

FaRM: Fair Reward Mechanism for Information Aggregation in Spontaneous Localized Settings

by

Moin Hussain Moti, Dimitris Chatzopoulos, Pan Hui, Sujit Prakash Gujar

in

*28th International Joint Conference on Artificial Intelligence
(IJCAI-2019)*

Macao, China

Report No: IIIT/TR/2019/-1



Centre for Visual Information Technology
International Institute of Information Technology
Hyderabad - 500 032, INDIA
August 2019

FaRM: Fair Reward Mechanism for Information Aggregation in Spontaneous Localized Settings

Moin Hussain Moti¹, Dimitris Chatzopoulos², Pan Hui^{2,3} and Sujit Gujar¹

¹International Institute of Information Technology Hyderabad, India

²The Hong Kong University of Science and Technology, Hong Kong

³University of Helsinki, Finland

moin.moti@research.iiit.ac.in, {dcab, panhui}@cse.ust.hk, sujit.gujar@iiit.ac.in

Abstract

Although peer prediction markets are widely used in crowdsourcing to aggregate information from agents, they often fail to reward the participating agents equitably. Honest agents can be wrongly penalized if randomly paired with dishonest ones. In this work, we introduce *selective* and *cumulative* fairness. We characterize a mechanism as fair if it satisfies both notions and present FaRM, a representative mechanism we designed. FaRM is a Nash incentive mechanism that focuses on information aggregation for spontaneous local activities which are accessible to a limited number of agents without assuming any prior knowledge of the event. All the agents in the vicinity observe the same information. FaRM uses (i) a *report strength score* to remove the risk of random pairing with dishonest reporters, (ii) a *consistency score* to measure an agent's history of accurate reports and distinguish valuable reports, (iii) a *reliability score* to estimate the probability of an agent to collude with nearby agents and prevents agents from getting swayed, and (iv) a *location robustness score* to filter agents who try to participate without being present in the considered setting. Together, report strength, consistency, and reliability represent a fair reward given to agents based on their reports.

1 Introduction

Information aggregation systems have been employed to gather data for a wide variety of applications. A few examples are pollution measurements in cities, product reviews on e-commerce websites, feedback on mobile applications on application markets, working hours and other details of a business, etc. Depending on the type of the information of interest and the available time for reporting, the devices that can collect the data and share them with the information aggregation systems vary. The focus of this work is on location specific queries that need to be answered within a short time in the so-called *spontaneous localized settings*. Such settings are challenging to handle since information regarding the *query subjects* can only be collected by mobile agents that are located in an area and, possibly, by pre-deployed static sensors.

For example, consider a food deli serving various snacks to be the subject of interest and let the query be the availability of a certain snack. Since the food deli is only relevant to the neighborhood and the queries can be very specific, Internet services, like Neighbourly¹, need to employ mobile agents.

In such scenarios, crowdsensing mechanisms can employ agents to collect information with spatiotemporal characteristics. The main problem of aggregating information using crowdsensing mechanisms is to ensure that the agents are truthful. A dishonest agent, e.g., may have some grudge against an agency and therefore reports its service as a lousy experience irrespective of the services of the agency. Agents can be categorized as honest, free-riders or misbehaving. Honest agents are expected to put effort into assessing the query subject and reporting truthful information. Free-riders do not put effort into the assessment and try to earn free credits by guessing the information requested for the subject. Misbehaving agents try to trick the mechanism by colluding among themselves and reporting false information. Information aggregation mechanisms must be robust against free-riders and misbehaving agents. For that, they incorporate incentive mechanisms to engage agents to submit accurate reports. Rational agents target on maximizing their utility (payments, reputation, etc.) while not sacrificing a substantial amount of resources (battery lifetime, dataplan usage, etc).

For traditional settings, there exist many popular peer prediction mechanisms inspired by the work of Prelec [2004], who proposed the *Bayesian Truth Serum* and Miller *et al.* [2005], who developed *Peer-to-Peer Prediction*. These mechanisms rate an agent's report against the report of another randomly selected agent and if they match, both agents are rewarded based on the prior probability of their answer. Unfortunately, when considering spontaneous localized settings, it is probable to not have readily available prior knowledge about the query subject. Also, there is a chance of matching an honest agent with a misbehaving one, which results in an unfair evaluation. Although these mechanisms are Bayesian Incentive Compatible, they do not ensure *fairness*. Goel and Faltings [2019], recently proposed a fair mechanism for information aggregation. However, they consider queries where the answers are known and hence, their mechanism is not applicable to spontaneous localized settings.

