A Mechanism to Optimally Balance Cost and Quality of Labeling Tasks Outsourced to Strategic Agents

Satyanath Bhat
Indian Institute of Science
Bangalore, India
satya.bhat@gmail.com

Swaprava Nath*
Indian Statistical Institute
New Delhi, India
swaprava@gmail.com

Sujit Gujar†
EPFL
Lausanne, Switzerland
sujit.gujar@epfl.ch

Onno Zoeter
Xerox Research Center
Meylan, France
onno.zoeter@xrce.xerox.com

Y. Narahari
Indian Institute of Science
Bangalore, India
hari@csa.iisc.ernet.in

Chris Dance
Xerox Research Center
Meylan, France
chris.dance@xrce.xerox.com

ABSTRACT

We consider an expert-sourcing problem where the owner of a task benefits from high quality opinions provided by experts. Execution of the task at an assured quality level in a cost effective manner 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. We propose a novel mechanism QUEST, that exploits the structure of our special class of problems and guarantees allocative efficiency, ex-post incentive compatibility, and strict budget balance. Our mechanism satisfies ex-post individual rationality for the experts and we also derive the weakest sufficient condition under which it is ex-post individual rationality for the center as well.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agents

General Terms

Algorithms, Economics, Theory

Keywords

Expert Sourcing, Strategic Agents, Interdependent Values, Mechanism Design

*This work was done when author was a graduate student in Indian Institute of Science. During this time he was supported by Tata Consultancy Services (TCS) Fellowship.

†This work was done when author was a part of Xerox Research Center, India.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

1. INTRODUCTION

For many high-stake decisions, it is interesting for decision makers to hire experts to provide additional information about the problem before a decision is made. For example, investors would seek information about a company before deciding to buy a stake, surgeons would ask for medical tests to be performed before deciding to operate a patient, and companies would do market research before deciding to launch a product.

Since such expert advice, often, is not free, a decision maker needs to decide whom to hire for advice, as a step before making the final decision. The optimal choice balances the cost of hiring experts with the benefit of increased accuracy. The balance depends on the expected return and costs of making a correct or incorrect final decision, the cost of hiring each expert, and the qualities of the experts' advice.

When all these crucial parameters are known, both the expert selection problem, and that of making the final decision given the advice, are relatively standard decision theory problems. However these problems become non-trivial if the qualities of the experts are private information. If the way experts are selected or paid off is not carefully designed, there is a risk that the experts will misreport their qualities thereby jeopardizing the decision maker's venture.

In this paper, we use the following abstraction of the above problem. A center (decision maker or task owner) wants an item (e.g. a document, company, patient, or product) to be binary labeled. A set of agents (experts) can each provide an independent noisy assessment of the label. Each agent has an associated quality (probability of observing the label without error) and cost of operating. Costs are common knowledge, but qualities need to be elicited from the agents. Depending on the quality reports of the agents, the mechanism selects a subset of the experts to label the item. The center observes their labels and then based on this information takes an action with maximal expected return. It then observes the true label and bases rewards to the allocated labelers accordingly.

This setting allows us to be concrete and precise. The formal model is presented in Section 2. Its relative simplicity also makes for an easy exposition. For many real decision making scenarios, the payment contingent on the realization of the true label is often realistic: the health state of a patient is observed during operation or after a reasonable window of time; similarly, the final performance of an investment or a new product can be observed. However, in other scenarios such as using experts to label large corpora of documents to train classifiers, the true label will not be observed. In such cases, extensions based on ideas such as peer prediction [18] can be worthwhile.

The assumptions of binary problems and known costs are restrictive. Binary (yes-no) decisions are common but extensions to multiple classes would be useful. Known costs are de-facto the case in platforms such as Mechanical Turk, and would be reasonable if the costs of running a particular tool (medical scanner, search in a financial database, etc.) can be assessed. An extension where both costs and qualities are private information is of great interest.

