensation is an example of *crowdsourcing* [5] which has recently become popular, and it provides an efficient way of task execution in business, healthcare, Government, legal matters, politics, etc. With the proliferation of the Internet-based crowdsourcing platforms like Amazon Mechanical Turk, Innocentive, oDesk, Topcoder, etc., which provide access to experts and skilled workforce, crowdsourcing has gained in popularity for such niche applications also.

All the scenarios above call for a strict quality assurance of the strategic experts. If the qualities were known to the designer, this problem reduces to a stochastic optimal control problem that can be solved using standard optimization techniques. However, the problem becomes non-trivial when the qualities are private information of the strategic experts and unknown to the designer, as the agents can potentially misreport their qualities in order to maximize their payoffs. We need a mechanism design approach that should ideally satisfy (a) allocative efficiency (allocation maximizes the total value of the agents), (b) incentive compatibility (experts report their true qualities), (c) individual rationality (the center and the experts obtain a non-negative utility through participation), and (d) budget balance (the payments by the center to the experts and the receipts by the experts balance out each other).

In this paper, we use the following abstraction of the above problems. A center (task owner) wants an item (e.g. a document) to be binary labeled. A set of agents (experts) with varied qualities are contacted using a crowdsourcing or any other platform. Depending on the quality reports of the agents, the mechanism selects a subset of the experts for labeling the item. The center observes their labels and finally observes the true label and rewards the allocated experts accordingly.

1.1 Overview and Main Results

In this paper, we pose crowdsourcing to strategic experts as a mechanism design problem. The labelers perfectly know their own qualities for labeling an item. If an expert is allocated the task of labeling, she incurs a cost (which is publicly known) to observe the noisy label of the document. On the other hand, the value to the center is the reward (or loss, which is also observable publicly) it earns after making a decision based on the reports of the labelers. The goal of the mechanism designer in this setting is to design an allocation rule and a payment rule that elicit the true qualities of the labelers, encourage participation of all players, maximize the social welfare, and ensure budget balance.

We show that this formulation belongs to an interdependent valuation setting, since the reward of the center is dependent on the qualities of all the labelers. The impossibility result by Jehiel and Moldovanu [7] makes mechanism design difficult in this setting. Though Mezzetti [13] circumvents this problem by proposing a two stage mechanism, that is not enough to guarantee all the desirable properties mentioned above. Our main contributions are as follows.

- We propose a novel mechanism QUEST (QUality Elicitation from STrategic agents) that is *ex-post incentive compatible* (EPIC) (Theorem 1) and *ex-post individually rational* (EPIR) under a sufficient condition (Theorem 2). In addition, it is *allocatively efficient* (AE) and *strictly budget balanced* (SBB) (Observation 1). The novelty and non-triviality of our mechanism lie in achieving the above properties in an interdependent value setting exploiting certain special characteristics of this problem.
- We show that the above four properties cannot be satisfied simultaneously if the sufficient condition is violated (Theorem 3). QUEST, therefore, delivers the properties with the weakest possible sufficient condition.
- We contrast our mechanism to a classic mechanism given by Mezzetti [13] for interdependent values and show that it does not guarantee SBB, while ours does (Section 4.2).

To the best of our knowledge, this is the first attempt in developing a quality assuring crowdsourcing mechanism in an interdependent value setting with quality levels held private by strategic agents. We propose a mechanism that satisfies all the four properties mentioned above with the minimal sufficient condition. The proposed mechanism overcomes the limitations of applying the VCG mechanism which cannot handle the interdependent setting.

1.2 Related Work

Mechanism design has been used in the literature as a tool to analyze crowdsourcing problems. Gao et al. [2] consider a crowdsourcing contest where the competing workers win a reward by exerting the most extra effort. They design a contest to maximize the expected quality at the center while trading it off with the risk (or variance). Lin et al. [11] propose a graphical model to represent the multiple workflow scenario and provide algorithms to learn the parameters of the model. They empirically show the superiority of their approach to existing single workflow models.

Stein et al. [18] look at the task of procuring services under a strict deadline as a mechanism design problem. A mechanism for determining near optimal prices for performing tasks in online labor markets that use crowdsourcing is presented by Singer and Mittal [17]. Jain et al. [6] develop incentive mechanisms for online question answer forums. Ramchurn et al. [16] propose trust based mechanisms for procurement scenarios where there exists uncertainty about agents successfully completing their assigned tasks. These mechanisms take into account the subjective measures of the probability of success of an agent and produce allocations that are efficient, incentive compatible, and individually rational. Jurca et al. [8] look at quality of service monitoring by a trusted monitor based on clients' truthful feedback on a service provider. Minder et al. [14] present a platform for crowdsourcing that assumes the worker abilities to be common knowledge and the costs are private.

Ho and Vaughan [4] look into the problem of online assignment in crowdsourcing markets and propose a two phase explore-exploit assignment algorithm. However, they assume honest agents and also that costs are the same for all the agents. Our paper overcomes these two limitations by offering a mechanism design solution with individual costs. We propose a model that is applicable to a certain sub-domain of the task outsourcing setting, and provide a mechanism that satisfies four very essential properties.

The rest of the paper is organized as follows. In Section 2, we present the formal model and the definitions. The proposed mechanism QUEST is presented in Section 3 and its properties are presented in Section 4. We discuss the applicability of VCG in Section 4.1, compare with the mechanism of Mezzetti in Section 4.2 and conclude the paper in Section 5.

2 The Model and Definitions

Let the set of players be denoted by N_p, which consists of a center (player 0) and n labelers $N = \{1, \ldots, n\}$, i.e., $N_p = \{0\} \cup N$.[4]

The center brings in a task where the final outcome y can take binary values in the set $\{0, 1\}$ according to a Bernoulli distribution with parameter θ, and this parameter is a common knowledge. The goal of the center is to improve the accuracy to predict y using experts' (the labelers') advice.

Labeler i has an intrinsic quality, given by q_i which is the probability of a correct observation. If the observed label is $\tilde{y}_i$, then $q_i = \mathbb{P}(\tilde{y}_i = y)$. The labelers also have a cost to make this observation, given by c_i, which is assumed to be

[4] For simplicity we have assumed a single center but the model can be extended to multiple centers as the payoffs of the players are additive for multiple tasks

common knowledge. However, the quality $q_i \in [0,1]$ is private information of the labeler, and that constitutes the type set of agent i.

We assume that the center's type set is a singleton, hence he does not report any private information. For brevity of notation, we suppress θ and c_i's when they are clear from the context.

