Collective Ratings for Online Communities With Strategic Users

Yu Zhang, *Student Member, IEEE*, and Mihaela van der Schaar, *Fellow, IEEE*

***Abstract*—Despite the success of emerging online communities, they face a serious practical challenge: the participating agents are strategic, and incentive mechanisms are needed to compel such agents to provide high-quality services. Traditional mechanisms based on pricing and direct reciprocity schemes are not effective in providing incentives in such communities due to their unique features: large number of agents able to perform diverse services, imperfect monitoring of agents' service quality, etc. To compel agents to provide high-quality services, we develop a novel game-theoretic framework for providing incentives using rating-based pricing schemes. In our framework, the service-providing agents are not rated individually; instead, they are divided into separate groups based on their expertise, location, etc., and are rated collectively, as a group. A collective rating is updated for each group based on the quality of service provided by all the agents appertaining to the group. Depending on whether a group of agents collectively contributes a sufficiently high level of services or not, the agents in the group are rewarded or punished through increased or decreased collective rating, which will lead to higher or lower payments they receive in the future. We systematically analyze how the group size and the rating scheme affect the community designer's revenue as well as the social welfare of the agents and, based on this analysis. We design optimal rating protocols and show that these protocols can significantly improve the social welfare of the community as compared to a variety of alternative incentive mechanisms.**

***Index Terms*—Collective rating, imperfect monitoring, online community, rating-based pricing, repeated games.**

I. Introduction

In recent years, online communities are emerging for exchanging services such as knowledge, content, expertise or physical (computing and communication) resources, as evidenced by the popularity of peer-to-peer (P2P) systems [1], [2], collaborative signal processing systems [3], online social networking systems [4], and crowdsourcing applications [7]. While such online communities differ in many ways, they will all need to deal with similar challenges if the agents are self-interested and strategic, i.e. they aim to maximize their own individual utilities. Since contributing high-quality services does not generate an immediate and direct benefit for themselves, self-interested agents tend to choose not to spend efforts or resources to provide high-quality services. In contract theory which studies the incentive of agents to exert high effort to provide services, this phenomenon is commonly known as moral hazard [27]. Given the self-interested nature of agents, it is thus crucial to design appropriate incentives that encourage agents to contribute high-quality services in order to sustain high-performance online communities.

A variety of incentive mechanisms have been explored to induce cooperation in service communities [8]–[25], among which the most popular forms are *pricing* and *reciprocity*.

Incentive mechanisms based on pricing incentivize agents to provide high-quality services by rewarding them with either money or fiat money (e.g. tokens), which are paid by the agents requesting services. However, most pricing schemes rely on the assumption that the individual behavior of each agent can be accurately identified and measured. This assumes perfect monitoring and assessment of agents' transactions, which, unlike in most service communities, is difficult to achieve in an online environment [12], [17]. For example on a crowdsourcing platform, perfect monitoring of an expert's exerted effort level and the quality of his/her solutions is impossible to implement [4]. Instead, the quality of the solutions is often determined by the reports/feedback submitted by the requesters, which are often imperfect since the requesters may not be able to accurately assess the received services or their reports may be missing or lost. Moreover, individualized reports/feedback for specific experts may not be available since experts may be anonymous or more than one expert may be involved in solving a task, and a requester may only assess and provide reports about the aggregated service he experiences. Given the infeasibility of perfect monitoring, traditional pricing schemes alone cannot be credibly implemented in online communities due to the following problem:

- If the payment is performed ex-ante (i.e. agents receive payments before they provide the services), an agent always has the incentive to take the payment in a transaction without providing the promised services;
- Whereas if the payment is ex-post/ex-interim, which means that agents cannot receive payments before providing services, a requesting agent always has the incentive to refuse payment by claiming that he received low quality services.

Also, a pricing schemes often requires a complex account infrastructure, which introduces substantial communication and computation overheads [10]. Hence, it is impractical for traditional pricing schemes and the corresponding infrastructure to be implemented in large-scale online communities.

Manuscript received July 02, 2013; revised November 22, 2013; accepted February 08, 2014. Date of publication April 24, 2014; date of current version May 16, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Stefano Marano.

Y. Zhang is with Microsoft Corporation, Sunnyvale, CA 94086 USA (e-mail: yuzhan@microsoft.com).

M. van der Schaar is with the Electrical Engineering Department, University of California, Los Angeles, CA 90024 USA (e-mail: mihaela@ee.ucla.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2014.2320457

In online communities where traditional pricing schemes are infeasible, it becomes important to exploit the fact that agents are part of a community in which they repeatedly interact, and use this knowledge to design more credible incentive mechanisms which reward or punish agents based on their past contributions.

Incentive mechanisms based on reciprocity treat the interactions among agents as repeated exchanges. They reward and punish individual agents by providing differential services, depending on their past service provision [2], [14]–[16]. Such incentive mechanisms can be classified into direct reciprocity and indirect reciprocity [9]. In direct reciprocity mechanisms, agents can identify each other, and they reward/punish each other based on their past personal experiences with each other [2]. Direct reciprocity is effective in sustaining cooperation if agents can identify each other and are interested in interacting with the same partner over time. However, in most online communities such as the emerging crowdsourcing platforms [7], the agents have asymmetric service needs and they are not interacting with the same partners over time and hence, a personal history of past reciprocation with the same partner cannot be established.

To encourage the self-interested agents to provide high-quality services in such communities, indirect reciprocity (also called "social" reciprocity) solutions have been proposed [14]–[16], which do not require the personal history of past reciprocation to enforce rewards and punishment. One of the most successful indirect reciprocity solutions is based on rating schemes [7], [17]–[24], in which agents are rated based on the quality of services that they provided in the past. The agents can then be rewarded or punished by the agents in the online community whom they will meet in the future, even though they have never interacted with each other in the past: a high/low rated agent will receive high/low payments from future requesting agents.

A key limitation of such existing rating schemes for online communities is that the rating of each individual agent is computed based on the assumption that his personal behavior, i.e. the quality of services he provides, can be accurately identified and measured [21]–[24]. However, such personal rating schemes are not appropriate for online communities which exhibit the following characteristics:

1) **Agent Anonymity**. In online communities, it is difficult, if not infeasible, to identify individual agents and keep track of their behavior. Moreover, agents often collectively provide services to a requesting agent and they may not want to be identified due to privacy reasons.
2) **High Costs of Maintaining Ratings**. The existing personal rating schemes require constant monitoring of all individual agents' service quality, in order to periodically update their personal ratings. The cost of maintaining, updating and communicating personal ratings grows rapidly as the agent population increases, thereby becoming prohibitively expensive in online communities with large populations.
3) **Whitewashing**. Since the online identity of each individual agent is not directly associated with his real-world identity, an agent can create multiple online (personal) identities by repeatedly entering and registering in the community. As a result, an agent who has been detected to be non-cooperative (e.g. avoiding the provision of high-quality services after receiving the payment in a transaction) may attempt to switch his online identity, by leaving and rejoining the community as a new member to avoid the punishments imposed by the personal rating scheme upon his old identity, a behavior commonly known as "whitewashing". Such whitewashing activities also complicate the deployment of personal rating schemes [25].
4) **Monitoring Errors**. In online communities, the monitoring of the agents' behavior is never perfect in practice (e.g. the reports provided by requesting agents about the service of providing agents is often inaccurate due to perception errors, incompetence etc.) and hence, the rating update is always not accurate. This makes the rating design problem more complicated: (1) If monitoring was perfect, the community designer could employ the strongest punishment (e.g. a trigger strategy [26]): expelling an agent from the community if he does not provide the requested service, thereby enforcing the strongest incentive on agents to provide the highest quality of services. However, when monitoring is imperfect, the monitoring results are not guaranteed to be accurate and thus relying on a strong punishment would lead to very low social welfare in the community if trigger strategies were employed. Instead, as we will show, choosing milder punishments when a deviation is observed is often more efficient. However, such punishments cannot be too mild either since sufficient incentives need to be provided to agents to compel them to deliver high quality services. Therefore, an optimally designed rating scheme needs to be designed as a tradeoff between efficiency and sustainability (incentive-compatibility).

In this paper, we address the abovementioned challenges by designing a new incentive mechanism that integrates differential pricing with a novel class of rating schemes, which we refer to as collective ratings. The agents in the community are classified as *requesters*, who request services from other agents, and *workers* who possess valuable resources (e.g. data, knowledge, information, etc.) and provide services to requesters. It is important to note that an agent can be both a worker and a requester. The workers are exogenously divided into multiple groups determined by the community designer, based on the expertise/resources possessed by the workers. Examples of such worker groups are: experts mastering the same skill on a crowdsourcing platform, sellers belonging to the same organization on an online trading platform, geographically-close peers residing in the same sub-network of a P2P network etc. Since the workers are divided exogenously into groups based on their expertise/resources, they cannot choose their group memberships (i.e. which groups to join) or switch the groups with which they are affiliated. Hence, the whitewashing problem that occurs in personal rating schemes is prevented. Each group is operated by one or multiple group operators (e.g. the trackers in P2P systems, the specialized operators in a crowdsourcing platform etc.), who forward service requests, match requesters and workers, implement and maintain the payment infrastructure,

TABLE I
COMPARISON OF OUR WORK AND THE EXISTING LITERATURE

	[8][9]	[2]	[14]-[16]	[17]-[18]	[21]-[23]	[20][24][25]	Our work
Incentive device	Pricing scheme	Direct reciprocity	Indirect reciprocity	Indirect reciprocity	Indirect reciprocity	Indirect reciprocity	Indirect reciprocity
Agent's interaction	One-time	Repeated	Repeated	Repeated	Repeated	Repeated	Repeated
Agent's identity	Known	Known	Known	Known	Known	Known	Anonymous
Monitoring error	No	No	No	No	No	Yes	Yes
Rating label	N/A	N/A	N/A	Personal	Personal	Personal	Collective
Agent's benefit and cost	Heterogeneous	Homogeneous	Homogeneous	Heterogeneous	Homogeneous	Homogeneous	Heterogeneous
Equilibrium analysis	Yes	No	Yes	No	Yes	Yes	Yes
Incentive mechanism optimization	No	No	Yes	No	No	Yes	Yes

etc. We assume that the payment in the online community is performed ex-ante and focus mainly on the incentives of workers on the service provision. However, our proposed methodology can be extended easily to the scenarios where the payments are performed ex-post, i.e. upon the completion of services.

Next, we briefly describe the operation of the proposed incentive mechanism. The community designer (e.g. the platform owner etc.) is in charge of keeping record on the collective rating of each group. A group earns its rating based on the history of service provision delivered collectively by its workers: a higher rating is assigned to groups whose workers have provided high-quality services in the past. In this way, the services of workers in a group are rated collectively rather than individually. The community designer uses the collective rating to implement (differential) rating-based pricing for the services of each group: it assigns higher payments to workers appertaining to groups having higher ratings.[1] Since the groups which possess high ratings are rewarded by higher payments, incentives are (indirectly) enforced to individual workers to provide high-quality services in order to help their groups to maintain high ratings.[2]

In this paper, we first model the repeated interactions among agents as a repeated games with anonymous matching. By designing an appropriate collective rating protocol to reward/punish workers' service provision, we rigorously analyze how the workers' service provision is influenced by the designed payments and incurred costs as well as the workers' valuation of their individual long-term utilities. Subsequently, we quantify the sufficient and necessary conditions for sustainable rating protocols (i.e. rating protocols which are equilibrium strategies), under which workers find in their own self-interest to play cooperatively and provide high-quality services. We then define and solve the community designer's problem to jointly optimize the rating scheme and the pricing scheme, maximizing the revenue that it obtains from the workers' service provision. Since workers in the same group share the same rating, the reward/punishment imposed on each of them and thus their incentives to provide high-quality services are not solely influenced by their own service quality, but also by the service quality of all the other group members. We quantify such externalities by explicitly analyzing how the size of the group impacts our design and the community performance.