The setting considered here provides a natural starting point since the type of each agent is one-dimensional. A multi-dimensional extension (more than one class, both costs and qualities unknown) is interesting but a fundamentally harder problem. We will discuss how our basic mechanism can be extended to such cases in the discussion section.

1.1 Overview and Main Results

In this paper, we formulate 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 a noisy (binary) label of the document. 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.

The above 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 [11] poses challenges to the mechanism design in this setting. Though Mezzetti [17] circumvents the 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), *ex-post individually rational* (EPIR) for all labelers (Theorem 2), *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 characteristics of this problem.
- Additionally we show that QUEST is also EPIR for the center under a weak sufficient condition (Theorem 3). We show that the above four properties cannot be satisfied simultaneously if the sufficient condition is violated (Theorem 4). QUEST, therefore, delivers the properties with the weakest possible sufficient condition.
- We emphasize that QUEST is strictly budget balanced while the classic mechanism given by Mezzetti [17] for interdependent values is not (Section 4.1).

To the best of our knowledge, this is the first attempt to develop a quality assuring expert-sourcing mechanism in an interdependent value setting with quality levels held private by the strategic experts. The proposed mechanism overcomes the limitations of applying the VCG mechanism which cannot handle the interdependent value setting.

An implicit assumption in our approach is non-manipulability of agent reports. This arises naturally in the problem setting we have considered. In a crowdsourced task, the agents hired, typically, are not aware of the reports of their colleagues. Thus, the only direction of manipulation of reports available to them is to work sub-optimally after they are hired. If everyone else is truthful, then this manipulation is not distinguishable from the case of overbid of quality followed by truthful report by that agent. The latter is not IC which will be clear from the proofs given in the paper, and hence it justifies our assumption.

1.2 Related Work

Quality assurance in crowdsourcing is a widely studied problem. The non-strategic version of this problem too poses interesting challenges and is a subject of active research. Lin et al. [15] 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. Minder et al. [19] present a platform for crowdsourcing that assumes the worker abilities to be common knowledge and the costs to be private. Ho and Vaughan [9] 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 the aforementioned four very essential properties.

Mechanism design has been used in the literature as a tool to provide solutions to crowdsourcing problems [6, 13, 15, 23, 19]. Gao et al. [6] consider a crowdsourcing contest where the competing workers win a reward by exerting the highest extra effort. They design a contest to maximize the expected quality at the center while trading it off with the risk (or variance). Singer and Mittal [24] present a mechanism for determining near optimal prices for performing tasks in online labor markets. Jain et al. [10] develop incentive mechanisms for online question answer forums.

Our work also closely relates to the service procurement problem [25, 21, 12, 8, 7] and mechanism design version of the principal agent problem [1, 26, 2]. Babaioff et al. [1] consider a team of workers employed by a principal where the principal benefits out of costly effort exerted by the team. While the effort is not observable, the final outcome is. The principal wishes to optimally incentivize his team to exert effort to increase the probability of a favorable outcome. Stein et al. [25] look at the task of procuring services under a strict deadline as a mechanism design problem. Ramchurn et al. [21] 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. [12] look at quality of service monitoring by a trusted monitor based on clients' truthful feedback on a service provider.

Our setting is closely related to the forecast elicitation problem but more involved. In the standard elicitation task, the selection of an agent is given. Proper scoring rules provide an elegant answer (e.g. [3]) and continue to receive a lot of interest (e.g. [5, 4, 22]). In our setting, the agents go through two phases, their reports of their accuracies in the first phase determine if they are selected for the second round where they incur a cost and receive a payment. So to provide proper incentives in this richer setting, a proper scoring rule doesn't suffice.

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 why VCG is not applicable in this setting, compare our mechanism with that of Mezzetti in Section 4.1 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 $N = \{1, \dots, n\}$, i.e., $N_p = \{0\} \cup N$. 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 which is 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 common knowledge. However, the quality $q_i \in [0, 1]$ is private information to the labeler, and that constitutes the type set of agent i. In this setting, center does not hold any private information.

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 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* agent i. 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 [14]. The utilities 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.