A direct revelation mechanism $M = \langle S, r, \mathcal{P} \rangle$, decides the following: (a) an allocation $S(\hat{q}) \subseteq N$ of the labelers given the quality reports of the labelers, given by $\hat{q}$, (b) the label $r(\tilde{y}^{S(\hat{q})}(q), \hat{q})$ from the binary set after the true observations $\tilde{y}^{S(\hat{q})}(q)$ are received from the selected labelers, where q is the true quality. Note that the observations come from the players that belong to $S(\hat{q})$, but are functions of the true quality, since that is the noise with which they observe y. We assume that the actual labels of the labelers $\tilde{y}^{S(\hat{q})}(q)$ are observable by the center and therefore cannot be misreported. (c) The payment is decided after the true y is realized. Each labeler $i \in S(\hat{q})$ receives $\mathcal{P}_i(S(\hat{q}), \tilde{y}^{S(\hat{q})}(q), \hat{q}, y)$ and the consolidated sum is charged to the center (player 0). We adopt the notation t_i to denote the transfer *to* the agent. Hence,

$$t_i = \begin{cases} \mathcal{P}_i(S(\hat{q}), \tilde{y}^{S(\hat{q})}(q), \hat{q}, y) & i \in S(\hat{q}) \\ 0 & i \in N \setminus S(\hat{q}) \\ -\sum_{i \in S(\hat{q})} \mathcal{P}_i(S(\hat{q}), \tilde{y}^{S(\hat{q})}(q), \hat{q}, y) & \text{for } i = 0 \end{cases}$$

The reward generated by the center after the true y is observed is given by the reward matrix R, which gives a reward of $R(r, y)$, when the label decided by the mechanism is r and the true label is y. We assume this reward is observable to all the participants and the mechanism designer.

The value of the agents in the mechanism M is given by,

$$v_i = \begin{cases} -c_i & i \in S(\hat{q}) \\ 0 & i \in N \setminus S(\hat{q}) \\ R(r(\tilde{y}^{S(\hat{q})}(q), \hat{q}), y) & i = 0 \end{cases}$$

Note that the valuation at the center (player 0) depends on the qualities of all the selected labelers, as the observed $\tilde{y}^{S(\hat{q})}$ is a function of the true q. This makes this problem fall under the interdependent valuation setting [10]. The utility of the agents are quasi-linear, and are given by,

$$u_i^M(\hat{q}, \tilde{y}^{S(\hat{q})}(q), y | q) = v_i + t_i,$$

where q denotes the true quality vector and $\hat{q}$ is the reported one. The dynamics of the mechanism is shown in Figure 1.

Let us now define the social welfare with and without an agent i, which will be useful in presenting the main mechanism of this paper.

Definition 1 (Social Welfare). *For a label selection rule r and a labeler selection rule S, when the true label y is observed, the center obtains a reward $R(r(\tilde{y}^S(q), \hat{q}), y)$ and each selected labeler $i \in S(\hat{q})$ incurs a cost c_i. Then, the social welfare is given by the net gain of the system,*

$$W(r, S(\hat{q}), \hat{q}, y | q) = R(r(\tilde{y}^{S(\hat{q})}(q), \hat{q}), y) - \sum_{j \in S(\hat{q})} c_j. \tag{1}$$

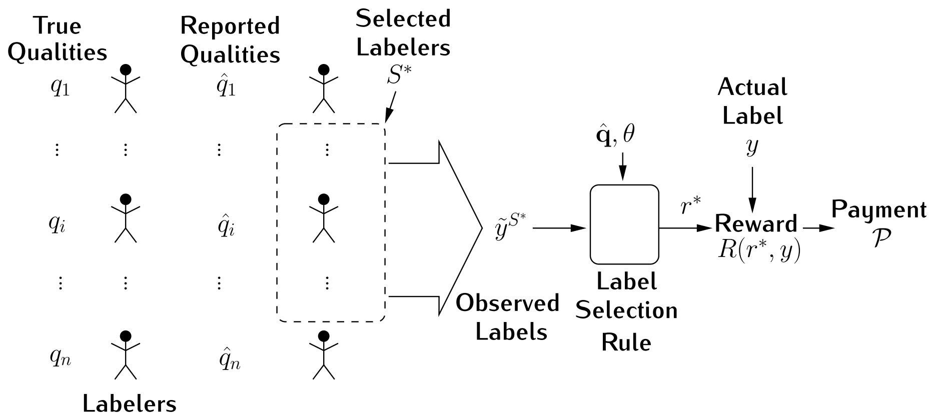

Fig. 1. Illustration of the mechanism design problem

Having chosen labeler set S, one can evaluate the worth of an agent $i \in S$ by calculating the expected social welfare in absence of i for the same observed y.

Definition 2 (Social Welfare in the Absence of i). *For a label selection rule r and a labeler selection rule S, when the true label y is observed, the social welfare in the absence of i is defined as,*

$$\begin{aligned} & W_{-i}(r, S(\hat{q}_{-i}), \hat{q}_{-i}, y|q_{-i}) \\ & = \mathbb{E}_X \left[R(r(\tilde{y}^{S(\hat{q}_{-i})}(q_{-i}), \hat{q}_{-i}), y) - \sum_{j \in S(\hat{q}_{-i})} c_j \right], \end{aligned} \tag{2}$$

where $X = \tilde{y}^{S(\hat{q}_{-i}) \setminus S(\hat{q})} | y, \hat{q}_{-i}$. We define $W_{-i} = 0$ for $i = 0$, i.e., the absence of the center yields no social welfare.

The set $S(\hat{q}_{-i})$ may contain labelers that are not present in the set $S(\hat{q})$, and hence the labels of those "missing" labelers cannot be observed. Hence, we take expectation w.r.t. their possible reports based on the reported qualities. The expression X essentially captures this.

2.1 Design Desiderata

We now look at a list of desirable properties which a mechanism in this setting should satisfy. Let us define the expected social welfare $Q(S, q)$ as follows.

$$Q(S, q) = \begin{cases} \sum_{\tilde{y}^S \in \{0,1\}^{|S|}} \left[R^*\left(\tilde{y}^S, q, \theta\right) \mathbb{P}\left(\tilde{y}^S | q, \theta\right) \right] - \sum_{i \in S} c_i, & \text{if } S \neq \emptyset \\ \max_r \mathbb{E}_{y|\theta} R(r, y), & \text{if } S = \emptyset. \end{cases} \tag{3}$$

Let us denote, $Q_\theta := \max_r \mathbb{E}_{y|\theta} R(r, y)$. In the above equation,

$$R^*\left(\tilde{y}^S, q, \theta\right) = \max_r \sum_{y \in \{0,1\}} \mathbb{P}\left(y | \tilde{y}^S, q, \theta\right) R(r, y);$$

$$\mathbb{P}\left(\tilde{y}^S|q,\theta\right) = \sum_{y\in\{0,1\}} \mathbb{P}\left(\tilde{y}^S|y,q\right)\mathbb{P}(y|\theta);$$

$\mathbb{P}\left(\tilde{y}^S|y,q\right) = \prod_i q_i^{\mathbb{I}(\tilde{y}^i-y)}(1-q_i)^{1-\mathbb{I}(\tilde{y}^i-y)}$, where $\mathbb{I}(x) = 1$, if $x = 0$, and 0 otherwise.