In Table I, we provide a comprehensive comparison between our proposed collective rating protocol and various existing incentive mechanisms.

The remainder of this paper is organized as follows. In Section II, an analytical framework is proposed to analyze the interactions emerging among individual agents. Section III formulates the utility functions for individual workers and their decision problems, and then defines the design problem of the optimal collective rating protocol. Section IV determines how to optimally design collective rating protocols. Section V analyzes the impact of the worker group size on the efficacy of the collective rating protocol. Section VI presents illustrative results to highlight the features of the proposed collective rating protocol. We conclude in Section VII.

II. SYSTEM MODEL

A. System Setting

We consider an online community where workers provide different types of services for monetary rewards. Without loss of generality, we consider the community consisting of N exogenously-formed worker groups. Let $\mathcal{N} = \{1, \ldots, N\}$ denote the set of worker groups with each group indexed by $i \in \mathcal{N}$. Each group i consists of M_i workers, represented by $\mathcal{M}_i = \{1, \ldots, M_i\}$, and is managed by a group operator, denoted as o_i. The worker j managed by the operator o_i is denoted as $h_{i,j}$. Each group operator maintains a service directory/map recording the services that can be provided by the workers it manages and announces this directory to all agents in the community.

The community designer (owner) administers the entire community with the self-interested goal of maximizing its revenue from the transactions among the agents. It designs and implements the pricing plan to reward workers for their services, implement the associated payment infrastructure as well as the rating scheme. The group operators are obedient entities deployed by the community designer to ensure that the community operates effectively. Hence, they are controlled by the community designer and thus are not strategic in their operations (i.e. they simply implement the strategies provided by the community designer). However, our proposed rating protocols can be easily extended to the scenario where the group operators are

[1]It should also be noted that the community designer does not directly pay the workers but only determines the pricing scheme, i.e. the amount of payments transferred among workers and requesters for service provision. The requesters themselves make direct payments for the services they request.

[2]Note that such collective rating protocols only require the groups to be identified, while allowing individual workers to remain anonymous.

TABLE II
SUMMARY OF VARIABLES

Variable	*Name*
M_i	Group i
$\mathcal{N}$	The set of groups
$\mathcal{M}_i$	The set of workers in group i
o_i	The operator of group i
$h_{i,j}$	Worker j in group i
$a_{i,j}$	The service rate of worker $h_{i,j}$
$\lambda_{i,j}$	The requested service rate of worker $h_{i,j}$
π_i	The aggregated service rate of group i
μ_i	The aggregated requested service rate from group i
c	Unit service cost
$Q_{i,j}$	One-period payment received by $h_{i,j}$
κ	The rating protocol
Θ	The set of rating labels
ω	The monitoring method
τ	The rating scheme
ψ	The pricing scheme
θ	The rating label
$\mathcal{S}$	The set of monitoring signals
s	The monitoring signal
p_{MIN}	The probability that a worker receives a negative signal upon compliance
p_{MAX}	The maximum probability that a worker receives a negative signal upon deviation
α	The punishment probability upon receiving a negative signal
ρ	The ratio between the transaction fee and the service fee
$v_{i,j}$	The expected one-period utility
$V_{i,j}$	The expected long-term utility
$\overline{q}$	The maximum benefit a requester receives from one unit service
U	The revenue of the community designer
J	The social welfare of the community
J^*	The social optimum

strategic entities who aim to optimize their own utilities. Moreover, in this paper, the community designer focuses on designing and deploying an incentive mechanism which is applied to all the groups. By employing the same design principle, the proposed rating protocol design can be easily extended to the scenario where differential rating schemes and pricing schemes are designed and optimized for each group.

Note also that in the proposed implementation, the groups are not necessarily mutually exclusive, i.e. a worker can appertain to multiple groups. For example, an expert on a crowdsourcing platform can possess multiple skills and be a member of multiple expertise groups. Also, he can join new groups when he acquires new expertise and exit affiliated groups when no longer possesses an expertise. Because the price the expert charges for one particular skill is determined by the rating of the corresponding expertise group and is independent of the prices which he charges for his other skills, such multi-group membership will not affect the design and efficacy of the incentive mechanism. Moreover, it is important to note that the set of workers, i.e. $\{\mathcal{M}_i\}_{i\in\mathcal{N}}$, does not necessarily include all agents in the community. It is also possible that an agent does not belong to any worker group, i.e. he is always a requester and does not provide any services to other agents. Our system and design can efficiently operate in such scenarios.

In the rest of this section, we formalize the interactions among the workers and requesters as a non-cooperative repeated game, where the time is divided into periods of the same length. For each period, we use $a_{i,j} \geq 0$ to denote the *service rate* of worker $h_{i,j}$, which represents the amount of resources/efforts he provides in one period to serve the requesters who request services from him. Also, we use $\lambda_{i,j} \geq 0$ to denote the **requested service rate** of $h_{i,j}$, which represents the amount of resources $h_{i,j}$ needs to provide in one period, in order to fulfill the services requested from him. It is important to note that $a_{i,j}$ and $\lambda_{i,j}$ are not necessarily the same since each worker can proactively determine the amount of resources he would like to contribute in each period. For each group i, we use $\pi_i \triangleq \sum_{j\in\mathcal{M}_i} a_{i,j}$ to denote its *aggregated service rate* per period and $\mu_i \triangleq \sum_{j\in\mathcal{M}_i} \lambda_{i,j}$ to denote its *aggregated requested service rate* per period.

At the beginning of each period, each worker $h_{i,j}$ strategically determines his service rate in this period, i.e. the value $a_{i,j} \in \mathbb{R}^+$. The one-period cost incurred by worker $h_{i,j}$ when providing services is $ca_{i,j}$, where c is the unit cost. Let $Q_{i,j}$ denote the payment received by $h_{i,j}$ in a period, his one-period utility can be written as $Q_{i,j} - ca_{i,j}$. Since $h_{i,j}$ receives his payment before providing services, choosing an action $a_{i,j} = 0$, i.e. to provide no service, myopically maximizes his utility. As a result, $h_{i,j}$ always has the incentive to not provide high-quality services, regardless of the amount of payment he receives (i.e. the value of $Q_{i,j}$). For this reason, incentive mechanisms are required in order to encourage individual workers to contribute their services.

B. Rating Protocol

The reason why workers find it optimal to not provide high-quality services is the absence of punishments for such behavior. Note that, since agents are interacting repeatedly in the community, incentive mechanisms can be designed by the community designer through the enforcement of punishments on the workers who do not provide high-quality services. The incentive mechanism proposed in this paper is based on rating protocols. A rating protocol is designed and implemented by the community designer, and is represented as a tuple $\kappa = (\Theta, \omega, \tau, \psi)$. It consists of four components: a set of rating labels Θ, a monitoring method ω, a rating scheme τ, and a pricing scheme ψ.

(1) **Rating label**: The proposed collective rating protocols incentivize the service provision of individual workers by adapting the rating of their group, which is denoted as $\theta \in \Theta = \{0, 1\}$, and characterizes the “quality” of the services

provided by the individual workers in a group. While in our paper we focus solely on binary rating sets—where $\theta = 1$ is the good rating and $\theta = 0$ is the bad rating—our analysis can be extended to a rating set with multiple rating levels. We focus on binary ratings in this paper to highlight what can be achieved even with *simple* collective rating protocols.

The ratings of groups are stored, updated, and published/broadcasted by the community designer. Hence, they represent the public information which can be accessed by all the agents (i.e. both workers and requesters) and group operators.

(2) **Monitoring method**: In order to update the ratings of groups, the community designer monitors the quality of service provided by each group to determine whether the workers in that group contribute sufficiently for high-quality service provision. The details of the monitoring method depend on the considered type of community. For example, in a P2P network where the services of data upload are performed, the quality of service can be determined by sampling the aggregated upload traffic rate from each group. In a crowdsourcing platform where the services of solving various types of tasks are performed, the quality of service can be determined by collecting the feedbacks from the requesters reviewing how well their tasks are solved.

For illustration, we adopt a simple monitoring method. Let $\mathcal{S} = \{0, 1\}$ denote the set of monitoring signals that the community designer gathers about a group i, and a signal $s = 0$ indicates that the services provided by this group are not adequate (e.g. the services provided by the group are, on average, of low quality) and thus, this group needs to be punished.[3] A monitoring method ω is defined as a measure-valued mapping $\omega : \bigcup_{j\in\mathcal{M}_i} \mathcal{A} \to \Delta(\mathcal{S})$ where $\mathcal{A} \in \mathbb{R}^+$ is the action space of each individual worker. In particular, $\omega(s \mid a_{i,1}, \ldots, a_{i,M_i})$ represents the conditional probability that the monitoring signal s is observed given that the actions $(a_{i,1}, \ldots, a_{i,M_i})$ are played by the workers $(h_{i,1}, \ldots, h_{i,M_i})$. The detailed formulation of $\{\omega(s \mid a_{i,1}, \ldots, a_{i,M_i})\}$ depends on the choice of the monitoring method ω. In this paper, we adopt the following formulation as an illustrative example:

$$\omega(s = 1 \mid a_{i,1}, \ldots, a_{i,M_i}) = p_{MAX} - (p_{MAX} - p_{MIN}) \min \left\{ \frac{\sum_{j\in\mathcal{M}_i} a_{i,j}}{\mu_i}, 1 \right\}. \quad (1)$$

In this example, the probability for a group i to receive a positive signal is minimized at the value p_{MIN} when all its workers serve at their requested service rates, i.e. $a_{i,j} \geq \lambda_{i,j}, \forall j \in \mathcal{M}_i$. Here it should be noted that group i still has a positive probability $p_{MIN} > 0$ to receive a negative signal even if its aggregated service rate is sufficiently high to support the requests towards it, i.e. $\pi_i = \sum_{j\in\mathcal{M}_i} a_{i,j} \geq \mu_i$. This is due to various monitoring errors during the implementation, and we use p_{MIN} to denote the probability that monitoring errors occur in the community. On the other hand, when the service rate of some worker is smaller than his requested service rate, the probability for his group to receive a negative signal linearly increases against the value $\mu_i - \pi_i$, with its maximum value achieved at $p_{MAX} > p_{MIN}$.

(3) **Rating scheme**: Given the monitoring signals, the community designer updates the ratings of all groups at the end of each period (or at the beginning of the next period). The rating scheme specifies the rule that the community designer adopts in the rating update and is represented by a mapping $\tau : \Theta \times \mathcal{S} \to \Delta(\Theta)$. Given a group of rating $\theta \in \Theta$ and the monitoring signal $s \in \mathcal{S}$ observed upon it, the rating scheme updates the rating to $\theta' \in \Theta$ with a probability $\tau(\theta'|\theta, s)$. We consider the following rating scheme:

$$\tau(\theta'|\theta, s) = \begin{cases} 1 & if\ \theta' = 1\ and\ s = 1 \\ \alpha & if\ \theta' = 0\ and\ s = 0 \\ 1 - \alpha & if\ \theta' = \theta\ and\ s = 0 \\ 0 & otherwise \end{cases}. \quad (2)$$

Briefly explained, the rating scheme (2) updates the rating of a group to be $\theta = 1$ if the monitoring signal $s = 1$, i.e. the group is observed to fulfill the aggregated requested service rate in this period. Once a negative signal is observed, the rating of this group is decreased to $\theta = 0$ with a probability α and remains unchanged with a probability $1 - \alpha$. Hence, α can be referred to as the strength of punishment imposed to workers when they do not provide high-quality services.