¹neighbourly.google.com/about

In this paper, we focus on information aggregation in spontaneous localized settings and propose a suitable mechanism. Spontaneous localized settings describe information aggregation requests in local environments about subjects which are accessible by only a small set of people. For example, the food deli scenario mentioned above, the queue length in a local coffee shop, and others. Any query in these settings refers to a query subject (QS) and is valid only for a bounded time. Also, there does not exist a prior distribution for the possible answers about QS. However, in these settings, the observation environment is assumed to be the same for all the agents in the vicinity of QS. This is similar to the output agreement (OA) setting where two agents observe the same signal and are scored based on the degree of similarity in their reports. OA setting was first introduced by Von Ahn and Dabbish [2004] and later studied experimentally by Robertson *et al.* [2009], and Huang and Fu [2012]. Waggoner and Chen [2014] were the first to provide a theoretical analysis of the output agreement setting. We introduce spontaneous localized settings in more detail in Section 3.

We leverage on the characteristic of a common observation environment to propose the *fair reward mechanism (FaRM)*, a mechanism that incentivizes agents to report honestly, marginalizes free-riders and misbehaving agents and ensures fairness in reward distribution without the need for a known prior distribution. FaRM is a novel Nash incentive compatible mechanism for information aggregation in spontaneous localized settings. Additionally, we introduce the notions of *selective fairness* and *cumulative fairness*. A mechanism satisfies selective fairness in cases where for two different agents that have submitted the same reports, they are evaluated similarly. Thus, all honest agents receive the same reward if they submitted the same report. We propose *report strength* as one component of the reward scheme integrated in FaRM to ensure selective fairness. Existing peer-prediction mechanisms consider every agent’s report equally valuable. FaRM, in contrast to this assumption, considers more valuable the agents with long history of submitting truthful reports compared to agents with lousy history of reporting due to their behavioral tendency to maintain the streak. Cumulative fairness mandates that a fair mechanism should acknowledge an agent’s history of submission and consider her consistency as part of final reward. We propose *consistency score* as the second component of FaRM’s reward scheme which defines an agent’s consistency so far in the system and ensures cumulative fairness. This also helps FaRM minimize the noise caused by false reports by considering the value of a report based on its reporter’s consistency. FaRM is designed to ensure selective fairness and cumulative fairness.

Since spontaneous localized settings focus information aggregation related to local query subjects, many participating agents will be located around the vicinity of the subject. Hence, it becomes imperative that the mechanism needs to be collusion resistant. To ensure this, we propose the *reliability score* to estimate agents probability of collusion with nearby agents. This score is designed to prevent the agents from colluding with nearby agents since it directly affects their final reward. In summary, our contributions are:

1. FaRM is Nash incentive compatible mechanism that en-

ures truthful reporting from agents and guarantees non-negative rewards to all agents (Theorem 4.15).

2. FaRM incorporates a payment scheme composed of: *report strength*, *consistency score* and *reliability score*.
3. FaRM uses report strength to ensure selective fairness by distributing the same immediate reward to all the agents with the same report (Section 4.1).
4. FaRM employs consistency score as an estimate of agent’s consistency in the system so far. It gets updated after each submission. It ensures cumulative fairness by rewarding the agent for her streak of honest submissions to the system (Section 4.2).
5. FaRM integrates reliability score as an estimate of agent’s collusion with nearby agents (Section 4.3).
6. FaRM filters reports from agents who are not close to the query subject by incorporating a *location robustness score* (Section 4.4).

2 Related Work

Information markets (IM), also known as prediction markets, are exchange-traded markets where agents with different beliefs trade on the outcome of events. The market prices act as an estimate of the probability of the occurring of an event. Therefore, prediction markets can be used to aggregate beliefs. Agents trade assets each of whose value is tied to a particular observation. We discuss several proposals in information aggregation related to the design of FaRM.

In their seminal work, Miller *et al.* [2005] use an IM setting to elicit information from the agents. The asset, in this case, was agents’ reports but the agents traded the reports with the system instead of other agents. This centralized system is an automated market maker which ensures the availability of the market for trade. Agents are rewarded using *strictly proper scoring rules*, which ascertain that the best response for each agent is to report truthfully. However, their mechanism relies on the *common prior assumption*, which is not applicable in localized settings. Prelec [2004] propose a mechanism which does not require knowledge of common prior information but is applicable on large number of agents. Lambert and Shoham [2008] also propose a mechanism does not require common prior information but in the equilibrium, agents are indifferent between being honest and misreporting.