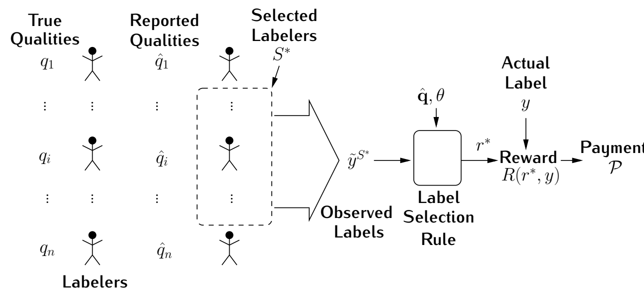

Figure 1: Illustration of the mechanism design problem

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,*

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

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,*

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

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.

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

Let us denote, $Q_\theta := \max_r \mathbb{E}_{y|\theta} R(r, y)$. This represents the maximum valuation that can be got by center without hiring any labeler. In the above equation,

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

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

$\mathbb{P}\left(\tilde{y}^S|y, q\right) = \prod_i q_i^{\mathbb{I}(\tilde{y}^i - y)} (1 - q_i)^{1 - \mathbb{I}(\tilde{y}^i - y)}$, 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,*

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

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 true profiles $q = (q_i, q_{-i})$, and for all i and $\hat{q}_i$,*

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

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) [16].

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 the agents, i.e.,*

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

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 the net monetary transfer in the system is non-positive.*

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

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

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 $= \langle S^*, r^*, \mathcal{P}^* \rangle$.

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 [11].

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

$$S^*(\hat{q}) \in \arg\max_{S \subseteq N} Q(S, \hat{q}), \quad (7)$$

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 this, 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,*

$$r^*(\tilde{y}^{S^*(\hat{q})}(q), \hat{q}) \in \arg\max_r \sum_y \mathbb{P}(y|\tilde{y}^{S^*(\hat{q})}(q), \hat{q}) R(r, y) \,. \quad (8)$$

Using the above setup, we define the payment rule as follows.

Definition 9 (Payment Rule)

$$\begin{aligned} t_i &= \mathcal{P}_i^*(S^*(\hat{q}), \tilde{y}^{S^*(\hat{q})}(q), \hat{q}, y) \\ &= \begin{cases} \alpha \times [W(r^*, S^*(\hat{q}), \hat{q}, y|q) \\ \quad -W_{-i}(r^*, S^*(\hat{q}_{-i}), \hat{q}_{-i}, y|q_{-i})] + \ c_i, & \textit{if } i \in S^*(\hat{q}) \\ 0, & \textit{otherwise} \end{cases} \\ t_0 &= -\sum_{i \in S^*(\hat{q})} \mathcal{P}_i^*(S^*(\hat{q}), \tilde{y}^{S^*(\hat{q})}(q), \hat{q}, y) \end{aligned} \quad (9)$$

This payment rule makes labelers partners in the center's venture by paying out a fraction $\alpha > 0$ of i's marginal contribution. Theorem 3 shows how the choice of α becomes crucial to ensure EPIR. Algorithm 1 shows the steps of QUEST using pseudo-code.

4. PROPERTIES OF QUEST

The proposed mechanism satisfies several important properties given by the following theorems. We denote $q = (q_i, q_{-i})$ to be the vector of true qualities of the agents. The following observation on the allocative efficiency and budget balance of QUEST follows from the definitions.

Observation 1 QUEST *is AE and SBB.*

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;

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

The proof is given in Appendix.

Theorem 2 (EPIR for all labelers) QUEST *is EPIR for all the labelers.*

PROOF. 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,

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

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, $Q(S^*(q), q) = \mathbb{E}_{\tilde{y}^{S^*},y|q}W^*$.

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

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,

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

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. □

Due to Theorem 2, we can now treat quality reports as truthful. 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 3 (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 be 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 the center.*

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

Recall the parameter α above from the payment rule given by Definition 9. We will prove this theorem via the following lemma on the monotonicity of Q.

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

The proof is given in Appendix.