(4) **Pricing scheme**: The pricing scheme defines the rules that the community designer uses to reward/punish workers in order to incentivize their service provision. The reward is realized by implementing differential prices for the services according to the ratings of groups. Formally, a rating-based pricing scheme is expressed as a mapping $\psi : \Theta \to \mathbb{R}^+$. For a worker $h_{i,j}$ from a group of rating θ, the total service fee he receives for service provision is $\psi(\theta)\lambda_{i,j}$, where $\psi(\theta)$ is the unit service fee. In order to enforce incentives on service provision, the community designer assigns workers from high-rating groups higher service fees as reward and hence, workers are encouraged to provide services to increase the ratings of their groups and thus receive higher service fees in return. Besides the service fee, a requester also pays the community designer certain transaction fee to cover the operation cost incurred by the community designer and the group operators. We assume that the transaction fee associated with a transaction is proportional to the corresponding service fee, at a predetermined ratio ρ [28], [29].

To determine the range of feasible prices, we assume that the maximum benefit which a requester obtains from receiving one unit service is $\bar{q}$. Therefore, the unit service fee should be restricted in the region $\left[c, \frac{\bar{q}}{1+\rho}\right]$, i.e. $\psi(\theta) \in \left[c, \frac{\bar{q}}{1+\rho}\right], \forall\theta$. If $\psi(\theta) > \frac{\bar{q}}{1+\rho}$, no requester will choose to request services since his received benefit is less than the total fee (the service fee and the transaction fee) that he pays which amounts to $(1 + \rho)\psi(\theta)$; whereas if $\psi(\theta) < c$, no worker will choose to provide services to the requesters since the payment he receives from the service provision is less than the total incurred cost.

It is important to note that under a collective rating protocol, the rating update and the monitoring are performed per group and hence, the total numbers of rating updates and monitoring during each period are both of linear order in the total number of groups in the community, i.e. $O(N)$. This is in sharp contrast to the personal rating protocols, e.g. the one proposed in [25],

[3]It should be noted that the group operators monitor the service collaboratively and hence, they observe a common signal for each individual group in each period.

under which the total numbers of rating updates and monitoring performed per period are both of linear order in the total number of individual workers, i.e. $O(\sum_{i=1}^{N} M_i)$. Therefore, the collective rating protocol significantly reduces the incurred computation cost as well as the communication overhead, and is easier to be implemented in online communities with large populations, where the values of $\{M_i\}$ are large.

With the rating protocol, a typical transaction in the online community takes the following steps:

1) A requester requests services and his request is forwarded to the group operators.
2) If a group operator, e.g. o_i, finds the requested service in its directory, it notifies the requester about its ability to perform this service and collects the payment from him according to the pricing scheme ψ.
3) o_i then selects a set of workers from the group i who can provide the requested service and distributes the received service fee to these workers.
4) The selected workers provide services collectively to fulfill the requested service.
5) After the selected workers provide services, the requester submits his feedback to the community designer, which reports the quality of service he received from group i.
6) The community designer aggregates the reports from all requesters who requested services from group i in this period and determines the monitoring signal s_i according to ω.
7) Given the monitoring signal s_i, the community designer updates the rating of group i according to τ.

To summarize, the design variables available to the community designer include the punishment probability α and the pricing scheme ψ. In the rest of this paper, we use the tuple of the available design parameters to denote the rating protocol, i.e. $\kappa \triangleq (\alpha, \psi)$.

C. Design Problem

We assume that the community designer is profit-seeking and designs the rating protocol κ to maximize its own revenue per period, which is the sum transaction fee it collects and is proportional to the sum service fee paid by the requesters in each period.[4] Let $\theta^{(t)}$ denote the rating of group i in period t. The sum service fee collected by group i in period t is $\mu_i \psi(\theta^{(t)})$. Given this, the rating protocol design problem can be formalized as follows:[5]

$$\max_{\alpha,\psi} U \triangleq \lim_{L\to\infty} \frac{1}{L} \sum_{t=0}^{L} \sum_{i\in\mathcal{N}} \mu_i \psi\left(\theta^{(t)}\right). \quad (3)$$

III. Sustainable Rating Protocol

This section first determines the long-term utility functions of individual workers and then analyzes their optimal strategies on service provision. Based on this analysis, we reformulate the optimal rating protocol design problem (3) for the community designer. This design problem will be solved in Section IV and V.

A. Utility Functions

This section discusses the long-term utility functions and service provision strategies of the individual workers. A worker $h_{i,j}$'s service provision strategy can be formalized as a mapping $\sigma : \Theta \to \mathbb{R}^+$. That is, the worker determines his service rate based on the rating of his group as $a = \sigma(\theta)$. Given this, the expected one-period utility of a worker $h_{i,j}$ can be expressed as:

$$v_{i,j}(\theta \mid \kappa, \sigma) = \psi(\theta)\lambda_{i,j} - c\sigma(\theta). \quad (4)$$

The expected long-term utility of a worker is the infinite-horizon discounted sum of his expected one-period utility with a discount factor $\delta > 0$, which can be expressed as:

$$V_{i,j}\left(\theta^{(t_0)} \mid \kappa, \sigma\right) = v_{i,j}\left(\theta^{(t_0)} | \kappa, \sigma\right) + \delta \sum_{\theta' \in \Theta} p_{i,j}(\theta' \mid \theta^{(t_0)}, \kappa, a) V_{i,j}(\theta' \mid \kappa, \sigma). \quad (5)$$

The first part of this utility $v_{i,j}(\theta^{(t_0)}|\kappa,\sigma)$ is the expected utility which $h_{i,j}$ receives in the current period t_0 when group i 's rating is $\theta^{(t_0)}$ and the second part $\delta \sum_{\theta}^{\prime} \in \Theta p_{i,j}(\theta'|\theta^{(t_0)},\kappa,\sigma) V_{i,j}(\theta'|\kappa,a)$ is $h_{i,j}$'s discounted utility in the future periods. $p_{i,j}(\theta'|\theta,\kappa,a)$ denotes the probability that the rating of group i changes from θ to θ' when the rating protocol is κ and $h_{i,j}$ chooses the action a. The value of $\{p_{i,j}(\theta'|\theta,\kappa,a)\}$ depends on the service provision strategies adopted by other workers from the same group. Since $h_{i,j}$ cannot observe the strategies of others, we assume that he maintains a simple belief that all the workers other than himself always choose to provide services at their requested service rates, i.e. $a_{i,j'} = \lambda_{i,j'}, \forall j' \in \mathcal{M}_i$ *and* $j' \neq j$. Given this assumption, $p_{i,j}(\theta'|\theta,\kappa,a)$ can be determined as follows:

$$p_{i,j}\theta'|\theta,\kappa,a) = \begin{cases} 1-\varepsilon(\lambda_{i,j},\mu_i,a)\alpha, & \theta=\theta'=1 \\ \varepsilon(\lambda_{i,j},\mu_i,a)\alpha, & \theta=1,\theta'=0 \\ 1-\varepsilon(\lambda_{i,j},\mu_i,a), & \theta=0,\theta'=1 \\ \varepsilon(\lambda_{i,j},\mu_i,a), & \theta=0,\theta'=0 \end{cases}. \quad (6)$$

Here $\varepsilon(\lambda_{i,j},\mu_i,a)$ denote the probability that the monitoring signal of group i is $s =$ when (1) its aggregated requested service rate is μ_i, (2) the requested service rate and the service rate of worker $h_{i,j}$ are $\lambda_{i,j}$ and a, respectively, and (3) all workers in group i other than $h_{i,j}$ fulfill their requested service rates. Similar to (1), we have

$$\varepsilon(\lambda_{i,j},\mu_i,a) = \max\left\{\frac{(\lambda_{i,j}-a)p_{MAX} + (\mu_i - \lambda_{i,j} + a)p_{MIN}}{\mu_i, p_{MIN}}\right\}. \quad (7)$$

The optimal service provision strategy maximizes a worker's expected long-term utility at any period, which can be formally defined as follows.

[4]It should be noted that the design methodology proposed in this paper also applies to the scenarios where the community designer are not profit-seeking and design the rating protocol to optimize some alternative objective functions other than its revenue.

[5]Since the revenue of the community designer is proportional to the sum service fee, it can be maximized by solving Eq. (3).

Definition 1 (Optimal Worker Strategy): The optimal strategy $\sigma_{i,j}^{\kappa}$ for worker $h_{i,j}$ is the strategy that satisfies the following inequalities:

$$V_{i,j}\left(\theta \mid \kappa, \sigma_{i,j}^{\kappa}\right) \geq V_{i,j}(\theta \mid \kappa, \sigma),\ \forall \theta \in \Theta. \tag{8}$$

■

It should be noted that since different workers receive different requested service rates, their optimal strategies are not necessarily the same even if they are from the same group (and thus share the same rating). By analyzing (4) and (5), it is easy to determine that given a rating protocol κ, a worker's decision problem can be formulated as a Markov decision process where the state is the rating θ of his group and the action is his service rate a. Hence, the optimal strategy complies with the one-shot deviation principle [26]. That is, if a strategy $\sigma_{i,j}^{\kappa}$ is optimal, a worker cannot benefit by deviating from $\sigma_{i,j}^{\kappa}$ in the current period (i.e. the one-shot deviation) and complying with $\sigma_{i,j}^{\kappa}$ in all subsequent periods. Using this idea, the optimal strategy $\sigma_{i,j}^{\kappa}$ can be computed using dynamic programming and we have the following proposition that characterizes its structure.

Proposition 1: Given the rating protocol κ, the optimal strategy $\sigma_{i,j}^{\kappa}$ is unique and preserves the following properties:

(i) $\sigma_{i,j}^{\kappa}(\theta) \in \{0, \lambda_{i,j}\}, \forall \theta$;
(ii) $\sigma_{i,j}^{\kappa}(0) \geq \sigma_{i,j}^{\kappa}(1), \forall i, j$;
(iii) for $h_{i,j_1}, h_{i,j_2} \in \mathcal{M}_i$ and if $\lambda_{i,j_1} \leq \lambda_{i,j_2}$, then $\sigma_{i,j_1}^{\kappa}(\theta) \leq \sigma_{i,j_2}^{\kappa}(\theta), \forall \theta$;
(iv) for any $h_{i,j} \in \mathcal{M}_i$ and any $h_{i',j'} \in \mathcal{M}_{i'}$ and if $\lambda_{i,j} = \lambda_{i',j'}$, then $\sigma_{i,j}^{\kappa}(\theta) \geq \sigma_{i',j'}^{\kappa}(\theta), \forall \theta$ if and only if $\mu_i \leq \mu_{i'}$.

Proof: Since the state space of the MDP for an individual worker is communicating and irreducible, the uniqueness of the optimal strategy directly follows.