Witkowski and Parkes [2011; 2012a] present mechanisms that do not assume any common prior information, do not need a large number of agents to function and are robust to private beliefs of agents but suffer from temporal separation of reports. Their improved mechanism, named Robust Bayesian Truth Serum (RBTS) [2012b], does not require temporal separation between two reports but works only for binary outcomes. Radanovic and Faltings [2013; 2014] improve RBTS to make it compatible with non-binary and continuous outcomes. Similarly, Zhang and Chen [2014] propose a mechanism which does not assume prior known information. Riley [2014] presents a mechanism which is minimalistic (requires only one report) under the assumption that all the agents with the same outcome have the same posterior

expectations. A common characteristic among the aforementioned mechanisms is the use of *Strictly Proper Scoring Rules* (SPSR) for distributing agents’ rewards. SPSR, however, do not guarantee fairness. The incentive compatible mechanism by Jurca and Faltings [2003] is not employing a SPSR and is designed for interactive reputation markets. The weakness of this mechanism is that it is susceptible to collusion and does not consider agents’ private beliefs. A new Robust version by Jurca and Faltings [2007] of the earlier mechanism makes it robust against private beliefs of agents to some extent but it assumes that common prior is known to the mechanism. Dasgupta and Ghosh [2013] also constructed a strong incentive compatible mechanism but it can only be applied to binary settings. Faltings *et al.* [2014] propose *Peer Truth Serum* (PTS) which is a minimalistic mechanism but assumes a prior belief model. Radanovic and Faltings [2015] improved PTS by eliminating the dependency on prior belief model, they called it *Logarithmic Peer Truth Serum* (LPTS). LPTS assumes that all agents observe same signal and therefore were able to come up with a minimalistic mechanism. Their mechanism is able to tackle colluders and focuses on effort elicitation to prevent freeloading by agents. Radanovic *et al.* [2016] later on propose an optimized version of PTS focusing on in multi-task settings.

FaRM, similarly to LPTS, assumes that all agents observe the same signal and does not require any common knowledge assumptions. It is a minimalistic and collusion resistant mechanism that prevents free-riding and guarantees non-negative rewards to the users.

3 Spontaneous Localized Settings

Existing mechanisms for information aggregation have been designed for the two following settings: (i) online reviews about products and services, and (ii) community sensing regarding prevalent topics like pollution level reporting over a wide region for a considerable duration of time. Additionally, to the best of our knowledge, existing literature on information markets has not discussed information aggregation of location-specific spontaneous tasks on local subjects. Most papers study an *independent value* model where the agents can independently observe varying signals. In this model, the truthful report of agents does not need to be the same. However, spontaneous localized settings follow a *common value* model where all agents observe the same signal.

More formally, considering a query subject QS , a set of nearby agents \mathcal{A} that choose to participate and asses QS , and a budget $B \in \mathbb{R}$, we want to estimate a function f . For example, QS can be a food deli and f the number of eateries available for purchase. S denotes the signal space for the query and is allowed to be non-binary. Every query in spontaneous localized settings focuses on only one task. The observation area is considered to be in the close vicinity of the QS . Since the query requires to be answered spontaneously, we assume the QS to not change its status for the time period the query remains active. Every agent $i \in \mathcal{A}$ observes the same signal $s \in S$ and reports a signal $r_i \in S$. r_i can be different to s . We define a general utility function for agent i as $\hat{u}_i(r_i, r_{-i})$ where r_i represents agent i ’s reported signal

Symbol	Description
QS	Query Subject
\mathcal{A}	Set of agents participating in the process.
S	Set of all the signals.
s	Observed signal by all agents.
r_i	Reported signal by agent i .
u_i	Reward of agent i .
B	Budget per crowdsensing query.
l_i	Context of a agent.
α_i	Consistency score of agent i .
β_i	Reliability score of agent i .
γ_i	Location robustness score of agent i .
\mathcal{I}_i	Nearby agents (internal peers) of agent i .
\mathcal{E}_i	External peers of agent i .
$\Phi()$	Strength function for a report type.
φ_1	Strength of strongest report.
φ_2	Strength of second strongest report.

Table 1: Notation Table.

and $r_{-i} (= (r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_{|\mathcal{A}|}))$ represents all other agents’ reported signal. Note that we assume $|\mathcal{A}| \geq 3$, i.e. at least 3 agents are participating for a particular query. Also, the sum of total utility (u_i) of all agents should not exceed the budget of the query i.e. $\sum u_i \leq B$.

In the following sections, we discuss the properties of FaRM. Due to space limitations we provide all the proofs (i.e., proofs of Theorem 4.15, Lemma 4.4, Lemma 4.8, Lemma 4.13, Proposition 4.16, Claim 4.6, Claim 4.7, Claim 4.10, Claim 4.12 and Observation 4.3) in an extended version [Moti *et al.*, 2019].