Remark: 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 3: Case 1: $Q_\theta \geq 0$. By definition, $Q(S^*(q), q) \geq Q_\theta$, via Equations (3) and (7). Now, if $S^*(q) = \varnothing$, 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 7 and 8), 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 $S^*(q) \neq \varnothing$, then the center's payoff is given by,

$$\mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\Big[R(r^*(\tilde{y}^{S^*(q)}, q, \theta), y) - \sum_{i \in S^*(q)} \mathcal{P}_i^*\Big]$$
$$= \mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\Big[R(r^*(\tilde{y}^{S^*(q)}, q, \theta), y) - \sum_{i \in S^*(q)} c_i$$
$$- \alpha|S^*(q)|W^* + \alpha \sum_{i \in S^*(q)} W^*_{-i}\Big]$$
$$= Q(S^*(q), q) - \alpha|S^*(q)|Q(S^*(q), q) + \alpha\mathbb{E}_{y,\tilde{y}^{S^*(q)}|q}\Big[\sum_{i \in S^*(q)} W^*_{-i}\Big]$$
$$\geq Q(S^*(q), q) - \alpha|S^*(q)|Q(S^*(q), q) + \alpha|S^*(q)|Q_\theta \quad (10)$$

The first equality comes by substituting Equation (9), 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 (10) 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, $S^*(q) \neq \varnothing$ when $q \succeq q_t$. Then one can show using Lemma 1 that the last term in Equation (10) 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. ■

Theorem 4 (Unachievability Result) *If the problem instance (R, c, θ, q_t) is such that the labeler quality vector $q \succeq q_t \succeq q_{0.5}$, but the sufficiency condition of Theorem 3 is violated, then no mechanism can satisfy AE, EPIC, EPIR, and SBB.*

Proof: Let us assume there exists a mechanism $M = \langle S^M, r^M, \mathcal{P}^M \rangle$ that 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 (10) for the mechanism M as follows.

$$\mathbb{E}_{y,\tilde{y}^{S^M(q)}|q}\Big[R(r^M(\tilde{y}^{S^M(q)}, q, \theta), y) - \sum_{i \in S^M(q)} \mathcal{P}_i^M\Big]$$

$$= \mathbb{E}_{y,\tilde{y}^{S^M(q)}|q}\Big[R(r^M(\tilde{y}^{S^M(q)}, q, \theta), y) - \sum_{i \in S^M(q)} c_i + \sum_{i \in S^M(q)} c_i - \sum_{i \in S^M(q)} \mathcal{P}_i^M\Big]$$

$$= Q(S^*(q), q) + \sum_{i \in S^M(q)} (c_i - \mathcal{P}_i^M)$$

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 3 is violated, it implies, $Q(S^*(q), q) = \epsilon < 0$. The $\mathcal{P}_i^M$ term indicates the payment to labeler i. For M to be EPIR for the labelers, $\mathcal{P}_i^M - c_i \geq 0$, for all i. Therefore, $\sum_{i \in S^M(q)}(\mathcal{P}_i^M - c_i) \geq 0$. For M to be EPIR for the center, $Q(S^*(q), q) + \sum_{i \in S^M(q)}(c_i - \mathcal{P}_i^M) \geq 0$, which implies, $\sum_{i \in S^M(q)}(\mathcal{P}_i^M - c_i) \leq Q(S^M(q), q) < 0$, which is a contradiction. ■

4.1 Comparison with Mezzetti's Mechanism

Let us compare QUEST vis-à-vis the classic mechanism given by Mezzetti [17] (we will call this MZT) which too is EPIC and AE in the interdependent value setting. Due to Theorem 4, it is sufficient for us to compare when sufficient condition of the Theorem 3 is met. 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 $(y, \tilde{y}^{S^*(q)})$. We now 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,

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

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

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

The inequality comes from the sufficient condition of Theorem 3. 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 $(y, \tilde{y}^{S^*(q)})$.