(i) By analyzing $V_{i,j}(\theta^{(t_0)}|\kappa, \sigma)$, it is easy to note that the first term $v_{i,j}(\theta^{(t_0)}|\kappa, \sigma)$ linearly decreases with the action a, while the second term $\delta \sum_{\theta' \in \Theta} p_{i,j}(\theta' \mid \theta^{(t_0)}, \kappa, a) V_{i,j}(\theta' \mid \kappa, \sigma)$ linearly increases with a. Hence, the optimal action of a worker $h_{i,j}$ in each period is $a = 0$ if

$$\frac{\partial v_{i,j}(\theta^{(t_0)}|\kappa, \sigma)}{\partial a} < -\frac{\partial \sum_{\theta' \in \Theta} p_{i,j}(\theta' \mid \theta^{(t_0)}, \kappa, a)}{\partial a} \delta V_{i,j}(\theta' \mid \kappa, \sigma)$$

and is $a = \lambda_{i,j}$ if

$$\frac{\partial v_{i,j}(\theta^{(t_0)}|\kappa, \sigma)}{\partial a} > -\frac{\partial \sum_{\theta' \in \Theta} p_{i,j}(\theta' \mid \theta^{(t_0)}, \kappa, a)}{\partial a} \delta V_{i,j}(\theta' \mid \kappa, \sigma).$$

(ii) To prove this statement, we only have to show that a strategy σ with $\sigma(0) = 0$ and $\sigma(1) = \lambda_{i,j}$ cannot be the optimal strategy. We prove this by contradiction. Suppose that σ is the optimal strategy for some worker $h_{i,j}$. Then according to the rating transition probability (6), the expected long-term utility of this worker can be constructed as follows:

$$\begin{aligned} V_{i,j}(1 \mid \kappa, \sigma) = {} & \psi(1)\lambda_{i,j} - c\lambda_{i,j} \\ & + \delta\left((1 - p_{MIN}\alpha) V_{i,j}(1 \mid \kappa, \sigma)\right. \\ & \left. + p_{MIN}\alpha V_{i,j}(0 \mid \kappa, \sigma)\right) \\ V_{i,j}(0 \mid \kappa, \sigma) = {} & \psi(0)\lambda_{i,j} \\ & + \delta\left((1 - \varepsilon(\lambda_{i,j}, \mu_i, 0)) V_{i,j}(1 \mid \kappa, \sigma)\right. \\ & \left. + \varepsilon(\lambda_{i,j}, \mu_i, 0) V_{i,j}(0 \mid \kappa, \sigma)\right). \end{aligned} \tag{9}$$

Since σ is optimal, $h_{i,j}$ cannot benefit by unilaterally deviate from σ when $\theta = 1$. We have the following inequality:

$$\begin{aligned} & V_{i,j}(1 \mid \kappa, \sigma) \\ & \quad \geq \psi(1)\lambda_{i,j} + \delta\left((1 - \varepsilon(\lambda_{i,j}, \mu_i, 0)\alpha) V_{i,j}(1 \mid \kappa, \sigma)\right. \\ & \qquad \left. + \varepsilon(\lambda_{i,j}, \mu_i, 0)\alpha V_{i,j}(0 \mid \kappa, \sigma)\right). \end{aligned} \tag{10}$$

Similarly, since $h_{i,j}$ cannot benefit by unilaterally deviate from σ when $\theta = 0$, the following inequality also holds:

$$\begin{aligned} & V_{i,j}(0 \mid \kappa, \sigma) \\ & \quad > \psi(0)\lambda_{i,j} - c\lambda_{i,j} + \delta\left((1 - p_{MIN}) V_{i,j}(1 \mid \kappa, \sigma)\right. \\ & \qquad \left. + p_{MIN} V_{i,j}(0 \mid \kappa, \sigma)\right). \end{aligned} \tag{11}$$

Substituting (9) into the inequalities (10), (11), we have that

$$\begin{aligned} & \delta\left(\varepsilon(\lambda_{i,j}, \mu_i, 0) - p_{MIN}\right)\left(V_{i,j}(1|\kappa, \sigma) - V_{i,j}(0|\kappa, \sigma)\right) \geq \frac{c\lambda_{i,j}}{\alpha} \\ & \delta\left(\varepsilon(\lambda_{i,j}, \mu_i, 0) - p_{MIN}\right)\left(V_{i,j}(1|\kappa, \sigma) - V_{i,j}(0|\kappa, \sigma)\right) < c\lambda_{i,j}. \end{aligned} \tag{12}$$

It is obvious that the two inequalities in (12) cannot hold simultaneously when $\alpha > 0$. Hence, this statement follows.

(iii) Consider two workers $h_{i,j_1}, h_{i,j_2} \in \mathcal{M}_i$ and $\lambda_{i,j_1} \geq \lambda_{i,j_2}$. Suppose $\sigma_{i,j_1}^{\kappa}(\theta) = \lambda_{i,j_1}$ and $\sigma_{i,j_2}^{\kappa}(\theta) = 0$. In the rest of this proof, we consider the scenario with $\theta = 1$. The analysis for the scenario with $\theta = 0$ follows in a similar manner. The expected long-term utilities of h_{i,j_1} and h_{i,j_2} can be expressed as follows:

$$\begin{aligned} & V_{i,j_1}\left(1 \mid \kappa, \sigma_{i,j_1}^{\kappa}\right) \\ & \quad = \psi(1)\lambda_{i,j_1} c\lambda_{i,j_1} + \delta\left((1 - p_{MIN}\alpha) V_{i,j_1}\left(1 \mid \kappa, \sigma_{i,j_1}^{\kappa}\right)\right. \\ & \qquad \left. + p_{MIN}\alpha V_{i,j_1}(0 \mid \kappa, \sigma_{i,j_1}^{\kappa})\right) \\ & V_{i,j_2}\left(1 \mid \kappa, \sigma_{i,j_2}^{\kappa}\right) \\ & \quad = \psi(1)\lambda_{i,j_2} + \delta\left((1 - \varepsilon(\lambda_{i,j_2}, \mu_i, 0)\alpha) V_{i,j_2}(1 \mid \kappa, \sigma_{i,j_2}^{\kappa})\right. \\ & \qquad \left. + \varepsilon\left(\lambda_{i,j_2}, \mu_i, 0\right)\alpha V_{i,j_2}\left(0 \mid \kappa, \sigma_{i,j_2}^{\kappa}\right)\right). \end{aligned} \tag{13}$$

From statement (ii), it is obvious that $\sigma_{i,j_1}^{\kappa}(0) = \lambda_{i,j_1}$. Given this, we have

$$\begin{aligned} & V_{i,j_1}\left(0 \mid \kappa, \sigma_{i,j_1}^{\kappa}\right) \\ & \quad = \psi(0)\lambda_{i,j_1} - c^d\lambda_{i,j_1} + \delta\left((1 - p_{MIN}) V_{i,j_1}(1 \mid \kappa, \sigma_{i,j_1}^{\kappa})\right. \\ & \qquad \left. + p_{MIN} V_{i,j_1}\left(0 \mid \kappa, \sigma_{i,j_1}^{\kappa}\right)\right). \end{aligned} \tag{14}$$

Since h_{i,j_1} has no incentive to unilaterally deviate either at rating $\theta = 0$ or at rating $\theta = 1$, then according to the one-shot deviation principle, the following inequality should hold:

$$\delta(\varepsilon(\lambda_{i,j_1}, \mu_i, 0) - p_{MIN}) \frac{\psi(1) - \psi(0)}{1 - \delta(1 - \alpha)p_{MIN}} \geq \frac{c}{\alpha}. \tag{15}$$

The LHS of (15) increases with $\varepsilon(\lambda,\mu,0)$. Note that $\varepsilon(\lambda_{i,j_1},\mu_i,0) < \varepsilon(\lambda_{i,j_2},\mu_i,0)$, the following inequality holds:

$$\delta(\varepsilon(\lambda_{i,j_2},\mu_i,0) - p_{MIN})\frac{\psi(1)-\psi(0)}{1-\delta(1-\alpha)p_{MIN}} \geq \frac{c}{\alpha}. \quad (16)$$

According to the one-shot deviation principle, (16) is a sufficient and necessary condition for σ^{κ}_{i,j_1} to be the optimal strategy of h_{i,j_2}. Given the fact that σ^{κ}_{i,j_2} is unique, we should have $\sigma^{\kappa}_{i,j_1}(1) = \sigma^{\kappa}_{i,j_2}(1)$, which leads to a contradiction. Hence, this statement follows.

(iv) The proof of this statement follows the same idea as statement (iii) and is omitted here. ∎

Remark: Proposition 1 provides several important insights. Statement (i) proves that by following the optimal strategy, a worker $h_{i,j}$'s choice at each rating is binary: he either fulfills his requested service rate with $a = \lambda_{i,j}$ or does not provide any service with $a = 0$. Statement (ii) proves that a worker's incentive to provide services monotonically decreases with his group's rating. Statement (iii) shows that for two workers within the same group, the worker who receives a larger requested service rate has stronger incentives to provide services, because his deviation has a higher probability to be observed by the community designer and thus get punished. With the same idea, statement (iv) shows that for two workers who receive the same requested service rate, the worker from the group with a smaller aggregated requested service rate has a stronger incentive to provide services. These properties significantly simplify the structure of the optimal service provision strategy of an individual worker, and can greatly facilitate our subsequent analysis of sustainable and optimal rating protocols.

B. Sustainable Rating Protocols

According to Proposition 1, an arbitrarily selected rating protocol κ may not effectively incentivize service provision. In this section, we discuss how to design a rating protocol under which every worker finds it within his self-interest to provide services. Such a rating protocol is called a *sustainable rating protocol* because self-interested strategic workers are willing to comply with it.

Definition 2 (Sustainable Rating Protocol): A rating protocol κ is sustainable if and only if $\sigma^{\kappa}_{i,j}(\theta) = \lambda_{i,j}, \forall\theta, \forall i, j$. ∎

Therefore, under a sustainable rating protocol, providing services is always individually optimal for each worker in any period. In this case, the *social welfare* of the community, which is defined as the average sum service rate provided by all workers in each period, i.e. $J \triangleq \frac{1}{L}\lim_{L\to\infty}\sum_{t=1}^{L}\sum_{i\in\mathcal{N}}\sum_{j\in\mathcal{M}_i} a^{(t)}_{i,j}$, achieves its optimum at $J^* = \sum_{i\in\mathcal{N}}\mu_i$.

In the rest of this section, we derive the sufficient and necessary conditions for a rating protocol to be sustainable.

When a rating protocol is sustainable, a worker $h_{i,j}$ adopts an optimal strategy $\sigma^{\kappa}_{i,j}$ with $\sigma^{\kappa}_{i,j}(1) = \sigma^{\kappa}_{i,j}(0) = \lambda_{i,j}$. His expected long-term utilities can be expressed as follows:

$$\begin{aligned} &V_{i,j}\left(1 \mid \kappa, \sigma^{\kappa}_{i,j}\right) \\ &\quad= \psi(1)\lambda_{i,j} - c\lambda_{i,j} + \delta\left((1-p_{MIN}\alpha)V_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right. \\ &\qquad\qquad \left.+p_{MIN}\alpha V_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right) \\ &V_{i,j}\left(0 \mid \kappa, \sigma^{\kappa}_{i,j}\right) \\ &\quad= \psi(0)\lambda_{i,j} - c\lambda_{i,j} + \delta\left((1-p_{MIN})V_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right. \\ &\qquad\qquad \left.+ p_{MIN}V_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right). \end{aligned} \quad (17)$$

Therefore, we have

$$V_{i,j}(1 \mid \kappa,\sigma^{\kappa}_{i,j}) - V_{i,j}(0|\kappa,\sigma^{\kappa}_{i,j}) = \frac{(\psi(1)-\psi(0))\,\lambda_{i,j}}{1-\delta p_{MIN}(1-\alpha)}. \quad (18)$$

According to the one-shot deviation principle, if worker $h_{i,j}$ unilaterally deviates from $\sigma^{\kappa}_{i,j}$ at rating 1, his expected long-term utility becomes

$$\begin{aligned} &V'_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right) \\ &\quad= \psi(1)\lambda_{i,j} + \delta\left((1-\varepsilon(\lambda_{i,j},\mu_i,0)\alpha)V_{i,j}(1 \mid \kappa,\sigma^{\kappa}_{i,j})\right. \\ &\qquad\qquad \left.+\varepsilon(\lambda_{i,j},\mu_i,0)\alpha V_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right). \end{aligned} \quad (19)$$

Similarly, his expected long-term utility when unilaterally deviating at rating 0 is

$$\begin{aligned} &V'_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right) \\ &\quad= \psi(0)\lambda_{i,j} + \delta\left((1-\varepsilon(\lambda_{i,j},\mu_i,0))\,V_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right. \\ &\qquad\qquad \left.+\varepsilon(\lambda_{i,j},\mu_i,0)V_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right)\right) \end{aligned} \quad (20)$$

By solving the following two inequalities:

$$\begin{aligned} V_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right) &\geq V'_{i,j}\left(1 \mid \kappa,\sigma^{\kappa}_{i,j}\right) \\ V_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right) &\geq V'_{i,j}\left(0 \mid \kappa,\sigma^{\kappa}_{i,j}\right), \end{aligned} \quad (21)$$

we obtain the region of the design parameters, with which the rating protocol can sustain the incentive of worker $h_{i,j}$ to provide services. Then by combining the regions obtained for all workers, the sufficient and necessary condition for a rating protocol to be sustainable is derived, which is formalized in the next theorem.