Definition 3 (Allocative Efficiency). *A labeler selection rule S^{AE} is allocatively efficient if,*

$$S^{AE}(q) \in \arg\max_{S\subseteq N} Q(S,q). \tag{4}$$

Notice that, $Q(S,q) = \mathbb{E}_{\tilde{y}^S|q,\theta}\max_r \mathbb{E}_{y|\tilde{y}^S,q,\theta}W(r,S,q,y|q)$. Hence, the efficient allocation maximizes the expected social welfare. Also, $Q(S^{\text{AE}}(q),q) \geq Q_\theta$, by the definition of the maximizing term (Equation (4)). This inequality holds for any number of agents, e.g., $Q(S^{\text{AE}}(q_{-i}),q_{-i}) \geq Q_\theta$.

Since the actual qualities are private to the agents, we need to elicit them *truthfully* as we are interested in maximizing the true social welfare realized. We use *Ex-Post Incentive Compatibility (EPIC)* as the notion of truthfulness.

Definition 4 (Ex-post Incentive Compatibility, EPIC). *A mechanism $M = \langle S, r, \mathcal{P}\rangle$ is ex-post incentive compatible, if for all q, and for all $\hat{q}_i$,*

$$\mathbb{E}_{X_1} u_i^M(q_i, q_{-i}, \tilde{y}^{S_1}, y|q) \geq \mathbb{E}_{X_2} u_i^M(\hat{q}_i, q_{-i}, \tilde{y}^{S_2}, y|q), \tag{5}$$

where, $S_1 = S(q_i, q_{-i})$, $S_2 = S(\hat{q}_i, q_{-i})$, and $X_1 = \tilde{y}^{S_1}, y|q,\theta$, $X_2 = \tilde{y}^{S_2}, y|q,\theta$.

EPIC is a stronger notion of truthfulness than Bayesian Incentive Compatibility (BIC), but is weaker than Dominant Strategy Incentive Compatibility (DSIC) [12].

To ensure that the labelers participate voluntarily in this labeling exercise, the mechanism has to make sure that the expected utility before observing $\tilde{y}_i$'s or y is non-negative for every agent. This desirable property is captured by individual rationality, defined as follows.

Definition 5 (Ex-post Individual Rationality, EPIR). *A mechanism M is called ex-post individually rational, if the expected utility is non-negative for all agents, i.e.,*

$$\mathbb{E}_{\tilde{y}^{S(q)},y|q} u_i^M(q, \tilde{y}^{S(q)}, y|q) \geq 0,\ \forall\ i \in N \tag{6}$$

It should be emphasized that the term *ex-post* refers to the fact that the decisions are taken after observing the types q. The nomenclature does not relate to the realization of y, as the labeler and label selection decisions are taken before the realization of y.

Definition 6 (Budget Balance). *A mechanism is* weakly *budget balanced if net monetary transfer in the system is non-positive.*

$$\sum_{i\in N_p} t_i \leq 0, \tag{7}$$

and when the inequality is met with equality, it is called strictly *budget balanced.*

In summary, the design question is to design an *allocatively efficient*, *strictly budget balanced*, *truthful* mechanism in this setting where the agents *participate voluntarily*.

3 The QUEST Mechanism

In this section, we present our mechanism QUEST (QUality Elicitation from sTrategic agents), that selects the set of labelers S^*, decides the label r^*, and the payment to the selected labelers $\mathcal{P}^*$. Therefore, QUEST $= \langle S^*, r^*, \mathcal{P}^* \rangle$.

Definition 7 (Labeler selection rule). *We can write the labeler selection rule in terms of an* expected social welfare *as,*

$$S^*(\hat{q}) \in \arg\max_{S \subseteq N} Q(S, \hat{q}), \tag{8}$$

where $Q(S, \hat{q})$ is defined in Equation (3)*.*

Note that when the reported types are $\hat{q}$, the labelers selected by the mechanism would be $S^*(\hat{q})$. Depending on that the mechanism selects a label that maximizes its reward based on the labels reported by the labelers in $S^*(\hat{q})$.

Definition 8 (Label selection rule). *Given the reported quality vector $\hat{q}$ and the observations of the labeler set $S^*(\hat{q})$, the optimal label r^* is selected by,*

$$r^*(\tilde{y}^{S^*(\hat{q})}(q), \hat{q}) \in \arg\max_{r} \sum_{y} \mathbb{P}\left(y | \tilde{y}^{S^*(\hat{q})}(q), \hat{q}\right) R(r, y) \; . \tag{9}$$

The idea for designing the payment is to pay i her cost, and in addition a fraction $\alpha > 0$ of i's marginal contribution. Let W be the realized social welfare, as stated in Definition 1 (based on S^*, r^*, $\hat{q}$, and the observed true label y), and W_{-i} be the social welfare excluding i, given by Definition 2. Using the above setup, we define the payment rule as follows.

Definition 9 (Payment rule).

$$\begin{aligned}
t_i &= \mathcal{P}_i^*(S^*(\hat{q}), \tilde{y}^{S^*(\hat{q})}(q), \hat{q}, y) \\
&= \begin{cases} \alpha \times [W(r^*, S^*(\hat{q}), \hat{q}, y | q) & \\ \qquad -W_{-i}(r^*, S^*(\hat{q}_{-i}), \hat{q}_{-i}, y | q_{-i})] + \; c_i, & \textit{if } i \in S^*(\hat{q}) \\ 0, & \textit{if } i \in N \setminus S^*(\hat{q}) \end{cases} \\
t_0 &= - \sum_{i \in S^*(\hat{q})} \mathcal{P}_i^*(S^*(\hat{q}), \tilde{y}^{S^*(\hat{q})}(q), \hat{q}, y)
\end{aligned} \tag{10}$$

This payment rule makes labelers partners in the center's venture. Theorem 2 shows how the choice of α becomes crucial to ensure EPIR.

Algorithm 1 shows the dynamics of QUEST using pseudo-code. Figure 2 shows the dependency of the different variables of this problem using a *multi-agent influence diagram* (MAID) [9].

Algorithm 1 QUEST

for agents $i = 1, \ldots, n$ **do**
  agent i observes q_i;
  agent i reports $\hat{q}_i$;
end for
select labelers $S^*(\hat{q})$ according to Definition 7;
for agents in $S^*(\hat{q})$ **do**
  center observes noisy label $\tilde{y}_i$ of labeler i;
end for
center reports $r^*(\tilde{y}^{S^*}(q), \hat{q})$ as per Definition 8;
true state of the document y is realized;
social welfare W is realized
make payment $\mathcal{P}_i^*$ to agent i, as per Definition 9;
charge an amount of $\sum_{i \in S^*} \mathcal{P}_i^*$ to the center;

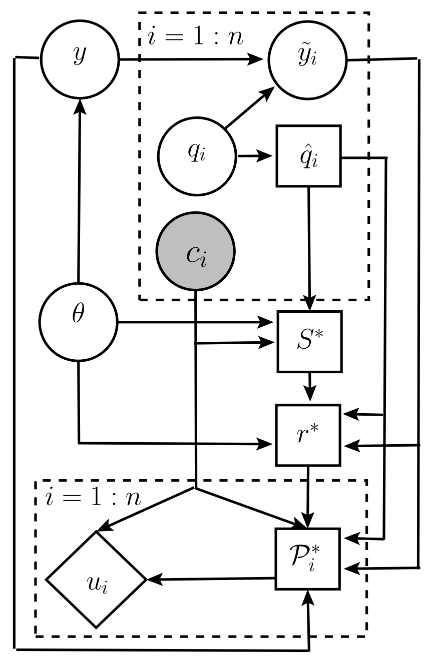

Fig. 2. Multi-agent influence diagram for QUEST

4 Properties of QUEST

The proposed crowdsourcing mechanism satisfies several important properties given by the following theorems. We denote $q = (q_i, q_{-i})$ to be the true quality of the agents. To improve readability, we move some of the proofs to Appendix. The following observation on the allocative efficiency and budget balance of QUEST follows from the definitions.