4.2 Comparison with VCPM Mechanism

A drawback of Mezzetti's mechanism is weak IC reports of the realized payoffs in second round. Nath and Zoeter [20] address this issue by proposing the VCPM mechanism which is applicable when the outcome is a subset allocation. VCPM also assures IR when free to choose pivot term h_i is chosen to be W_{-i}. In the setting considered, as valuations of all the agents are revealed, the second stage report is truthful and so the value consistency term g is vacuously 0. Thus, with the choice of g and h as discussed above, the expectation of monetary transfer w.r.t. $(y, \tilde{y}^{S^*(q)})$ is given by

$$nQ(S^*(q), q) - \sum_{i=1}^{n} Q(S^*(q_{-i}), q_{-i}) \geq 0.$$

So, there is a budget deficit when VCPM is used and so VCPM fails to be budget balanced.

5. DISCUSSION AND FUTURE WORK

In this paper, we studied the problem of a decision maker wanting to hire experts to provide more information about a problem before a decision is made. This paper introduced a novel mechanism, QUEST, to coordinate the hiring of experts by such a decision maker. The mechanism is allocatively efficient, ex-post incentive compatible, ex-post individually rational for labelers, and strict budget balanced. The ex-post individual rationality for the decision maker is also achieved under a weak assumption that the social welfare is non-negative.

The mechanism makes the experts partner in the decision maker's venture. In addition to the compensation of the labeling costs, a part of the decision maker's return is redistributed to the selected experts (or in the case of an unfortunate roll of the dice, a part of the cost). Since the return to the decision maker depends on the quality of the experts, standard mechanisms such as VCG that rely on the independence between valuations cannot be used. QUEST leverages special properties of the outsourcing setting such that it ensures strong budget balance, whereas the classic mechanism for general dependent value settings [17] doesn't.

QUEST provides a mechanism with several strong properties for a class of important outsourcing problems. As discussed in the introduction, the studied setting is applicable in several real-world scenarios. Furthermore, there are many clear directions for future work by relaxing assumptions.

If the true label y is not observed before making payments, a weaker solution concept and ideas from peer-prediction [18] can possibly be used to base payments on how well the reported labels correspond to the reports made by others. The assumption that errors are made independently can be a good enough approximation if experts use very different techniques (e.g. different types of scans in a hospital). However it is often a dangerous assumption to make. Since individual experts are not likely to have perfect knowledge about correlations between errors, such a scenario would require a form of learning, to learn such correlations from data. It is interesting to relax the assumption of the critical event being binary, as well as relaxing the assumption of costs being common knowledge. Such more general settings lead to multi-dimensional mechanism design problems. For practical applications, even with a small set of labelers, this would require reasonable approximations.

Platforms such as Mechanical Turk, oDesk, etc., make it easier to outsource and crowdsource tasks. We expect that QUEST and its future extensions can provide a fundamental basis to balance the cost and final quality of outsourced tasks, thereby increasing the usefulness of such platforms.

6. ACKNOWLEDGEMENTS

This work is part of a collaborative project between Xerox Research and Indian Institute of Science.

References

[1] M. Babaioff, M. Feldman, and N. Nisan. Combinatorial agency. In *EC*, pages 18–28. ACM, 2006.

[2] Y. Bachrach and J. S. Rosenschein. Incentives in effort games. In *AAMAS*, pages 1557–1560, 2008.

[3] G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.

[4] Y. Chen and I. Kash. Information elicitation for decision making. In *AAMAS*, pages 175–182, 2011.

[5] V. Conitzer. Prediction markets, mechanism design, and cooperative game theory. In *UAI*, pages 101–108. AUAI Press, 2009.

[6] X. Gao, Y. Bachrach, P. Key, and T. Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *AAAI*, pages 38–44, 2012.

[7] E. Gerding, S. Stein, K. Larson, A. Rogers, and N. R. Jennings. Scalable mechanism design for the procurement of services with uncertain durations. In *AAMAS*, pages 649–656, 2010.