Theorem 1: A rating protocol κ is sustainable if and only if

$$\frac{c(1-\delta\underline{p}(1-\alpha))}{\alpha\delta\left(\psi(\underline{1})-\psi(0)\right)} \leq \min_{i,j}\frac{(p_{MAX}-p_{MIN})\lambda_{i,j}}{u_i}. \quad (22)$$

Proof: Solving the inequality equation set (21), we have that a worker $h_{i,j}$ always provides services in his optimal strategy if the following inequality holds

$$(\varepsilon(\lambda_{i,j},\mu_i,0)-p_{MIN}) \geq \frac{c\left(1-\delta p_{MIN}(1-\alpha)\right)}{\alpha\delta\left(\psi(0)-\psi(1)\right)}. \quad (23)$$

Substituting (7) into the LHS of (23) and taking the intersection among all workers, we have (22). ■

Theorem 1 reveals how the design parameters influence the sustainability of a rating protocol, which is further highlighted in the following corollary.

Corollary 1: If a rating protocol $\kappa = (\alpha, \psi)$ is sustainable,

(i) a rating protocol $\kappa' = (\alpha', \psi)$ is also sustainable when $\alpha' \geq \alpha$.

(ii) a rating protocol $\kappa'' = (\alpha, \psi'')$ is also sustainable when $\psi''(0) - \psi''(1) \geq \psi(0) - \psi(1)$.

Proof:

(i) If $\kappa = (\alpha, \psi)$, we have $\frac{c(1-\delta p_{MIN}(1-\alpha))}{\alpha\delta(\psi(1)-\psi(0))} \leq \min_{i,j} \frac{(p_{MAX}-p_{MIN})\lambda_{i,j}}{\mu_i}$ according to Theorem 1. Since $\alpha' \geq \alpha$, then $\frac{c(1-\delta p_{MIN}(1-\alpha))}{\alpha'\delta(\psi(1)-\psi(0))} \leq \min_{i,j} \frac{(p_{MAX}-p_{MIN})\lambda_{i,j}}{\mu_i}$ also holds, which gives the conclusion.

(ii) The proof of this statement is similar to statement (i) and omitted here.

■

Remark: Corollary 1 proves that, given a larger punishment probability α, a worker's expected future loss monotonically increases. Since the threat of future loss is higher, stronger incentives are provided to individual workers. Meanwhile, the differential pricing scheme also determines the incentive on service provision. Let $\Delta\psi \triangleq \psi(1) - \psi(0)$ denote the marginal service fee paid to workers who have good ratings, Corollary 1 proves that a higher $\Delta\psi$ also provides stronger incentives to work hard/contribute more high-quality services to an individual worker. This is because the service fee he expects to receive in the future increases if he fulfills his requested service rate in the current period.

Theorem 1 and Corollary 1 prove that in order to provide the strongest incentive to provide services, the rating protocol should choose the largest punishment probability as $\alpha = 1$ and the largest marginal service fee as $\Delta\psi = \bar{q} - c$. However, a sustainable rating protocol still induces significant efficiency loss without careful choices of the design parameters, as discussed in the next section. Hence, $\alpha = 1$ and $\Delta\psi = \bar{q} - c$ are not necessarily the optimal designs.

C. Design Problem Reformulation

In this section, we revisit the optimal rating protocol design problem. Given the optimal strategies of individual workers, the rating of each group i, denoted by θ_i, evolves as a Markov chain, whose transition probability $p_{\sigma_i^\kappa}(\theta'|\theta)$ is determined by the optimal strategies of all workers residing in it, i.e. $\sigma_i^\kappa \triangleq (\sigma_{i,j}^\kappa)_{h_{i,j}\in\mathcal{M}_i}$. Let $\{\eta_i(\theta)\}_{\theta\in\Theta}$ denote the stationary distribution of this Markov chain, the average revenue that this group obtains in one period is $u_\kappa(i) \triangleq \mu_i \sum_{\theta\in\Theta} \eta_i(\theta)\psi(\theta)$. We prove in the following proposition that for each individual group, a sustainable rating protocol always achieves a higher revenue than a non-sustainable rating protocol.

Proposition 2: Given a sustainable rating protocol κ and a non-sustainable rating protocol κ', the following inequality holds for all groups:

$$u_\kappa(i) > u_{\kappa'}(i), \forall i. \tag{24}$$

Proof: This proof can be directly obtained from Proposition 1 and is omitted here. ■

Proposition 2 proves that the sum service fee U of all workers is always maximized by a sustainable rating protocol. Therefore, by optimizing its revenue, the community designer also maximizes the sum service rate in the community at the value $J^* = \sum_{i\in\mathcal{N}} \mu_i$. As a result, solving the protocol design problem (3) is equivalent to finding the sustainable rating protocol that maximizes U. Given the stationary rating distribution $\{\eta_i(\theta)\}_{\theta\in\Theta}$, (3) can be rewritten as follows:

$$\begin{aligned} &\max_{\alpha,\psi} U_\kappa \triangleq \sum_{i\in\mathcal{N}} \mu_i \sum_{\theta\in\Theta} \eta_i(\theta)\psi(\theta) \\ &s.t.\ \sigma_{ij}^\kappa(0) = \sigma_{ij}^\kappa(1) = \lambda_{i,j},\ \forall i, j. \end{aligned} \tag{25}$$

IV. Optimal Rating Protocols

In this section we investigate the design of the optimal rating protocol, i.e. the optimal rating scheme and the optimal pricing scheme, in order to solve the problem (23). For an improved presentation, we first derive the stationary distribution $\{\eta_i(\theta)\}_{\theta\in\Theta}$ of a group's rating.

Given a sustainable rating protocol κ, each worker always choose to provide services to fulfill his requested service rate in any period, i.e. $\sigma_{i,j}^\kappa(\theta) = \lambda_{i,j}, \forall i, j, \forall\theta$. The transition probabilities of a group's rating can be expressed as follows:

$$\begin{aligned} &p_{\sigma_i^\kappa}(0 \mid 0) = p_{MIN} p_{\sigma_i^\kappa}(1 \mid 0) = 1 - p_{MIN}, \\ &p_{\sigma_i^\kappa}(0 \mid 1) = p_{MIN}\alpha p_{\sigma_i^\kappa}(1 \mid 1) = 1 - p_{MIN}\alpha. \end{aligned} \tag{26}$$

With simple manipulations, the stationary distribution $\{\eta_i(\theta)\}_{\theta\in\Theta}$ is derived as follows:

$$\begin{aligned} \eta_i(1) &= \frac{1 - p_{MIN}}{1 - p_{MIN}(1-\alpha)} \\ \eta_i(0) &= \frac{p_{MIN}\alpha}{1 - p_{MIN}(1-\alpha)}. \end{aligned} \tag{27}$$

It is interesting to note that under a sustainable rating protocol, the stationary distribution is independent of a group's aggregated requested service rate and hence, the ratings of all groups evolve following the same Markov chain. That is, $\eta_i(\theta) = \eta_{i'}(\theta), \forall i, i' \in \mathcal{N}$ and $\forall\theta$.

The design of the optimal rating protocol consists of two steps. In the first step, we investigate the optimal design of the punishment probability α, given a fixed pricing scheme ψ. The result is denoted by $\alpha^{\#}(\psi)$. In the second step, we optimize the pricing scheme to obtain the optimal rating protocol $\kappa^* = (\alpha^*, \psi^*)$, where $\alpha^* = \alpha^{\#}(\psi^*)$.

Substituting (27) into $u_\kappa(i)$, we have

$$\begin{aligned} u_\kappa(i) &= \mu_i \left(\frac{1-p_{MIN}}{1-p_{MIN}(1-\alpha)}\psi(1) + \frac{p_{MIN}\alpha}{1-p_{MIN}(1-\alpha)}\psi(0) \right) \\ &= \mu_i \left(\frac{1 - p_{MIN}}{1 - p_{MIN}(1-\alpha)}\psi(1) \right. \\ &\qquad \left. + \frac{p_{MIN}\alpha}{1 - p_{MIN}(1-\alpha)}(\psi(1) - \Delta\psi) \right) \\ &= \mu_i \left(\psi(1) - \frac{1 - p_{MIN}}{1 - p_{MIN}(1-\alpha)}\Delta\psi \right). \end{aligned} \tag{28}$$

In the above equation, $\psi(1)$ represents the basic service fee, while $\frac{1-p_{MIN}}{1-p_{MIN}(1-\alpha)}\Delta\psi$ represents the punishment imposed by the community designer in order to incentivize the service provision from individual workers. From Corollary 1, it is known that the designed punishment, which monotonically increases with both $\Delta\psi$ and α, needs to be sufficiently large in order to provide sufficient incentives. However, this punishment also cannot be too large since it reduces the revenue of the community (by reducing the service fee that workers of rating 0 receive). With this idea, we first derive in the next theorem the optimal punishment probability α, given a fixed pricing scheme.

Theorem 2: Given a fixed pricing scheme ψ, the punishment probability α that maximizes the revenue of the community can be solved as follows:

$$\alpha^{\#}(\psi) = \max\left\{\frac{1-\delta p_{MIN}}{\delta}\frac{c}{Y\Delta\psi - p_{MIN}c, 0}\right\}. \tag{29}$$

Proof: By manipulating the inequality (22), it can be derived that in order to make the resulting rating protocol sustainable, the following inequality should hold for α:

$$\alpha \geq \frac{1-\delta p_{MIN}}{\delta}\frac{c}{Y\Delta\psi - p_{MIN}c}. \tag{30}$$

Also, it is can observed from (28) that u_i monotonically decreases with α and hence, the optimal α takes the value that makes the inequality (30) binding. That is,

$$\alpha^{\#}(\psi) = \frac{1-\delta p_{MIN}}{\delta}\frac{c}{Y\Delta\psi - p_{MIN}c}. \tag{31}$$

Since α also cannot be negative, (29) is derived. ■

Substituting (29) into the design problem (25), the optimal rating protocol that maximizes the revenue of the community is derived in the next theorem.

Theorem 3: The optimal rating protocol that solves (25) can be designed as follows: $\psi^*(1) = \frac{\bar{q}}{1+\rho}$, $\psi^*(0) = \max\left\{c^d, \max\left[\begin{array}{l}\frac{\bar{q}}{1+\rho} - \frac{c}{\delta},\\ \frac{\bar{q}}{1+\rho} - p_{MIN}cY\left(1-p_{MIN}+\frac{1}{\delta}\right)\end{array}\right]\right\}$, $\alpha^* = \alpha^{\#}(\psi^*)$.