4 FaRM: Fair Reward Mechanism

The payment scheme designed for FaRM is the product of three sub-utility functions, namely *report strength*, *consistency score* and *reliability score*. The report strength is calculated at the end of every report collection process based on the submitted reports. The consistency score is updated for each agent after every participation. For a new agent i , her consistency score, (α_i), is initialized to 0. The reliability score for each agent which is an estimate of an agent’s collusion with nearby agents. The consistency score motivates agents to participate and report truthfully while the reliability score prevents agents from colluding with her neighboring agents. We discuss these metrics in Section 4.1, Section 4.2 and Section 4.3 respectively where we analyze and prove that truthful reporting forms pure strategy Nash equilibrium (PSNE) for each metric.

Definition 4.1 (Pure Strategy Nash Equilibrium). *Given a strategic form game $\Gamma = \langle \mathcal{A}, (\Pi_i) : S \rightarrow S, (\hat{u}_i) \rangle$, the strategy profile $\pi^* = (\pi_1^*, \pi_2^*, \dots, \pi_n^*)$ is called a PSNE of Γ if*

$$\pi_i^* = \operatorname{argmax}_{\pi_i \in \Pi_i} \hat{u}_i(\pi_i, \pi_{-i}^*) \quad \forall i \in N$$

That is, each player’s Nash equilibrium strategy is a best response to the Nash equilibrium strategies of the other players [Narahari, 2014].

Definition 4.2 (Nash Incentive Compatibility). A social choice function $u : S_1 \times \dots \times S_n$ is said to be Nash Incentive Compatible (or truthfully implementable in Nash equilibrium) if the direct revelation mechanism $\mathcal{D} = ((S_i)_{i \in \mathcal{A}}, \hat{u}(\cdot))$ has a pure strategy Nash equilibrium $\pi^*(\cdot) = (\pi_1^*(\cdot), \dots, \pi_n^*(\cdot))$ in which $\pi_i^*(s_i) = s_i \forall s_i \in S_i, \forall i \in \mathcal{A}$.

That is, directly asking the agents to report their types and using this information in $\hat{u}(\cdot)$ to get the social outcome will solve both the problems, namely, preference elicitation and preference aggregation [Garg et al., 2008; Narahari, 2014].

4.1 Report Strength

We define the strength of a report as the count of agents who have reported the same signal as agent i . If there is a wide spectrum of possible signals, we discretize the range in buckets and consider two signals as equal, if they belong in the same bucket. Report strength is the first of the three scores we compute for a report. It only depends on the current performance of the agent and hence can be considered as the immediate reward for an agent in a particular query.

Computation of Report Strength

Every $s \in S$ represents a report type and each report type receives certain number of reports. Let Φ be a function that counts the number of reports of the same type. Φ_{r_i} is defined as the strength of agent i 's report by measuring how many agents are in agreement with agent i .

$$\Phi(r_i) = \sum_{j \in \mathcal{A}} 1_{r_i=r_j} \quad (1)$$

$\Phi(\cdot)$ is a function on a generic report type and is independent of the agent. Thus, for computational efficiency, strength score should be pre-computed for every signal which can be done by maintaining a counter for every signal while iterating over all agent reports. That is for n agents, total computational complexity is $O(n)$. Individual strength score for any agent can then be obtained by just referring to the strength score of the corresponding signal reported.

Observation 4.3. The report strength of the report of an agent is always positive.

Lemma 4.4. Consider a game induced by u^Φ where $u^\Phi = \Phi$ is the sub-utility function corresponding to the report strength (Φ); and let \hat{s} be the observed signal by all agents and r_i be the report submitted by any agent i . Then under Nash equilibrium,

$$u_i^\Phi(r_i = \hat{s}, r_{-i} = \hat{s}) > u_i^\Phi(r_i \neq \hat{s}, r_{-i} = \hat{s}) \forall r_i \in S, \forall i \in \mathcal{A}$$

That is, if all other agents were to report truthfully, the best response for agent i in order to maximize her sub-utility u_i^Φ is also to report truthfully.

Selective Fairness

Information aggregation involves evaluating reported signals from agents. It is difficult to evaluate an agent since there is no ground truth available in most scenarios. Existing peer prediction based mechanisms discussed in Section 2 rate an agent's report against the report of a random peer agent p_i . This, however, exposes agent i to unfair evaluation because

if the peer agent p_i is dishonest or misbehaving and submits false report, it will reflect in worse reward for agent i even if she submits truthful report. [Dwork et al., 2012] state that an algorithm is fair if it generates similar results for agents with similar attributes. This is not possible with peer prediction based mechanisms.

Existing mechanisms in the literature mainly discuss settings which follow an *independent-value model* where agents observe varying signals for the same query. This makes it difficult to rate all agents against equivalent reports. In our case, we follow a *common-value model* where all agents observe the same signal. This makes it easier to rate all agents against the same outcome.