[8] E. H. Gerding, K. Larson, A. Rogers, and N. R. Jennings. Mechanism design for task procurement with flexible quality of service. In *Service-Oriented Computing: Agents, Semantics, and Engineering*, pages 12–23. Springer, 2009.

[9] C. Ho and J. Vaughan. Online task assignment in crowdsourcing markets. In *AAAI*, pages 45–51, 2012.

[10] S. Jain, Y. Chen, and D. C. Parkes. Designing incentives for online question-and-answer forums. *Games and Economic Behavior*, 2012.

[11] P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5), 2001.

[12] R. Jurca, B. Faltings, and W. Binder. Reliable qos monitoring based on client feedback. In *WWW*, pages 1003–1012, 2007.

[13] E. Kamar and E. Horvitz. Incentives for truthful reporting in crowdsourcing. In *AAMAS*, pages 1329–1330, 2012.

[14] V. Krishna. *Auction Theory*. Academic Press, 2009.

[15] C. Lin, Mausam, and D. Weld. Dynamically switching between synergistic workflows for crowdsourcing. In *AAAI*, 2012.

[16] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Micoreconomic Theory*. Oxford University Press, 1995.

[17] C. Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, 72(5):1617–1626, 2004.

[18] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[19] 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.

[20] S. Nath and O. Zoeter. A strict ex-post incentive compatible mechanism for interdependent valuations. *Economics Letters*, 121(2):321 – 325, 2013. ISSN 0165-1765.

[21] 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. *JAIR*, 35(1):119–159, 2009.

[22] R. Ray, R. D. Vallam, and Y. Narahari. Eliciting high quality feedback from crowdsourced tree networks using continuous scoring rules. In *AAMAS*, pages 279–286, 2013.

[23] Y. Sakurai, T. Okimoto, M. Oka, M. Shinoda, and M. Yokoo. Quality-control mechanism utilizing worker's confidence for crowdsourced tasks. In *AAMAS*, pages 1347–1348, 2013.

[24] Y. Singer and M. Mittal. Pricing tasks in online labor markets. In *Proceedings of Workshop on Human Computation (AAAI 2011)*, 2011.

[25] 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.

[26] A. Zohar and J. S. Rosenschein. Robust mechanisms for information elicitation. In *AAMAS*, pages 1202–1204. ACM, 2006.

APPENDIX

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 (7)). Hence we need to prove only the first inequality. In fact, we show that the inequality holds for any S, i.e., $Q(S, q) \leq Q(S, \tilde{q})$. For $S = \varnothing$ the result is true vacuously.

Case 1: $|S| = 1$. 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,

$$Q(S, q) = \sum_{\tilde{y} \in \{0,1\}} \max_{r \in \{0,1\}} \Big[\sum_{y \in \{0,1\}} R_{ry} \mathbb{P}(\tilde{y}|y, q, \theta) \mathbb{P}(y|\theta) \Big] - \sum_{i \in S} c_i. \tag{11}$$

The cost term on the RHS of Equation (11) appears in both $Q(S, q)$ and $Q(S, \tilde{q})$. Hence, WLOG assume $c_i = 0$.

$$\begin{aligned} Q(S, q) &= \max_{r_1 \in \{0,1\}} \{\underbrace{R_{r_1 0}(1-\theta)q + R_{r_1 1}(1-q)\theta}_{=:f(r_1,q)}\} \\ &+ \max_{r_2 \in \{0,1\}} \{\underbrace{R_{r_2 0}(1-q)(1-\theta) + R_{r_2 1} q\theta}_{=:g(r_2,q)}\} \\ &= \max_{(r_1,r_2) \in \{0,1\}^2} (f(r_1, q) + g(r_2, q)) \\ &= \max\{R_1(q), R_2(q), R_3(q), R_4(q)\}, \end{aligned} \tag{12}$$

where,

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

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 $b = R_{00}(1-\theta) + R_{11}\theta$. As R_1 through R_4 are affine functions of q, their maximum given by Equation (12) is convex in q. The lines $R_2(q)$, $R_3(q)$ intersect at $(q = 0.5\ , d)$, as shown in Figure 2 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 get the desired result.