Observation 1 QUEST *is AE and SBB.*

The following theorem considers the truthfulness property of QUEST.

Theorem 1 (EPIC). QUEST *is EPIC for all agents.*

The proof is given in Appendix A.1.

Let us call quality vector q to be *Pareto better* than quality vector q_t, and denote it by $q \succeq q_t$ if, $q_i \geq q_{i,t}, \forall i$, where q_i is the i-th component of q. Observe that a labeler with $q = 0.1$ is as good as the labeler with $q = 0.9$ when the labels from the former are flipped. Hence it is enough to consider qualities only above 0.5. Let us denote the quality vector such that $q_{i,0.5} = 0.5,\ \forall i$ as $q_{0.5}$. The following theorem provides a sufficient condition for which the expected utility of each agent is non-negative under QUEST.

Theorem 2 (EPIR). *Let the problem instance (R, c, θ, q_t) be such that the labelers' qualities are Pareto better than q_t, i.e., $q \succeq q_t \succeq q_{0.5}$, and the expected social welfare at q_t is non-negative, i.e., $Q(S^*(q_t), q_t) := \epsilon(R, c, \theta, q_t) \geq 0$, then the following choice of $\alpha(\epsilon(R, c, \theta, q_t))$ ensures that* QUEST *is EPIR for all agents including center.*

$$\alpha(\epsilon(R, c, \theta, q_t)) \leq \begin{cases} \frac{1}{n} & \textit{if } Q_\theta \geq 0 \\ \frac{\epsilon(R,c,\theta,q_t)}{n(\epsilon(R,c,\theta,q_t) - Q_\theta)} & \textit{otherwise.} \end{cases}$$

Recall the parameter α above from the payment rule given by Definition 9. The condition of having a lower bound on the qualities of the agents is not a limitation since in the worst case, q_t can be a vector of 0.5 as that is the worst possible quality the mechanism designer can expect. Often the designer can have a better idea about the qualities of the agents and therefore can set a better q_t, and the α accordingly. Later, we show that if the qualities are lower bounded as given in the above theorem, no mechanism can achieve all the properties satisfied by QUEST when the expected social welfare $Q(S^*(q_t), q_t)$ is negative. This shows the tightness of the sufficient condition given by the above theorem. We will prove this theorem via the following lemma on the monotonicity of Q.

Lemma 1 (Monotonicity of Q). *If $q \succeq q_t \succeq q_{0.5}$, $Q(S^*(q), q) \geq Q(S^*(q_t), q_t)$.*

Proof sketch: We prove the lemma by proving the fact that $Q(S^*(q), q) \geq Q(S^*(q), q_t) \geq Q(S^*(q_t), q_t)$. The second inequality is true by definition of S^* and we show the first inequality for any S, which includes $S^*(q)$ as well. We argue that it is enough to consider just one labeler $i \in S$ whose quality improves when others' qualities are held fixed and show that the expected social welfare Q is non-decreasing. The detailed proof is given in Appendix A.1.

Remark 1. Lemma 1 agrees with the intuition that better quality labeler set cannot hurt the expected social welfare. However it is not very obvious given that the result holds for any reward matrix.

Proof of Theorem 2: First, let us prove that under the sufficient condition, QUEST is EPIR for the center.

Case 1: $Q_\theta \geq 0$. By definition, $Q(S^*(q), q) \geq Q_\theta$, via Equations (3) and (8). Now, if $S^*(q) = \emptyset$, center's expected payoff is $Q_\theta \geq 0$. To shorten the notation here and elsewhere in the paper, we will use W^* to denote the fact that the labeler selection and label selection were done according to the rule of QUEST (Equations 8 and 9), with only the relevant arguments inside. In the following, we use W^* to denote $W(r^*, S^*(q), q, y|q)$ and W^*_{-i} to denote $W_{-i}(r^*, S^*(q_{-i}), q_{-i}, y|q_{-i})$. If $S^*(q) \neq \emptyset$, then the center's payoff is given by,

$$\begin{aligned}
&\mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\left[R(r^*(\tilde{y}^{S^*(q)}, q, \theta), y) - \sum_{i \in S^*(q)} \mathcal{P}^*_i\right] \\
&= \mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\left[R(r^*(\tilde{y}^{S^*(q)}, q, \theta), y) - \sum_{i \in S^*(q)} c_i - \alpha|S^*(q)|W^* + \alpha \sum_{i \in S^*(q)} W^*_{-i}\right] \\
&= Q(S^*(q), q) - \alpha|S^*(q)|Q(S^*(q), q) + \alpha\mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\left[\sum_{i \in S^*(q)} W^*_{-i}\right] \\
&\geq Q(S^*(q), q) - \alpha|S^*(q)|Q(S^*(q), q) + \alpha|S^*(q)|Q_\theta \qquad (11)
\end{aligned}$$

The first equality comes by substituting Equation (10), and the second is by the fact that $Q(S^*(q), q) = \mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}[W^*]$. The inequality is due to the fact that $\mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}W^*_{-i} = Q(S^*(q_{-i}), q_{-i}) \geq Q_\theta$, and S^* is AE. Now, the last term in Equation (11) can be made non-negative by setting $\alpha \leq 1/n$, since $Q_\theta \geq 0$.

Case 2: $Q_\theta < 0$. We are given that $Q(S^*(q_t), q_t) = \epsilon \geq 0$. Therefore, $S^*(q) \neq \emptyset$ when $q \succeq q_t$. Then one can show using Lemma 1 that the last term in Equation (11) is non-negative when $\alpha \leq \frac{\epsilon}{n(\epsilon - Q_\theta)}$ and $q \succeq q_t$. Hence, we have shown that QUEST is EPIR for the center.