Definition 4.5 (Selective Fairness). Let \mathcal{A} be a set of agents in the vicinity of subject QS and agents $i, j \in \mathcal{A}$ be any two arbitrary agents who submit two identical reports r_i and r_j such that $r_i = r_j$. The utility function of a payment scheme \tilde{u} , admits selective fairness if,

$$\tilde{u}_i(r_i, r_{-i}) = \tilde{u}_j(r_j, r_{-j}), \forall i, j \in \mathcal{A}$$

Claim 4.6. Report strength is a selectively fair computation score for agents.

4.2 Consistency Score (α)

The report of an agent with a history of reporting truthfully is more valuable to FaRM than the report of an agent with a lousy history of reporting. Therefore, instead of distributing rewards based only on the performance of agents in the last query, we keep a metric of every agent's performance so far in the mechanism and take it into consideration while paying the agents. We call this metric *consistency score*.

Consistency score depends on the strength of agent i 's report. After every report of agent i , the FaRM updates agent's consistency score. The score increases for accurate reporting and gets penalized for inaccurate reporting. It should be noted, however, that the accuracy is relative since the ground truth is not known and the highest reported signal is used as a proxy for ground truth. Consistent agents receive higher rewards than less consistent ones for the same report submitted to the system because of the higher value of their report. We, therefore, motivate agents to maintain the quality of their reports in order to maintain high rewards. This also encourages agents who have a higher rate of submitting truthful reports to the system to continue participating and at the same time discourages free-riders and dishonest reporters.

Computation of Consistency Score

The consistency score is incremented with respect to the second highest reported signal (denoted by ϕ_2) and is decremented based on the highest reported signal (denoted by ϕ_1).

$$\varphi_1 = \max_{s \in S} (\Phi(s)) \quad (2)$$

$$\varphi_2 = \begin{cases} \max_{2s \in S} (\Phi(s)) & \text{if } \max_{2s \in S} (\Phi(s)) > 0 \\ \frac{\varphi_1 - 1}{\varphi_1} & \text{if } \max_{2s \in S} (\Phi(s)) = 0 \end{cases} \quad (3)$$

where $\max 2$ gives the second highest value and $\varphi_1 \geq \varphi_2 > 0$. φ_1 and φ_2 can be pre-computed before computing individual agent scores. For n agents, the complexity is $O(n)$. Only

agents with the strongest reports will have their consistency score incremented. However, it is not enough to submit one of the strongest reports, the report type must hold percentage strictly greater than other report types. Other agents will have their consistency score decremented. It is updated as follows,

$$\alpha_i^t = \begin{cases} \alpha_i^{t-1} - \frac{\alpha_i^{t-1}}{k} \times \frac{(\varphi_1 - \Phi(r_i))}{|\mathcal{A}|} & \text{if } \Phi(r_i) < \varphi_1 \\ \alpha_i^{t-1} + \frac{1 - \alpha_i^{t-1}}{k} \times \frac{(\varphi_1 - \varphi_2)}{|\mathcal{A}|} & \text{if } \Phi(r_i) = \varphi_1, \end{cases} \quad (4)$$

where $k \geq 1$. Note that k is just an arbitrary constant which can be used to tweak the reward rate. t represents the sequence of rounds agent i has participated. Once φ_1 and φ_2 are pre-computed, updating α has constant computational complexity per agent (i.e., the overall complexity is $O(n)$).

Claim 4.7. *Consistency score is bounded in the range $[0,1]$.*

Notice that the consistency score increases rapidly when it is closer to zero and it decreases fast when it is closer to one. This property motivates new agents to join the system and prevent existing agents from downgrading the quality of their reports.

Lemma 4.8. *Consider a game induced by u^α where $u^\alpha = \alpha$ is the sub-utility function corresponding to consistency score (α); and let \hat{s} be the observed signal by all agents and r_i be the report submitted by any agent i . Then considering the property of Nash equilibrium,*

$$u_i^\alpha(r_i = \hat{s}, r_{-i} = \hat{s}) > u_i^\alpha(r_i \neq \hat{s}, r_{-i} = \hat{s}) \forall r_i \in S, \forall i \in \mathcal{A}$$

That is, if all other agents were to report truthfully, the best response for agent i in order to maximize her sub-utility u_i^α is also to report truthfully.

Cumulative Fairness

Report from an agent who is consistently reporting truthful information is more valuable than the report of an agent with a lousy history of reporting.

Example: In her summer vacation, Anna joins professional cooking classes out of interest. She likes to frequently practice recipes at home and get reviews from her family and neighbors. Once she makes Turkish delights and has Barbara and Chris independently assess the delicacy. However, they submit contrasting reviews. But Anna recalls Barbara's previous reports and finds her to be misreporting on various occasions in order to make her happy, while Chris has been almost always critical but truthful in his reports. Thus, based on the history of reports of Barbara and Chris, Anna identifies Chris' reports to be more useful than that from Barbara.