Proof: First, it is obvious that $\psi^*(1) = \frac{\bar{q}}{1+\rho}$ such that the basic service fee is maximized. Substituting (29) into (28) and taking the first-order derivative over $\Delta\psi$, we have the optimal marginal service fee to be

$$\Delta\psi^* = p_{MIN}cY\left(1 - p_{MIN} + \frac{1}{\delta}\right). \tag{32}$$

Meanwhile, since $\alpha \leq 1$, we should have $\Delta\psi^* \leq \frac{c}{\delta}$. Also noticing that $\psi^*(0) \geq c$, we have

$$\psi^*(0) = \max\left\{c, \max\left[\begin{array}{l}\frac{\bar{q}}{1+\rho} - \frac{c}{\delta},\\ \frac{\bar{q}}{1+\rho} - p_{MIN}cY\left(1-p_{MIN}+\frac{1}{\delta}\right)\end{array}\right]\right\}$$

and thus Theorem 3 follows. ■

Remark: Theorem 3 derives the optimal rating protocol design. It is proven that in order to maximize the revenue of the community, both the marginal service fee $\Delta\psi$ and the punishment probability α should be selected as the smallest values that enforce sufficient incentives for individual workers to provide services. If $\Delta\psi$ and α are too small, workers lose their incentives to provide services since the punishment enforced on their deviations is not sufficiently large, whereas if $\Delta\psi$ and α are too large, significant revenue loss is incurred due to the imperfect monitoring of service provision quality.

V. Impact of the Group Size

Different from the existing personal rating protocols (e.g. [21], [22]), the size of a group (i.e. the number of workers managed by a group operator) plays an important role in the design of collective rating protocols, which determines the incentive of individual workers as well as the revenue of the community. In this section, we specifically investigate how the group size impacts the efficacy of the rating protocol.

First, we analyze how the group size impacts the sustainability of the rating protocol. The following proposition derives an upper bound on the maximum group size that any sustainable rating protocol can support.

Proposition 3: No sustainable rating protocol exists if there is a group i with $M_i \geq \frac{\delta}{c}(p_{MAX} - p_{MIN})(\frac{\bar{q}}{1-\rho} - c)$.

Proof: From Theorem 1, it is known that given a rating protocol $\kappa = \{\alpha, \psi\}$, it is sustainable if and only if the following inequality holds for every worker $h_{i,j}$:

$$\frac{(p_{MAX} - p_{MIN})\lambda_{i,j}}{\mu_i} \geq \frac{c\,(1-\delta p_{MIN}(1-\alpha))}{\alpha\delta\Delta\psi}. \tag{33}$$

It is easy to observe that the RHS of (33) monotonically decreases with α and $\Delta\psi$ and hence achieves its minimum at $\frac{c}{\delta(\bar{q}/(1-\rho)-c)}$ when $\alpha = 1$ and $\Delta\psi = \frac{\bar{q}}{1-\rho} - c$. Meanwhile, there is at least one worker $h_{i,j} \in \mathcal{M}_i$ such that $\frac{\lambda_{i,j}}{\mu_i} \leq \frac{1}{M_i}$. Suppose $M_i \geq \frac{\bar{\lambda}e^{-\beta T}}{cT}(p_{MAX} - p_{MIN})(\frac{\bar{q}}{1-\rho} - c)$, for the considered worker $h_{i,j}$, we have

$$\begin{aligned}\frac{(p_{MAX} - p_{MIN})\lambda_{i,j}}{\mu_i} &\leq (p_{MAX} - p_{MIN})\frac{1}{M_i}\\ &\leq (p_{MAX} - p_{MIN})\frac{c}{\delta}\\ &\quad\times\frac{1}{(p_{MAX} - p_{MIN})\left(\frac{\bar{q}}{(1-\rho)-c}\right)}\\ &= \frac{c}{\delta}\frac{1}{\left(\frac{\bar{q}}{(1-\rho)-c}\right)}.\end{aligned} \tag{34}$$

Therefore, the inequality of (34) cannot hold for the worker $h_{i,j}$ and the rating protocol is not sustainable. Because the rating protocol selected in the above analysis is arbitrary, it can be concluded that no sustainable rating protocol can be designed and hence this proposition follows. ■

Remark: Proposition 3 reveals an important principle for the design of online communities: when the size of a group is too large, there are always some workers who do not have sufficient incentive to provide services and thus no sustainable rating protocol can be designed in this case. The size of a group thus cannot be arbitrarily large. The intuition behind this proposition is as follows: with a larger group size and thus a larger aggregated requested service rate from this group, it becomes more difficult to detect by the community designer if a worker from this group does not provide high-quality services, thereby giving workers in this group stronger incentives to do so.

Proposition 3 illustrates the negative effect of a large group size on the incentive provision. This result is further highlighted in the next corollary.

Corollary 2: (i) If a rating protocol κ is sustainable in an online community with a worker collection $\{\mathcal{M}_i\}_{i=1}^N$, then κ is also sustainable in an online community with a worker collection $\{\mathcal{M}_i'\}_{i=1}^N$, where $\mathcal{M}_i' \subseteq \mathcal{M}_i, \forall i$.

(ii) If a rating protocol κ is not sustainable in an online community with a worker collection $\{\mathcal{M}_i\}_{i=1}^N$, then κ is also not sustainable in an online community with a worker collection $\{\mathcal{M}_i''\}_{i=1}^N$, where $\mathcal{M}_i'' \supseteq \mathcal{M}_i, \forall i$.

Proof: Both statements are the straightforward results of Theorem 1 and the proofs are omitted here. ■

Remark: Corollary 2 provides an important guideline in designing the sustainable rating protocol for online communities. It shows that, if a rating protocol is sustainable in an online community, it remains sustainable by removing existing workers from the community. On the contrary, if a rating protocol is unsustainable, it remains unsustainable by adding new workers to the community. Therefore, the incentive of an individual worker to provide service monotonically decreases with the group size.

In the rest of this section, we use an example to show how the group size impacts the revenue of the community. Specifically, we consider a community where each group has the same size $\tilde{M}$, while each worker receives the same requested service rate $\tilde{\lambda}$. The next proposition illustrates how the optimal revenue obtained from this exemplary community changes with $\tilde{M}$.

Proposition 4: When $M_i = \tilde{M}, \forall i \in \mathcal{N}, \lambda_{i,j} = \tilde{\lambda}, \forall i, j$,

(i) the revenue obtained under the optimal rating protocol is

$$U^*(\tilde{M}) = N\left(\frac{\tilde{\lambda}\bar{q}}{1-\rho}\tilde{M} - \frac{(1-p_{MIN})c}{(p_{MAX}-p_{MIN})\delta}\tilde{M}^2\right). \quad (35)$$

(ii) the revenue is maximized when $\tilde{M} = \frac{\tilde{\lambda}\bar{q}(p_{MAX}-p_{MIN})}{2(1-\rho)(1-p_{MIN})c}$.

Proof:

(i) When $M_i = \tilde{M}, \forall i \in \mathcal{N}$ and $\lambda_{i,j} = \tilde{\lambda}, \forall i, j$, we have $Y = \frac{(p_{MAX}-p_{MIN})\tilde{\lambda}}{\tilde{M}}$. Substituting this into the optimal rating protocol (32), we have $\psi^*(0) = \frac{\bar{q}}{1-\rho} - \frac{c\tilde{M}}{(p_{MAX}-p_{MIN})\tilde{\lambda}\delta}$. Hence, the revenue under the optimal rating protocol can be expressed as:

$$\begin{aligned}
&U^*(\tilde{M})\\
&= \sum_{i\in\mathcal{N}}\mu_i\left(\frac{\psi^*(1)-(1-p_{MIN})}{1-p_{MIN}(1-\alpha^*)(\psi^*(1)-\psi^*(0))}\right)\\
&= N\tilde{M}\tilde{\lambda}\left(\frac{\bar{q}}{1-\rho} - (1-p_{MIN})\frac{c\tilde{M}}{(p_{MAX}-p_{MIN})\tilde{\lambda}\delta}\right)\\
&= N\left(\tilde{\lambda}\frac{\bar{q}}{1-\rho}\tilde{M} - \frac{(1-p_{MIN})c}{(p_{MAX}-p_{MIN})\delta}\tilde{M}^2\right). \quad (36)
\end{aligned}$$

(ii) By taking the first-order derivative over $U^*(\tilde{M})$, it is derived that $U^*(\tilde{M})$ achieves its maximal value when $\tilde{M} = \frac{\lambda\bar{q}(p_{MAX}-p_{MIN})\delta}{2(1-\rho)(1-p_{MIN})c}$. ■

Remark: Proposition 4 proves that the revenue of an online community cannot arbitrarily increase with the group size. The underlying intuition is that when the group size is large, a stronger punishment probability α and a larger marginal service fee $\Delta\psi$ need to be imposed in order to meet the incentive compatibility constraints of workers. Hence, there is always an upper bound on the group size to maximize the revenue—when the group size is too large, the revenue of the community starts to decrease.

VI. Illustrative Results

In this section, we provide numerical results to illustrate the performance of our proposed rating protocol in compelling individual workers to provide their services. We consider an online community for data sharing with $N = 5$ groups, where workers provide data uploading services to requesters. Throughout the experiments, we fix the following parameters: $p_{MAX} = 0.4$ and $\bar{q} = 3$. We also set that each group contains 50 workexrs. In Section VI-A–VI-D, we set that each worker receives a constant requested service rate (i.e. the size of the requested download data per period) across periods, which is uniformly distributed in the range $[1MB, 5MB]$. However, it should be noted that different workers can receive different requested service rates.

A. Comparison With Benchmark Protocols

In the first experiment, we examine the performance of the optimally design rating protocol, i.e. the rating protocol $\kappa^* = (\alpha^*, \psi^*)$ designed according to Theorem 3 which solves the optimization problem (25). It should be noted that in the experiments, the workers are playing their optimal strategies against the rating protocol, i.e. $\sigma_{i,j}^{\kappa^*}$. For the comparison, the following four benchmark protocols are considered:[6]

1) *Collective Tit-for-Tat strategy (CTFT)*:[7] The rating of a group is set to 0 once a negative signal $s = 0$ is observed and is set to 1 once a signal $s = 1$ is observed. The community designer adopts a fixed pricing scheme: $\psi(1) = \bar{q}/(1+\rho)$ and $\psi(0) = c$.
2) *Collective Tit-for-K-Tat strategy (CTFKT)*: The rating of a group is set to 1 once a signal $s = 1$ is observed and is set to 0 once the signal $s = 0$ is observed in K consecutive periods. The pricing scheme adopted is the same as that in Tit-for-Tat.
3) *Collective Trigger strategy (CTS)*: The rating of a group starts at 1. When K negative signals are observed, the group's rating is set to 0 for all the subsequent periods. The pricing scheme adopted is the same as that in Tit-for-Tat.
4) *Rating independent pricing (RIP)*: The community designer adopts a flat-rate pricing scheme: the service fee for

[6]It should be noted that Tit-for-Tat, Tit-for-K-Tat, and Trigger strategies are originally proposed as personal rating protocols [26]. In this experiment, however, we extend these strategies to operate as collective rating protocols to enable a direct comparison with our proposed incentive protocol. Without our extensions, these protocols suffer from severe limitations when applied as incentive schemes in the considered online communities.

[7]Here we use the name Tit-for-Tat to represent the fact that a punishment is given to a group every time a negative signal is observed upon it and vice versa. It should be noted that the Tit-for-Tat strategy used here is different from the Tit-for-Tat strategy employed in the BitTorrent system [2], which is based on direct reciprocity.

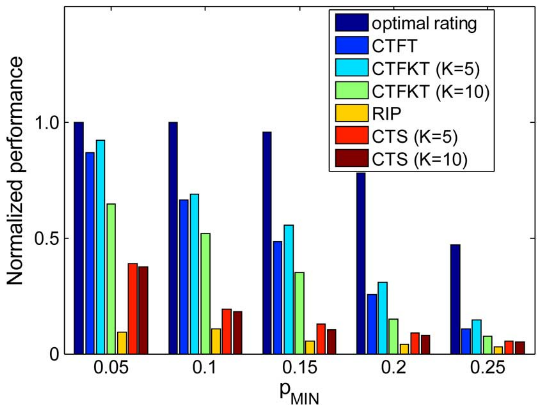

Fig. 1. Normalized performances against the error proability p_{MIN}.

providing one unit service is always $\bar{q}/(1+\rho)$, regardless of the ratings.