Next we show that QUEST is EPIR for the labelers too. If labeler i is not selected in $S^*(q)$, the payoff and cost are both 0 and EPIR holds. So we consider a q such that i is a part of $S^*(q)$. We use the shorthand S^* to denote $S^*(q)$ and S^*_{-i} to denote $S^*(q_{-i})$. Then,

$$\frac{1}{\alpha}\mathbb{E}_{\tilde{y}^{S^*},y|q}[u_i^{\text{QUEST}}(q_i, q_{-i}, \tilde{y}^{S^*}, y|q)] = \mathbb{E}_{\tilde{y}^{S^*},y|q}(\left[W^* - W^*_{-i}\right]) + (c_i - c_i)/\alpha$$

We are done if we show that $\mathbb{E}_{\tilde{y}^{S^*},y|q}W^* \geq \mathbb{E}_{\tilde{y}^{S^*},y|q}W^*_{-i}$. By EPIC, $Q(S^*(q), q) = \mathbb{E}_{\tilde{y}^{S^*},y|q}W^*$.

$$\mathbb{E}_{\tilde{y}^{S^*},y|q}W^*_{-i} = \mathbb{E}_{\tilde{y}^{S^*},y|q}\mathbb{E}_{\tilde{y}^{S^*_{-i}\setminus S^*}|y,q}\left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j\right]$$

Writing $S_1 = S^*_{-i} \cap S^*$ and observing that $\tilde{y}^{S^*_{-i} \setminus S^*}$ is independent of $\tilde{y}^{S_1}$, we get,

$$\begin{aligned}
\mathbb{E}_{\tilde{y}^{S^*},y|q} W^*_{-i} &= \mathbb{E}_{\tilde{y}^{S^*},y|q} \mathbb{E}_{\tilde{y}^{S^*_{-i} \setminus S^*}|y,\tilde{y}^{S_1},q} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= \mathbb{E}_{\tilde{y}^{S^*},\tilde{y}^{S^*_{-i} \setminus S^*},\tilde{y}^{S_1},y|q} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= Q(S^*(q_{-i}), q_{-i})
\end{aligned}$$

Now, $Q(S^*(q_{-i}), q_{-i})$ is the expected welfare of an AE outcome when i is not a part of labeler pool. The labeler selection rule S^* has the property that the alternatives $S^*(q) \in A$ contain the alternatives $S^*(q_{-i}) \in A_{-i}$. This is because the available choices of $S^*(q_{-i})$ are contained in the possible choices of $S^*(q)$. Therefore we conclude that,

$$Q(S^*(q), q) \geq Q(S^*(q_{-i}), q_{-i})$$

This concludes the proof. ■

Theorem 3 (Unachievability Result). *If the problem instance (R, c, θ, q_t) be such that the labeler quality vector $q \succeq q_t \succeq q_{0.5}$, but the sufficiency condition of Theorem 2 is violated, then no mechanism can satisfy AE, EPIC, EPIR, and SBB.*

Proof: Let us assume there exists a mechanism $M = \langle S^M, r^M, \mathcal{P}^M \rangle$ satisfies EPIR, EPIC, AE and SBB simultaneously. The labeler and label selection rule are the same as QUEST since they are AE by definition, i.e., $S^M \equiv S^*, r^M \equiv r^*$. Since M is EPIC, we can work with the true qualities q. Now, we can write the expected utility to the center by rewriting the first term in Equation (11) for the mechanism M as follows.

$$\begin{aligned}
&\mathbb{E}_{y,\tilde{y}^{S^M(q)}|q} \left[R(r^M(\tilde{y}^{S^M(q)}, q, \theta), y) - \sum_{i \in S^M(q)} \mathcal{P}^M_i \right] \\
&= \mathbb{E}_{y,\tilde{y}^{S^M(q)}|q} \left[R(r^M(\tilde{y}^{S^M(q)}, q, \theta), y) - \sum_{i \in S^M(q)} c_i + \sum_{i \in S^M(q)} c_i - \sum_{i \in S^M(q)} \mathcal{P}^M_i \right] \\
&= Q(S^*(q), q) + \sum_{i \in S^M(q)} (c_i - \mathcal{P}^M_i)
\end{aligned}$$

The last equality comes as $S^M \equiv S^*, r^M \equiv r^*$, hence the expected welfare under M is same as $Q(S^*(q), q)$, the welfare under QUEST. As the sufficiency condition of Theorem 2 is violated, it implies, $Q(S^*(q), q) = \epsilon < 0$. The $\mathcal{P}^M_i$ term indicates the payment to labeler i. For M to be EPIR for the labelers, $\mathcal{P}^M_i - c_i \geq 0$, for all i. Therefore, $\sum_{i \in S^M(q)} (\mathcal{P}^M_i - c_i) \geq 0$. For M to be EPIR for the center, $Q(S^*(q), q) + \sum_{i \in S^M(q)} (c_i - \mathcal{P}^M_i) \geq 0$, which implies, $\sum_{i \in S^M(q)} (\mathcal{P}^M_i - c_i) \leq Q(S^M(q), q) < 0$, which is a contradiction. Hence proved. ■

4.1 Why is VCG not Applicable?

Though the proposed mechanism resembles a VCG mechanism it operates under a different setting. VCG is applicable to an independent private value setting, whereas this setting is that of interdependent values. VCG does not guarantee truthfulness in an interdependent value setting [7].

4.2 Comparison with Mezzetti's Mechanism

In this paper we investigate a special sub-class of the interdependent value setting that is relevant for crowdsourcing setting and the proposed mechanism QUEST satisfies AE, EPIC, EPIR, and SBB under a sufficient condition given by Theorem 2. Let us compare QUEST vis-à-vis the classic mechanism given by Mezzetti [13] (we will call this MZT) which too is EPIC and AE in the interdependent value setting. We note that in the first stage, MZT determines the allocation based on the type reports $\hat{q}_i$'s, and the allocation rule is the same as in QUEST (Definition 7). However, the payment in the second round is different. QUEST is SBB even after observing $(y, \tilde{y}^{S^*(q)})$. However, this is not guaranteed by MZT even *ex-ante* observing $(y, \tilde{y}^{S^*(q)})$. Let us explain why. The center's valuation after observing $(y, \tilde{y}^{S^*(q)})$ is,

$$v_0 = R(r^*(\tilde{y}^{S^*(q)}, q), y) \ .$$

We consider true q since MZT is EPIC. The value of a labeler $i \in S^*(q)$ is given by, $v_i = -c_i$. Therefore, $t_0^{\text{MZT}} = -\sum_{i \in S^*(q)} c_i$. The transfer to the labeler i is given by,

$$t_i^{\text{MZT}} = \mathcal{P}_i^{\text{MZT}} = R(r^*(\tilde{y}^{S^*(q)}, q), y) - \sum_{j \in S^*(q) \setminus \{i\}} c_j \ .$$

Therefore, the net monetary transfer is given by,

$$\sum_{i \in N_p} t_i^{\text{MZT}} = n \left(R(r^*(\tilde{y}^{S^*(q)}, q), y) - \sum_{j \in S^*(q)} c_j \right) \ .$$

If we take the expectation of the net monetary transfer w.r.t. $(y, \tilde{y}^{S^*(q)})$, the expression on the RHS becomes,

$$n\mathbb{E}_{y, \tilde{y}^{S^*(q)} | q} \left(R(r^*(\tilde{y}^{S^*(q)}, q), y) - \sum_{j \in S^*(q)} c_j \right) = nQ(S^*(q), q) \geq 0.$$

The inequality comes from the sufficient condition of Theorem 2. Hence MZT is *ex ante* BB only when the expected social welfare is zero. In the more interesting scenario, where the system generates a positive social welfare, MZT may run into a budget deficit. However, QUEST is SBB even *ex-post* observing $(y, \tilde{y}^{S^*(q)})$. The above expression also says that the *ex-ante* budget deficit in MZT is n-fold the net *ex-ante* social welfare of QUEST.