It's only fair to take the value of an agent's report into consideration while distributing reward. Motivated by similar analogy, we incorporated consistency score (α) in our final reward structure.

Definition 4.9 (Cumulative Fairness). *Let \mathcal{A} be a set of agents in the vicinity of local subject QS and agents $i, j \in \mathcal{A}$ be any two arbitrary agents submitting report r_i and r_j such that $r_i = r_j$. Then the utility function $\hat{u}(\cdot)$ of an arbitrary*

payment scheme is said to admit cumulative fairness if,

$$\hat{u}_i^t(r_i = \hat{s}, r_{-i}, \hat{u}_i^{t-1}) > \hat{u}_j^t(r_j = \hat{s}, r_{-j}, \hat{u}_j^{t-1}) \\ \forall i, j \in \mathcal{A} \text{ s.t. } \hat{u}_i^{t-1} > \hat{u}_j^{t-1}$$

Claim 4.10. *Consistency score is a cumulatively fair computation score for agents.*

4.3 Reliability Score (β)

Reliability score provides incentives to agent i to not collude with her nearby agents \mathcal{I}_i . It is defined as the ratio of *external agreement* by *internal agreement*. The internal agreement is the percentage of nearby agents (formally identified as *internal peers* \mathcal{I}_i) that share the same report as agent i and the external agreement is the percentage of agents other than nearby agents (formally addressed as *external peers* \mathcal{E}_i) who share the same report as agent i .

Definition 4.11 (Internal and External Peers). *The internal peers of agent i (denoted by \mathcal{I}_i) are the ones with which i can exchange a message using a device-to-device communication protocol like Wi-Fi direct or Bluetooth. The internal peers combined with the agent's location compose her context l_i . Every agent j that is not directly accessible from i (i.e., $j \notin \mathcal{I}_i$) belongs to her external peers denoted by \mathcal{E}_i .*

Computation of Reliability Score

For agent i , her reliability score is computed as:

$$\beta_i = \left(\frac{\sum_{j \in \mathcal{E}_i} 1_{r_i=r_j}}{|\mathcal{E}_i|} \right) / \left(\frac{\sum_{j \in \mathcal{I}_i} 1_{r_i=r_j}}{|\mathcal{I}_i|} + 1 \right) \quad (5)$$

For n agents, the computational complexity for calculating β is $O(n)$ per agent.

Claim 4.12. *Reliability score is bounded in the range $[0,1]$*

Lemma 4.13. *Consider a game induced by u^β where $u^\beta = \beta$ is the sub-utility function corresponding to reliability score (β); and let \hat{s} be the observed signal by all agents and r_i be the report submitted by any agent i . Then considering the property of Nash equilibrium,*

$$u_i^\beta(r_i = \hat{s}, r_{-i} = \hat{s}) > u_i^\beta(r_i \neq \hat{s}, r_{-i} = \hat{s}) \forall r_i \in S, \forall i \in \mathcal{A}$$

That is, if all other agents were to report truthfully, the best response for agent i , in order to maximize her sub-utility u_i^β , is also to report truthfully.

Claim 4.14. *Reliability score prevents agents from colluding with nearby agents.*

Proof Sketch. Reliability score (β) can be seen as,

$$\beta = \frac{\text{external agreement}}{\text{internal agreement} + 1}$$

where external agreement is the percentage of the external peers that agree with the agent and internal agreement is the percentage of the internal peers that agree with the agent. For the sake of convenience we use variables l_i and g_i as,

$$l_i = \text{internal agreement} = \frac{1}{|\mathcal{I}_i|} \sum_{j \in \mathcal{I}_i} 1_{r_i=r_j} \\ g_i = \text{external agreement} = \frac{1}{|\mathcal{E}_i|} \sum_{j \in \mathcal{E}_i} 1_{r_i=r_j}$$

Let's consider a malicious colluding agent m . We assume agents around QS are reasonably scattered and that the majority of agents are honest. In that case, agent m will have high internal agreement and less external agreement i.e. $l_m > g_m$. Since, $\beta_m \propto g_m$ and $\beta_m \propto 1/(l_m + 1)$, the more agent m colludes with internal peers the lower her reliability score drops. In case where nearby agents are also honest, it is still in favor of agent m to submit truthful report since she will want to maintain her reporting agreement with external peers.