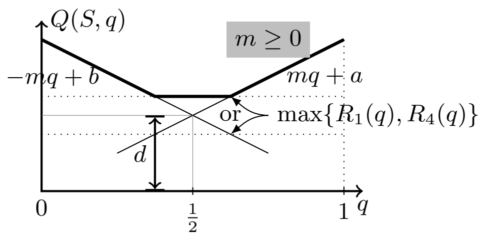

Figure 2: Expected Welfare versus q when $|S| = 1$

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. The expected social welfare is,

$$Q(q_i, q_{-i}, S) = \sum_{\tilde{y}^{S-i}} \Big\{ \sum_{\tilde{y}_i \in \{0,1\}} \max_{r \in \{0,1\}} \Big[\sum_{y \in \{0,1\}} \Big(R_{ry} \mathbb{P}(\tilde{y}^{S-i} | y, q_{-i}, \theta) \Big) \mathbb{P}(\tilde{y}_i | y, q_i, \theta) \mathbb{P}(y|\theta) \Big] \Big\} - \sum_{i \in S} c_i.$$

The term in curly braces for a fixed $\tilde{y}^{S-i}$ resembles Equation (11) and hence is non-decreasing in $q_i \in [0.5, 1]$ by a similar argument as used when $|S| = 1$. Since the total welfare is the sum of such functions, the result follows. ■

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. 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 $q = (q_i, q_{-i})$ and reported profile $\hat{q} = (\hat{q}_i, q_{-i})$. 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.

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

The third inequality comes from $(S_2 \cup S_4 \supseteq S^*_{-i}(q_{-i}))$. ■

Lemma 3 *For any $S \subseteq N$, 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:

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

Write $r_1 = r^*(\tilde{y}^S, (q_i, q_{-i})), y)$, $r_2 = r^*(\tilde{y}^S, (\hat{q}_i, \hat{q}_{-i})), y)$.

$$\begin{aligned}
&\mathbb{E}_{\tilde{y}^S, y|q} W^*(q_i, q_{-i}) - \mathbb{E}_{\tilde{y}^S, y|q} W^*(\hat{q}_i, \hat{q}_{-i}) \\
&= \mathbb{E}_{\tilde{y}^S|q} \Big[\textstyle\sum_y \mathbb{P}\left(y|\tilde{y}^S, q\right) R(r_1, y) - \sum_y \mathbb{P}\left(y|\tilde{y}^S, q\right) R(r_2, y) \Big] \\
&\geq 0. \quad \text{(from Definition 8)}
\end{aligned}$$
■

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:

$$\begin{aligned}
&\mathbb{E}_{\tilde{y}^{S_2}, y|q} W^*(q_i, q_{-i}) \\
&= \mathbb{E}_{\tilde{y}^{S_2}|q} \mathbb{E}_{y|\tilde{y}^{S_2}, q} \Big[R(r^*(\tilde{y}^{S_2}, q), y) - \sum_{j \in S_2} c_j \Big] \\
&\geq \mathbb{E}_{\tilde{y}^{S'}|q} \mathbb{E}_{y|\tilde{y}^{S'}, q} \Big[R(r^*(\tilde{y}^{S'}, q), y) - \sum_{j \in S'} c_j \Big] (\because \textit{Definition } 7) \\
&\geq \mathbb{E}_{\tilde{y}^{S'}|q} \mathbb{E}_{y|\tilde{y}^{S'}, q} \Big[R(r^*(\tilde{y}^{S'}, (\hat{q}_i, q_{-i})), y) - \sum_{j \in S'} c_j \Big] (\because \textit{Lemma } 3)
\end{aligned}$$

The last inequality holds true even for $S' = S_1$. ■

Proof of Theorem 1: The payment under QUEST is given by Equation (9). 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 $S_1 = S^*(\hat{q}_i, q_{-i}), S_2 = S^*(q_i, q_{-i})$ which is same as

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

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. ■