The reason of the above phenomenon is because Mezzetti [13] provides one possible implementation in the interdependent valuation setting, and does not provide a characterization result. Other mechanisms may exist that satisfy the same or more properties as MZT in a specific setting, and QUEST is one such example.

5 Conclusions and Future Work

This paper was motivated by several real world problems in which a center benefits from the opinions provided by experts but the experts themselves could exhibit strategic behavior. Achieving an assured quality level in a cost effective manner in such situations becomes a mechanism design problem when the individual qualities are private information of the experts. We proposed a mechanism QUEST, that ensures four desirable properties: allocative efficiency, ex-post incentive compatibility, ex-post individual rationality, and strict budget balance. QUEST is able to achieve the above properties in a subclass of interdependent value setting with the weakest sufficiency condition. We believe this is the first attempt in developing a quality assuring crowdsourcing mechanism in an interdependent value setting with quality levels held private by strategic agents.

Considering settings where both cost and quality are privately held by labelers is an important and interesting future work. Also extending the mechanism to a more general setting than binary labels would be a useful extension.

Acknowledgments. This work is part of a collaborative project between Xerox Research and Indian Institute of Science on incentive compatible learning for e-services.

References

[1] Associated-Press. Some U.S. hospitals outsourcing work. http://www.msnbc.msn.com/id/6621014, June 2004.

[2] X. Gao, Y. Bachrach, P. Key, and T. Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 38–44, 2012.

[3] A. Gupta, R. K. Goyal, K. A. Joiner, and S. Saini. Outsourcing in the healthcare industry: Information technology, intellectual property, and allied aspects. *Information Resources Management Journal (IRMJ)*, 21(1): 1–26, 2008.

[4] C. Ho and J. Vaughan. Online task assignment in crowdsourcing markets. In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, pages 45–51, 2012.

[5] J. Howe. *Crowdsourcing: Why the Power of the Crowd is Driving the Future of Business*. Three Rivers Press, 2009.

[6] S. Jain, Y. Chen, and D. C. Parkes. Designing incentives for online question-and-answer forums. *Games and Economic Behavior*, (0):–, 2012. ISSN 0899-8256. doi: http://dx.doi.org/10.1016/j.geb.2012.11.003. URL http://www.sciencedirect.com/science/article/pii/S0899825612001674.

[7] P. Jehiel and B. Moldovanu. Efficient Design with Interdependent Valuations. *Econometrica*, (69):1237–1259, 2001.

[8] R. Jurca, B. Faltings, and W. Binder. Reliable qos monitoring based on client feedback. In *Proceedings of the 16th international conference on World Wide Web (WWW 2007)*, pages 1003–1012. ACM, 2007.
[9] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, Volume 45(Issue 1): 181–221, October 2003.
[10] V. Krishna. *Auction Theory*. Academic Press, 2009.
[11] C. Lin, Mausam, and D. Weld. Dynamically switching between synergistic workflows for crowdsourcing. In *Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI 2012)*, 2012.
[12] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Micoreconomic Theory*. Oxford University Press, 1995.
[13] C. Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, 72(5):1617–1626, 2004.
[14] P. Minder, S. Seuken, A. Bernstein, and M. Zollinger. CrowdManager - combinatorial allocation and pricing of crowdsourcing tasks with time constraints. In *Proceedings of Workshop on Social Computing and User Generated Content, Valencia, Spain*, 2012.
[15] P. J. Nicholson. Occupational health services in the UK - challenges and opportunities. *Occupational Medicine*, 54:147–152, 2004. doi: 10.1093/occmed/kqg125.
[16] S. Ramchurn, C. Mezzetti, A. Giovannucci, J. Rodriguez-Aguilar, R. Dash, and N. Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research*, 35(1):119–159, 2009.
[17] Y. Singer and M. Mittal. Pricing tasks in online labor markets. In *Proceedings of Workshop on Human Computation (AAAI 2011)*, 2011.
[18] S. Stein, E. H. Gerding, A. Rogers, K. Larson, and N. R. Jennings. Algorithms and mechanisms for procuring services with uncertain durations using redundancy. *Artificial Intelligence*, 175(14):2021–2060, 2011.

A Appendix

A.1 Proofs not Included in the Main Text

Proof of Lemma 1

Proof: Consider two quality levels $q = \{q_1, \dots q_n\}$ and let $\tilde{q} = \{\tilde{q_1}, \dots, \tilde{q_n}\}$ such that $\forall i \in \{1, 2, \dots, n\}, q_i \leq \tilde{q_i}$. We will show that,

$$Q(S^*(q), q) \leq Q(S^*(q), \tilde{q}) \leq Q(S^*(\tilde{q}), \tilde{q}).$$

The second inequality above is true by the definition of labeler selection rule (Equation (8)). Hence we need to prove only the first inequality. In fact, we show that the inequality holds for any S, i.e., $Q(S, q) \leq Q(S, \tilde{q})$. We see that for $S = \emptyset$, $Q(S, q)$ is independent of q (Equation (3)). Hence we need to prove the result only when $S \neq \emptyset$.

Case 1: $|S| = 1$. It is clear that only the quality q of the labeler $i \in S$ influences the expected welfare $Q(S, q)$. Hence we only work with q of that labeler. For notational convenience, we will use shorthand R_{ry} to denote $R(r, y)$. For the single labeler set S, the expected social welfare is given by,

$$\begin{aligned} Q(S,q) &= \sum_{\tilde{y}\in\{0,1\}} \max_{r\in\{0,1\}} \left[\sum_{y\in\{0,1\}} R_{ry}\mathbb{P}(y,\tilde{y}|q,\theta) \right] - \sum_{i\in S} c_i \\ &= \sum_{\tilde{y}\in\{0,1\}} \max_{r\in\{0,1\}} \left[\sum_{y\in\{0,1\}} R_{ry}\mathbb{P}(\tilde{y}|y,q,\theta)\mathbb{P}(y|\theta) \right] - \sum_{i\in S} c_i. \end{aligned} \tag{12}$$