On the other hand, if agent m were to be honest while her internal peers collude, the outcome is $l_m < g_m$. This leads to increase in reliability score. Hence, if internal peers of agent m are colluding, it is profitable for agent m to remain honest. Thus, reliability score prevents agents from getting swayed by nearby agents (internal peers). \square

4.4 Location Robustness Score (γ)

Apart from the aforementioned scores, we introduce the location robustness score of agent i , γ_i , to detect whether the location of a mobile agent is close to the location of the query subject. Every mobile agent can exchange messages with her internal peers in order to justify that she is within a distance from the query subject [Chatzopoulos *et al.*, 2016]. In detail, considering that the context of agent i , \mathcal{I}_i is composed of her location l_i^{Loc} and her nearby agents \mathcal{I}_i , agent i can send a message to every node $j \in \mathcal{I}_i$ with l_i^{Loc} in it using a communication technology like Bluetooth or WiFi-direct. Every agent $j \in \mathcal{I}_i$ will calculate the distance with agent i , D_{ji} and will respond to the message by including her location l_j^{Loc} and the calculated distance. Agent j can sign her message with a private key in order to make sure that agent i will not alternate her estimation of the difference between the two locations. Based on these messages, agent i will calculate the robustness of her location.

$$\gamma_i = \frac{1}{|\mathcal{I}_i|} \sum_{j \in \mathcal{I}_i} 1_{D_{ji} \leq D^T} \quad (6)$$

Which is practically the fraction of their neighbors who are within a distance D^T that depends on the communication technology. For example, two mobile agents can communicate with Bluetooth if $D^T < 20$ meters and with WiFi-direct if $D^T < 50$ meters. Depending on the query subject and the predetermined allowed distance from which a mobile agent can produce a report, D^T may be smaller than the values achieved by the employed communication technology. FaRM uses location robustness to detect fraudulent agents.

Every participating mobile agent, has to submit γ_i together with her report. Given that a fraudulent agent is not able to alternate γ_i , since every message from her internal peers is signed, FaRM filters reports with low γ values.

4.5 Rewards

Once FaRM filters reports using the location robustness score, it updates consistency and reliability scores for each agent, then finally it calculates every agent's reward for that query. The location robustness score (γ) is only used for filtration and not considered in the final reward. The final reward of agent i depends on the product of $\Phi(r_i)$, α_i and β_i .

The strength of agent's report factors in the immediate reward for the respective query. The consistency score takes into account agent's tendency to report truthfully and therefore motivate a more consistent agent to participate more. The reliability score keeps the agent from colluding with other agents. An agent who has a history of colluding with nearby agents will have less motivation to participate in the system, hence preventing noise generated by misleading reports. The final utility score (u_i) is computed as:

$$u_i(r_i, r_{-i}) = \frac{u_i^\Phi(r_i, r_{-i}) \cdot u_i^\alpha(r_i, r_{-i}) \cdot u_i^\beta(r_i, r_{-i})}{|\mathcal{A}|^2} \times B, \quad (7)$$

where B is the reward budget per query. Also, since final utility is the product of all three scores, even if some closely located agents collude in order to increase report strength, they risk decreasing reliability score which can mask the benefits of collusion and result in even worse payoffs.

Theorem 4.15. *FaRM is Nash incentive compatible with guaranteed non-negative payoffs and weak budget balanced.*

Proposition 4.16. *FaRM admits selective fairness and cumulative fairness and hence is a fair mechanism.*

The proofs of Theorem 4.15, Lemma 4.4, Lemma 4.8, Lemma 4.13, Proposition 4.16, Claim 4.6, Claim 4.7, Claim 4.10, Claim 4.12 and Observation 4.3 are provided in an extended version [Moti *et al.*, 2019].

5 Conclusion

Existing information aggregation mechanisms are focused on settings relevant to mass population, e.g., online reviews about products and services, and community sensing regarding prevalent societal topics (pollution, global warming, etc.). FaRM focuses on location-specific queries of people in their everyday life. It can be challenging to obtain prior distribution models for the signal space and hence, prior-free mechanisms are preferable for these set of queries. In spontaneous localized settings all the agents are assumed to observe the same signal. FaRM leverages this property to provide fair rewards to all agents. FaRM's payment scheme consists of three sub-utility structures, namely *report strength*, *consistency score* and *reliability score*. Report strength ensures *selective fairness* for the immediate query by evaluating all agents equivalently against the same result. Consistency score ensures *cumulative fairness* by accounting for an agent's consistency in the system. The report of an agent consistently reporting truthful reports to the system is more valuable than the report of an agent with a lousy reporting history. Consistency score takes the value of a report into consideration. Therefore, it helps in preventing agents from free-riding as it can affect their consistency score negatively. Reliability score estimates an agent's collusion with the nearby agents and reflects the same in the final reward. It motivates agents to earn a higher reward by not colluding with nearby agents and instead stick with their observed signal. Together these scores ensure fair reward to the agents, not just for their immediate contribution but also for their consistent and reliable service to the system. In summary, FaRM is a Nash incentive compatible mechanism which rewards agents fairly for their services.