We assume that the cost for unit data sharing is $c = 10$. The performance of the community is measured by the social welfare J. For better illustrations, the plotted performances are all normalized by the maximum social welfare J^*, which is achieved when all workers fulfil their requested service rates.

Fig. 1 plots the performances after running the community for $L = 10000$ periods, when p_{MIN} ranges from 0.05 to 0.25.[8] First, it can be noted that the optimal rating protocol outperforms all four benchmarks by achieving, on average, a significantly higher sum service rate that is close to the social optimum J^* (i.e. the normalized performance is close to 1 in Fig. 1). The main reason behind this phenomenon is that, without properly designing the rating protocol, individual workers do not have sufficient incentives to provide their services. Meanwhile, the performance gap between the optimal rating protocol and the benchmarks becomes more significant as p_{MIN} decreases. This is due to the fact that a higher monitoring error probability p_{MIN} introduces a higher probability for an individual worker to be punished even if he fulfills his requested service rate, which reduces his incentive on the service provision. Importantly, it can be observed that although the collective Tit-for-tat, Tit-for-K-Tat, and Trigger strategies implement the idea of collective rating, their performance is significantly worse than that obtained by the optimal rating protocol. This shows the importance of jointly and rigorously designing the rating scheme and the pricing scheme and the loss in performance due to inadequate punishments in the presence of imperfect monitoring.

Fig. 2(a) illustrates the impact of the discount factor δ on the performances of various protocols. As δ increases, a worker puts a higher weight on his future utility relative to his instant utility (at the current period). Hence, with a larger δ, it is easier to give incentives to provide service using future rewards and punishments through the rating protocol, which is proven by the significantly better performance of the optimally designed rating protocol, compared with the others. Fig. 2(b) plots the performances of various protocols against the service cost c. As c increases, the one-period utility received by a worker monotonically decreases and so is his incentive. Therefore, it becomes more difficult to give incentives to provide service with a larger c.

[8]Note that in the figures, we remove the prefix "collective" from the names of all benchmark protocols for simpler illustration.

B. The Performance in Heterogeneous Communities

So far, we have assumed that the workers in the community are homogeneous and have the same unit service cost c and the same discount factor δ. In the next experiment, we investigate how the proposed rating protocol performs in a heterogeneous community where workers incur different unit service costs and have different discount factors.

We set $p_{MIN} = 0.1$, and consider that the unit service cost c follows a truncated Gaussian distribution with a mean value 10 and a variance σ_c. Also, we assume that δ follows a Gaussian distribution with a mean value 0.8 and a variance σ_δ. Fig. 3 portrays the result. It can be observed that our proposed rating protocol maintains a high performance in heterogeneous communities. This is because both α^* and ψ^* in the optimally designed protocol satisfy the sustainability condition (22) with the inequality being strict. Hence, the optimal rating protocol can "tolerate" small variations on the parameters without significantly degrading its performance.

To further illustrate the "robustness" of the rating protocol against the parameters, we design the optimal rating protocol given $c = 10$ and $\delta = 0.8$ and investigate how the performance of this protocol changes when c and δ vary. It could be noted from Fig. 4(a) that although the rating protocol is optimally designed for $c = 10$, it remains sustainable when $c < 10$ and when $c \in (10, 12]$. Therefore, the social welfare of the community can be sustained at its optimum value J^* within this region. A similar phenomenon can be observed from Fig. 4(b), in which the optimal rating protocol remains sustainable when $\delta \in [0.6, 1.0]$.

C. The Impact of Design Parameters

In the next experiment, we investigate how the design parameters α and ψ impact the community's performance. Here we set $p_{MAX} = 0.25$, $c = 10$ and $\delta = 0.8$. Fig. 5(a) shows how the social welfare (normalized by J^*) changes against the value of α, when the pricing scheme is fixed with $\psi(1) = 3$ and $\psi(0) = 1$. When α is small, the punishment is weak and hence, individual workers do not have sufficient incentives to provide services. As α increases, the incentive on the service provision is constructed and the sum service rate gradually increases. When α is sufficiently large, the rating protocol becomes sustainable with all workers providing services. A similar phenomenon is observed in Fig. 5(b), which plots the change in the sum service rate against $\Delta\psi$ for a fixed $\alpha = 0.8$. The social welfare monotonically increases with $\Delta\psi$, as the workers' incentives to provide service increase.

D. The Impact of the Group Size

In this part, we consider a homogeneous community where each group operator manages M workers and each worker receives a requested service rate $\lambda = 2$. Here we set the error probability $p_{MIN} = 0.1$. Fig. 6(a) plots the change of the (normalized) social welfare against the group size M. As shown in Proposition 3, when M is small, the optimal rating protocol can give workers' incentives to provide services. In this case, the social welfare linearly grows with the group size. However, when $M > 25$, the workers lose the incentive to provide services and the social welfare significantly decreases. Fig. 6(b)

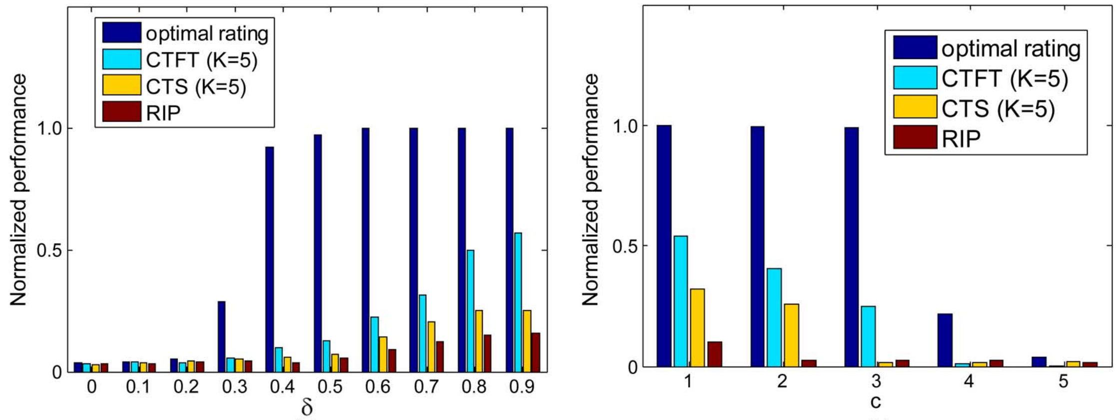

(a) (b)

Fig. 2. Normalized performances against (a) the discount factor δ; (b) the service cost c.

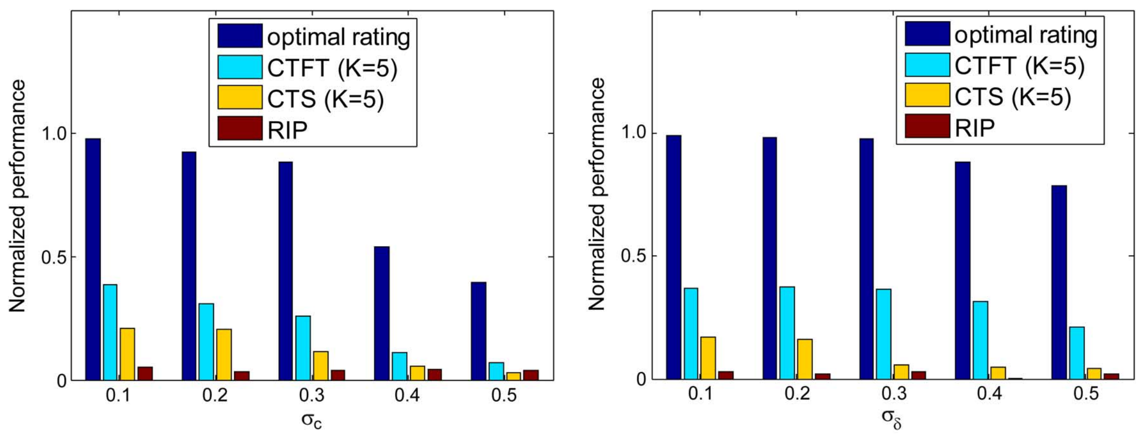

(a) (b)

Fig. 3. Normalized performances against the variances σ_c and σ_δ.

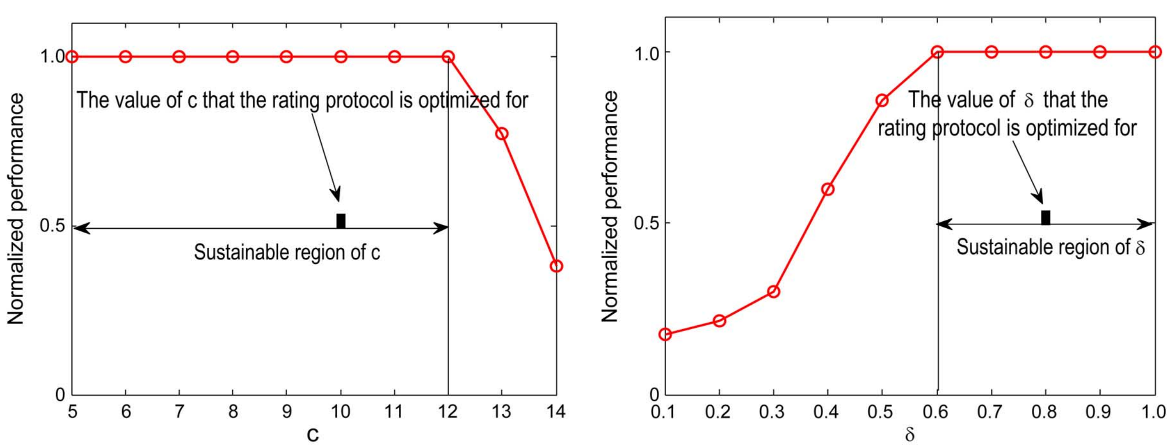

(a) (b)

Fig. 4. The sustainable regions of c and δ.

illustrates how the community's revenue changes against M. It is interesting to notice that the optimal value of M (i.e. the group size that maximizes the community's revenue) does not equal the value where the workers start to lose their incentive (i.e. $M = 25$): when $M \in [15, 25]$, even though the optimal rating protocol is still able to sustain workers' incentive on the service provision, it has to deploy a large marginal service fee $\Delta\psi$, which reduces the revenue collected from groups at rating 0 and thus the total revenue of the community.

From Fig. 6, it can be observed that the group size significant impacts the performance of the rating protocol. In the previous analysis, we have assumed that the group size is known by both the workers and the community designer when designing the optimal rating protocol. In the next experiment, we investigate what is the impact on the rating protocol's performance when this assumption does not hold.