The cost term on the R.H.S. of Equation 12 appears in both $Q(S, q)$ and $Q(S, \tilde{q})$ and therefore cancels out while comparing. Hence, we can WLOG assume $c_i = 0$ for $i \in S$ to prove the lemma. Expanding, we obtain,

$$Q(S,q) \tag{13}$$

$$\begin{aligned} &= \max_{r_1\in\{0,1\}} \{\underbrace{R_{r_1 0}(1-\theta)q + R_{r_1 1}(1-q)\theta}_{=:f(r_1,q)}\} + \max_{r_2\in\{0,1\}} \{\underbrace{R_{r_2 0}(1-q)(1-\theta) + R_{r_2 1}q\theta}_{=:g(r_2,q)}\} \\ &= \max_{r_1\in\{0,1\}} f(r_1,q) + \max_{r_2\in\{0,1\}} g(r_2,q) \\ &= \max_{(r_1,r_2)\in\{0,1\}^2} (f(r_1,q) + g(r_2,q)) \\ &= \max\{R_1(q), R_2(q), R_3(q), R_4(q)\}, \end{aligned} \tag{14}$$

where,

$$\begin{aligned} R_1(q) &= f(0,q) + g(0,q) = R_{01}\theta + R_{00}(1-\theta), \qquad \textbf{invariant with } q, \\ R_2(q) &= f(0,q) + g(1,q) = mq + a, \\ R_3(q) &= f(1,q) + g(0,q) = -mq + b, \\ R_4(q) &= f(1,q) + g(1,q) = R_{11}\theta + R_{10}(1-\theta), \qquad \textbf{invariant with } q, \end{aligned}$$

where, $m = \big(R_{00}(1-\theta) - R_{01}\theta - R_{10}(1-\theta) + R_{11}\theta\big)$, $a = R_{01}\theta + R_{10}(1-\theta)$ and $b = R_{00}(1-\theta) + R_{11}\theta$. We notice that the functions R_1 through R_4 are affine functions of q and hence their maximum given by Equation (14) is convex in q. All that we need to show now is that $Q(S, q)$ attains a minima at $q = 0.5$, and together with the convexity of $Q(S, q)$, it is going to be non-decreasing within the interval $q \in [0.5, 1]$.

It can be seen that at $q = 0.5$, the two lines $R_2(q)$, $R_3(q)$ intersect and the expected welfare is given by $0.5\big(R_{00}(1-\theta) + R_{01}\theta + R_{10}(1-\theta) + R_{11}\theta\big)$. Let us denote this by d, as shown in Figure 3 for $m \geq 0$ (the complementary plot for $m < 0$ would be similar with the lines $R_2(q)$ and $R_3(q)$ flipped around $q = 0.5$). The maximizer of the two lines is given by the equation $|m(q-1/2)|+d$. Combined with the max of $R_1(q)$ and $R_4(q)$, we conclude that the expected social welfare $Q(S, q)$ does not decrease when q increases from 0.5 to 1, and Figure 3 shows one such case when $\max\{R_1(q), R_4(q)\} > d$ (shown in bold lines).

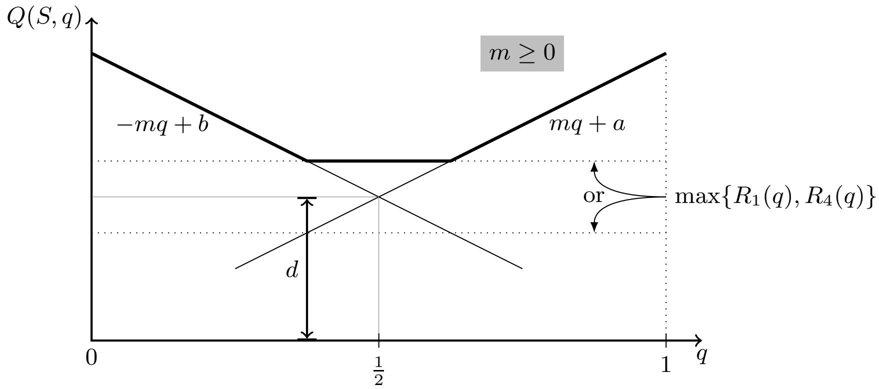

Fig. 3. Expected Welfare versus q when $|S| = 1$

Case 2: $|S| > 1$. It is enough to consider the case when for only one particular player i the quality is increased, i.e., $\tilde{q}_i \geq q_i$ and the other players' qualities are held fixed. This is because if we show that welfare increases monotonically with q_i is this case, the argument can be repeated for another player j, with i's quality at $\tilde{q}_i$, j's quality increasing from q_j to $\tilde{q}_j$ and other agents' quality held fixed. Also, if $i \notin S$, the expected welfare remains the same, hence it is enough to consider only when $i \in S$. So we find the the expected social welfare is to be,

$$\begin{aligned}
&Q(q_i, q_{-i}, S) \\
&= \sum_{\tilde{y}^{S_{-i}}} \left\{ \sum_{\tilde{y}_i \in \{0,1\}} \max_{r \in \{0,1\}} \left[\sum_{y \in \{0,1\}} R_{ry} \mathbb{P}(y, \tilde{y}_i, \tilde{y}^{S_{-i}} | q, \theta) \right] \right\} - \sum_{i \in S} c_i \\
&= \sum_{\tilde{y}^{S_{-i}}} \left\{ \sum_{\tilde{y}_i \in \{0,1\}} \max_{r \in \{0,1\}} \left[\sum_{y \in \{0,1\}} \left(R_{ry} \mathbb{P}(\tilde{y}^{S_{-i}} | y, q_{-i}, \theta) \right) \mathbb{P}(\tilde{y}_i | y, q_i, \theta) \mathbb{P}(y|\theta) \right] \right\} \\
&\qquad - \sum_{i \in S} c_i.
\end{aligned}$$

The term in curly braces for a particular fixed $\tilde{y}^{S_{-i}}$ resembles Equation (12) and hence it is non-decreasing in $q_i \in [0.5, 1]$ by a similar argument used when $|S| = 1$. Since the total welfare is sum of such functions we get the desired result. ■

Proof of Theorem 1

Since the center has a singleton type space, which is common knowledge, the EPIC result is required only for the labelers. We prove this theorem for the labelers with the aid of three lemmas. To show that QUEST is EPIC, let us assume, WLOG, that only agent i is a potential misreporter. We assume that the true type profile is given by $q = (q_i, q_{-i})$. Therefore, $\hat{q} = (\hat{q}_i, q_{-i})$. For notational simplicity, we will use the shorthands $W^*(\hat{q}_i, q_{-i})$ to denote $W(r^*, S^*(\hat{q}), \hat{q}, y|q)$ and $W^*_{-i}(\hat{q}_i, q_{-i})$ to denote $W_{-i}(r^*, S^*(\hat{q}_{-i}), \hat{q}_{-i}, y|q)$.

Lemma 2. *Let* $S_1 = S^*(\hat{q}_i, q_{-i}), S_2 = S^*(q_i, q_{-i})$ *then,* $\mathbb{E}_{\tilde{y}^{S_1},y} W^*_{-i}(\hat{q}_i, q_{-i}) = \mathbb{E}_{\tilde{y}^{S_2},y} W^*_{-i}(q_i, q_{-i})$ *for all* $\hat{q}_i$.