References

- [Chatzopoulos *et al.*, 2016] Dimitris Chatzopoulos, Sujit Gujar, Boi Faltings, and Pan Hui. Localcoin: An ad-hoc payment scheme for areas with high connectivity: Poster. In *MobiHoc '16*, pages 365–366. ACM, 2016.
- [Dasgupta and Ghosh, 2013] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330, 2013.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226. ACM, 2012.
- [Faltings *et al.*, 2014] Boi Faltings, Jason Jingshi Li, and Radu Jurca. Incentive mechanisms for community sensing. *IEEE Transactions on Computers*, 63(1):115–128, 2014.
- [Garg *et al.*, 2008] Dinesh Garg, Y Narahari, and Sujit Gujar. Foundations of mechanism design: A tutorial part 1-key concepts and classical results. *Sadhana*, 33(2):83, 2008.
- [Goel and Faltings, 2019] Naman Goel and Boi Faltings. Deep bayesian trust : A dominant and fair incentive mechanism for crowd. In *To appear in Proceedings of the 33rd AAI Conference on Artificial Intelligence*, 2019.
- [Huang and Fu, 2012] Shih-Wen Huang and Wai-Tat Fu. Systematic analysis of output agreement games: Effects of gaming environment, social interaction, and feedback. In *Workshops at the Twenty-Sixth AAI Conference on Artificial Intelligence*, 2012.
- [Jurca and Faltings, 2003] Radu Jurca and Boi Faltings. An incentive compatible reputation mechanism. In *E-Commerce, 2003. CEC 2003. IEEE International Conference on*, pages 285–292. IEEE, 2003.
- [Jurca and Faltings, 2007] Radu Jurca and Boi Faltings. Robust incentive-compatible feedback payments. In *Agent-Mediated Electronic Commerce. Automated Negotiation and Strategy Design for Electronic Markets*, pages 204–218. Springer, 2007.
- [Lambert and Shoham, 2008] Nicolas Lambert and Yoav Shoham. Truthful surveys. In *International Workshop on Internet and Network Economics*, pages 154–165. Springer, 2008.
- [Miller *et al.*, 2005] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.
- [Moti *et al.*, 2019] Moin Hussain Moti, Dimitris Chatzopoulos, Pan Hui, and Sujit Gujar. Farm: Fair reward mechanism for information aggregation in spontaneous localized settings (extended version), 2019.
- [Narahari, 2014] Yadati Narahari. *Game theory and mechanism design*, volume 4. World Scientific, 2014.
- [Prelec, 2004] Dražen Prelec. A bayesian truth serum for subjective data. *science*, 306(5695):462–466, 2004.
- [Radanovic and Faltings, 2013] Goran Radanovic and Boi Faltings. A robust bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAI Conference on Artificial Intelligence (AAAI' 13)*, pages 833–839, 2013.
- [Radanovic and Faltings, 2014] Goran Radanovic and Boi Faltings. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the 28th AAI Conference on Artificial Intelligence*, pages 770–776, 2014.
- [Radanovic and Faltings, 2015] Goran Radanovic and Boi Faltings. Incentive schemes for participatory sensing. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1081–1089. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [Radanovic *et al.*, 2016] Goran Radanovic, Boi Faltings, and Radu Jurca. Incentives for effort in crowdsourcing using the peer truth serum. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(4):48, 2016.
- [Riley, 2014] Blake Riley. Minimum truth serums with optional predictions. In *Proceedings of the 4th Workshop on Social Computing and User Generated Content (SC14)*, 2014.
- [Robertson *et al.*, 2009] Stephen Robertson, Milan Vojnovic, and Ingmar Weber. Rethinking the esp game. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, pages 3937–3942. ACM, 2009.
- [Von Ahn and Dabbish, 2004] Luis Von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 319–326. ACM, 2004.
- [Waggoner and Chen, 2014] Bo Waggoner and Yiling Chen. Output agreement mechanisms and common knowledge. In *Second AAI Conference on Human Computation and Crowdsourcing*, 2014.
- [Witkowski and Parkes, 2011] Jens Witkowski and David C Parkes. Peer prediction with private beliefs. In *Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC 2011)*, 2011.
- [Witkowski and Parkes, 2012a] Jens Witkowski and David C Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981. ACM, 2012.
- [Witkowski and Parkes, 2012b] Jens Witkowski and David C Parkes. A robust bayesian truth serum for small populations. In *AAAI*, volume 12, pages 1492–1498, 2012.
- [Zhang and Chen, 2014] Peter Zhang and Yiling Chen. Elicibility and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 245–252. International Foundation for Autonomous Agents and Multiagent Systems, 2014.