For the experiment, we assume that the group size $M = 25$ for each group. The workers and the community designer main-

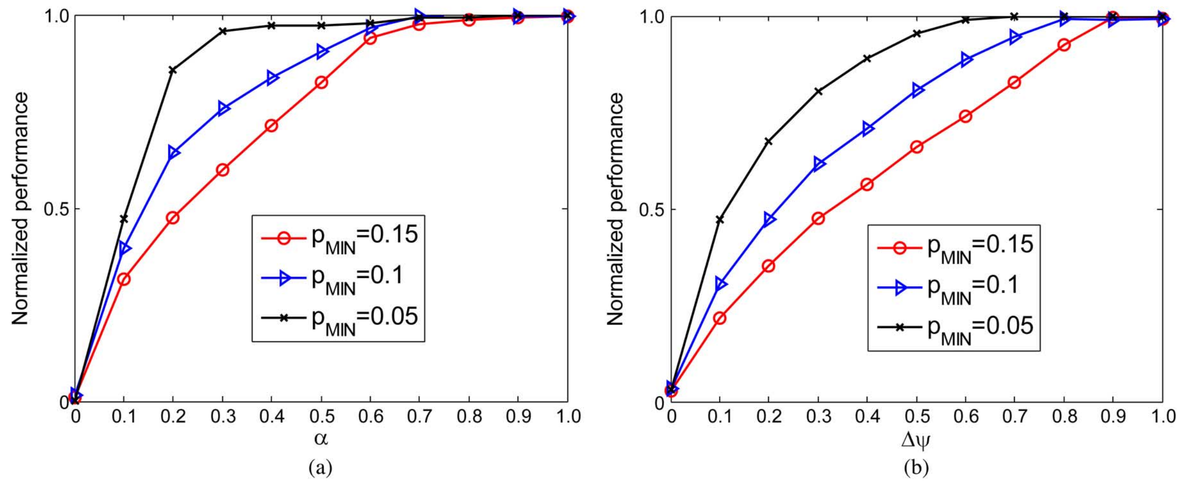

Fig. 5. The community's performance against (a) the punishment probability α; (b) the marginal service fee $\Delta\psi$.

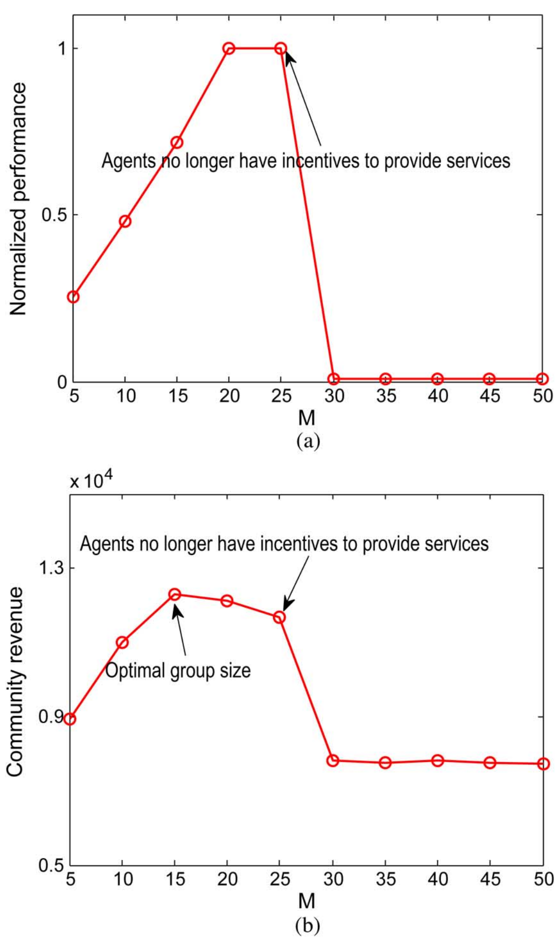

Fig. 6. The community's normalized performance and revenue against the group size M.

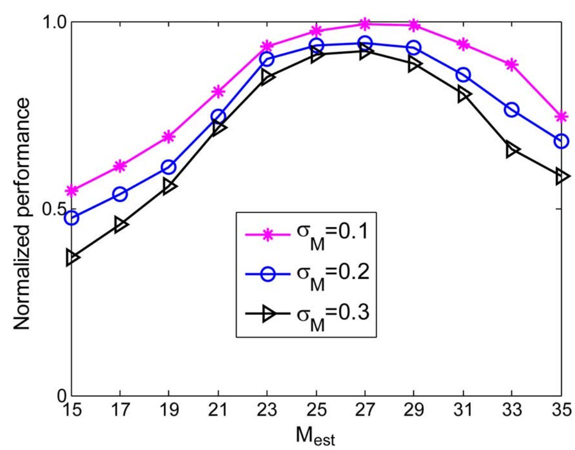

Fig. 7. The impact of inaccurate group size estimation.

tain inaccurate estimates about the value of$M$. Specifically, the group size estimation of an individual worker follows a truncated Gaussian distribution with a mean value 25 and a variance σ_M. Meanwhile, the estimate of the community design is denoted by a variable M_{est} and can be any positive integer. Fig. 7 illustrates the performance of the rating protocol, which is optimally designed based on the (inaccurate) community designer's group size estimation M_{est}. To better highlight the effect of not knowing the exact value of M, we normalize the performance against the social welfare that is achieved when both the workers and the community designer know the exact value of M. It can be observed that the rating protocol delivers close-to-optimal performances when M_{est} is close to 25. That is, when the community designer's estimation is close to the actual group size, the resulting optimal rating protocol is still capable of incentivizing the workers to provide services and thus remains sustainable. Therefore, the rating protocol can perform well even when the exact value of the group size is not perfectly known, as long as the estimate on the group size is sufficiently accurate.

VII. Conclusions and Future Research

In this paper, we designed collective rating protocols to encourage self-interested agents in anonymous online communities with large populations to provide high quality services. By exploiting the ongoing nature of the agents' interaction, we rigorously prove that even though the community designer is profit-seeking and designs the rating protocol to optimize its own revenue (but not the social welfare of the community), the optimal rating protocol (in the sense of revenue maximization) is able to provide individual self-interested agents sufficient incentives to provide services.

This work can be extended in various directions, among which we mention two.

1) This work studies the design and operation of rating protocols in equilibrium. While such assumption is commonly adopted in all the research work rigorously studying rating protocols, e.g. [20], [22]–[24], an important direction for future research is investigating the scenario where the workers continuously learn and form heterogeneous beliefs about the multi-agent environment in which they operate. However, unlike existing work in multi-agent systems in the signal processing area, the agents are now strategic and they can build higher order beliefs about each other (i.e. beliefs about beliefs of others etc.) and aim to influence the decision making of other agents. This can result in very complicated hierarchies of beliefs and complex learning, which can lead to an important but complicated new research agenda.
2) In this work, we assume that the formation of groups is determined exogenously, by the community designer. However, in many environments—mobile networks, femto-cell networks, cognitive radio users, sensor networks etc.—the agents are self-organizing and the groups are formed endogenously by the group formation decisions made by agents. Studying the interplay between the optimal design of the rating protocol and endogenous group formation forms another interesting research direction. In this paper we went one step in that direction by considering the impact of the group size on the rating protocol. However, an entire research agenda in this direction lies ahead.

REFERENCES

[1] H. Park and M. van der Schaar, "Evolution of resource reciprocation strategies in P2P Networks," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1205–1218, 2010.
[2] B. Cohen, "Incentives build robustness in BitTorrent," presented at the P2P Econ. Workshop, Berkeley, CA, USA, 2003.
[3] S. Fazeli-Dehkordy, S. Panahi, and S. Gazor, "Multiple peer-to-peer communications using a network of relays," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 3053–3062, 2009.
[4] L. Adamic, J. Zhang, E. Bakshy, and M. Ackerman, "Knowledge sharing and yahoo answers: Everyone knows something," in *Proc. Int. Conf. World Wide Web (WWW)*, 2008, pp. 655–674.
[5] B. Viswanath, A. Mislove, M. Cha, and K. Gummadi, "On the evolution of user interaction in Facebook," in *Proc. ACM Workshop Online Social Netw.*, 2009, pp. 37–42.
[6] W. Wang, V. Srinivasan, and K. Chua, "Using mobile relays to prolong the lifetime of wireless sensor networks," in *Proc. 11th Annu. Int. Conf. Mobile Comput. Netw.*, 2005, pp. 270–283.
[7] Y. Zhang and M. van der Schaar, "Reputation-based incentive protocols in crowdsourcing applications," in *Proc. IEEE INFOCOM*, 2012, pp. 2140–2148.
[8] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," *Lecture Notes in Comput. Sci.*, vol. 2232, pp. 75–87, 2001.
[9] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster, "Incentive mechanism for large collaborative resource sharing," in *Proc. IEEE Int. Symp. Cluster Comput. Grid*, 2004, pp. 1–8.
[10] V. Vishnumurthy, S. Chandrakumar, and E. Gun Sirer, "KARMA: A secure economic framework for P2P resourcing sharing," presented at the P2P Econ. Workshop, Berkeley, CA, USA, 2003.
[11] Y. Zhang and M. van der Schaar, "Collective rating for online labor markets," in *Proc. Allerton Conf.*, 2012, pp. 371–378.
[12] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proc. ACM SIGCOMM*, 2001, pp. 149–160.
[13] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker, "A scalable content-addressable network," in *Proc. ACM SIGCOMM*, 2001, pp. 161–172.
[14] S. Takahashi, "Community enforcement when players observe partners' past play," *J. Econom. Theory*, vol. 145, no. 1, pp. 42–62, 2010.
[15] G. Ellison, "Cooperation in the prisoner's dilemma with anonymous random matching," *Rev. Econom. Studies*, vol. 61, no. 3, pp. 567–588, 1994.
[16] J. Deb, "Cooperation and Community Responsibility: A Folk Theorem for Random Matching Games With Names," mimeo, 2007.
[17] A. Singh and L. Liu, "TrustMe: Anonymous management of trust relationships in decentralized P2P systems," *Peer-to-Peer Comput.*, pp. 142–149, 2003.
[18] R. Gupta, V. Sekhri, and A. Somani, "CompuP2P: An architecture for internet computing using peer-to-peer networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 11, pp. 1306–1320, 2006.
[19] A. Blanc, Y. Liu, and A. Vahdat, "Designing incentives for peer-to-peer routing," in *Proc. IEEE INFOCOM*, 2005, vol. 1, no. 1, pp. 374–385.
[20] C. Dellarocas, "Reputation mechanism design in online trading environments with pure moral hazard," *Inf. Syst. Res.*, vol. 16, no. 2, pp. 209–230, 2005.
[21] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman, "Reputation systems," *Commun. ACM*, vol. 43, pp. 45–48, 2000.
[22] M. Kandori, "Social norms and community enforcement," *Rev. Econom. Studies*, vol. 59, no. 1, pp. 63–80, 1992.
[23] M. Okuno-Fujiwara and A. Postlewaite, "Social norms and random matching games," *Games Econ. Behav.*, vol. 9, no. 1, pp. 79–109, 1995.
[24] C. Dellarocas, "How often should reputation mechanisms update a trader's reputation profile?," *Inf. Syst. Res.*, vol. 17, no. 3, pp. 271–285, 2006.
[25] Y. Zhang, J. Park, and M. van der Schaar, "Rating protocols for online communities," *ACM Trans. Econom. Comput.*, accepted for publication.
[26] G. Mailath and L. Samuelson, *Repeated Games and Reputations: Long-Run Relationships*. Oxford, U.K.: Oxford Univ. Press, 2006.
[27] P. Bolton and M. Dewatripont, *Contract Theory*. Cambridge, MA, USA: MIT Press, 2004.
[28] B. Caillaud and B. Jullien, "Chicken & egg: Competition among intermediation service providers," *RAND J. Econom.*, vol. 34, no. 2, pp. 309–328, 2003.
[29] D. Lucking-Reiley, "Auctions on the internet: What's being auctioned, and how?," *J. Indust. Econom.*, vol. 48, no. 3, pp. 227–252, 2010.

Yu Zhang (S'08) received the bachelor's and master's degrees from Tsinghua University, Beijing, China, in 2006 and 2008, respectively.

He received the Ph.D. degree from the University of California, Los Angeles, in 2013.

Mihaela van der Schaar (SM'04–F'10) received the Ph.D. degree from Eindhoven University of Technology, The Netherlands, in 2001.

She is now a Professor with the Electrical Engineering Department, University of California, Los Angeles.