This lemma shows that the expected social welfare in the absence of i is independent of i's reported quality.
Proof: Write $S_3 = S^*(q_{-i}) \setminus S_1$ and $S_4 = S^*(q_{-i}) \setminus S_2$. We will use the shorthand S^*_{-i} to denote $S^*(q_{-i})$ from now on. Let us consider the following term,

$$\begin{aligned}
\mathbb{E}_{\tilde{y}^{S_2},y} W^*_{-i}(q_i, q_{-i}) &= \mathbb{E}_{\tilde{y}^{S_2},y} \mathbb{E}_{\tilde{y}^{S_4}|y} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= \mathbb{E}_{\tilde{y}^{S_2},y} \mathbb{E}_{\tilde{y}^{S_4}|\tilde{y}^{S_2},y} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= \mathbb{E}_{\tilde{y}^{S_2},\tilde{y}^{S_4},y} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= \mathbb{E}_{\tilde{y}^{S^*_{-i}(q_{-i})},y} \left[R(r^*(\tilde{y}^{S^*_{-i}}, q_{-i}), y) - \sum_{j \in S^*_{-i}} c_j \right] \\
&= \mathbb{E}_{\tilde{y}^{S_1},y} W^*_{-i}(\hat{q}_i, q_{-i}) \quad \text{(following similar steps).}
\end{aligned}$$

The first equality arises since $\tilde{y}^{S_2}$ is independent of $\tilde{y}^{S_4}$ given the true label y. The third equality is true since the term in expectation depends on $\tilde{y}^{S^*_{-i}}$ and $S_2 \cup S_4 \supseteq S^*_{-i}(q_{-i})$. ■

Lemma 3. *For any* $S \subseteq N$ *fixed, the expected social welfare is maximal when every agent in* S *reports truthfully. In other words, with true quality profile* $q = (q_i, q_{-i})$*, we have* $\mathbb{E}_{\tilde{y}^S,y|q} W^*(q_i, q_{-i}) \geq \mathbb{E}_{\tilde{y}^S,y|q} W^*(\hat{q}_i, \hat{q}_{-i})$.

Proof: Consider,

$$\begin{aligned}
&\mathbb{E}_{\tilde{y}^S,y|q} W^*(q_i, q_{-i}) - \mathbb{E}_{\tilde{y}^S,y|q} W^*(\hat{q}_i, \hat{q}_{-i}) = \mathbb{E}_{\tilde{y}^S|q} \mathbb{E}_{y|\tilde{y}^S,q} \left[W^*(q_i, q_{-i}) - W^*(\hat{q}_i, \hat{q}_{-i}) \right] \\
&= \mathbb{E}_{\tilde{y}^S|q} \mathbb{E}_{y|\tilde{y}^S,q} \left[R(r^*(\tilde{y}^S, (q_i, q_{-i})), y) - \sum_{j \in S} c_j - R(r^*(\tilde{y}^S, (\hat{q}_i, \hat{q}_{-i})), y) + \sum_{j \in S} c_j \right]
\end{aligned}$$

Let us write, $r_1 = r^*(\tilde{y}^S, (q_i, q_{-i})), y)$, $r_2 = r^*(\tilde{y}^S, (\hat{q}_i, \hat{q}_{-i})), y)$. So, we get,

$$\begin{aligned}
&\mathbb{E}_{\tilde{y}^S,y|q} W^*(q_i, q_{-i}) - \mathbb{E}_{\tilde{y}^S,y|q} W^*(\hat{q}_i, \hat{q}_{-i}) \\
&= \mathbb{E}_{\tilde{y}^S|q} \left[\sum_y \mathbb{P}\left(y|\tilde{y}^S, q\right) R(r_1, y) - \sum_y \mathbb{P}\left(y|\tilde{y}^S, q\right) R(r_2, y) \right] \\
&\geq 0.
\end{aligned}$$

The inequality arises from the label selection rule (Definition 8), as $r_1 = r^*(\tilde{y}^S, (q_i, q_{-i})), y)$ maximizes the expression $\sum_y \mathbb{P}\left(y|\tilde{y}^S, q\right) R(r_1, y)$ for all $\tilde{y}^S$. ■

Lemma 4. *Suppose* S_1, S_2 *are as defined in Lemma 2. For any* i, $\mathbb{E}_{\tilde{y}^{S_2},y|q} W^*(q_i, q_{-i}) \geq \mathbb{E}_{\tilde{y}^{S_1},y|q} W^*(\hat{q}_i, q_{-i})$, *that is, the expected social welfare for the center is maximal when everyone reports truthfully.*

Proof: Consider,

$$\begin{aligned}
\mathbb{E}_{\tilde{y}^{S_2},y|q} W^*(q_i, q_{-i}) &= \mathbb{E}_{\tilde{y}^{S_2}|q} \mathbb{E}_{y|\tilde{y}^{S_2},q} \left[R(r^*(\tilde{y}^{S_2}, q), y) - \sum_{j \in S_2} c_j \right] \\
&\geq \mathbb{E}_{\tilde{y}^{S'}|q} \mathbb{E}_{y|\tilde{y}^{S'},q} \left[R(r^*(\tilde{y}^{S'}, q), y) - \sum_{j \in S'} c_j \right] \\
&\geq \mathbb{E}_{\tilde{y}^{S'}|q} \mathbb{E}_{y|\tilde{y}^{S'},q} \left[R(r^*(\tilde{y}^{S'}, (\hat{q}_i, q_{-i})), y) - \sum_{j \in S'} c_j \right]
\end{aligned}$$

The first inequality arises from the fact that the set $S_2 = S^*$ was chosen to maximize the term inside the bracket, and so, the inequality holds for all $S' \subseteq N$. The second inequality arises from Lemma 3. In particular, the last inequality holds true even for $S' = S_1$. This completes the proof. ■

Proof of Theorem 1: The payment under QUEST is given by Equation (10). The utility of agent i is $u_i^{\text{QUEST}} = \mathcal{P}_i - c_i = \alpha(W^*(\hat{q}_i, q_{-i}) - W^*_{-i}(\hat{q}_i, q_{-i}))$. To show the mechanism is EPIC, we need to show that,

$$\mathbb{E}_{\tilde{y}^{S_2},y|q}[u_i^{\text{QUEST}}(q_i, q_{-i}, \tilde{y}^{S_2}, y|q)] \geq \mathbb{E}_{\tilde{y}^{S_1},y|q}[u_i^{\text{QUEST}}(\hat{q}_i, q_{-i}, \tilde{y}^{S_1}, y|q)],$$

where $S_1 = S^*(\hat{q}_i, q_{-i}), S_2 = S^*(q_i, q_{-i})$. This is the same as showing,

$$\mathbb{E}_{\tilde{y}^{S_2},y|q}\left[W^*(q_i, q_{-i}) - W^*_{-i}(q_i, q_{-i})\right] \geq \mathbb{E}_{\tilde{y}^{S_1},y|q}\left[W^*(\hat{q}_i, q_{-i}) - W^*_{-i}(\hat{q}_i, q_{-i})\right].$$

Now, the W^*_{-i} terms on either side of the inequality cancel out due to Lemma 2, so to show EPIC, we need to show,

$$\mathbb{E}_{\tilde{y}^{S_2},y} W^*(q_i, q_{-i}) \geq \mathbb{E}_{\tilde{y}^{S_1},y} W^*(\hat{q}_i, q_{-i}).$$

The above follows directly from Lemma 4. Hence proved. ■