Reposting Service in Online Social Networks: Modeling and Online Incentive Protocols

Shiyuan Zheng*, Xudong Liu*, Hong Xie†, John C.S. Lui*
*Department of Computer Science and Engineering, The Chinese University of Hong Kong
†College of Computer Science, Chongqing University
*{syzheng,xdliu21,cslui}@cse.cuhk.edu.hk, †xiehong2018@foxmail.com

2022 IEEE 30th International Conference on Network Protocols (ICNP) | 978-1-6654-8234-9/22/$31.00 ©2022 IEEE | DOI: 10.1109/ICNP55882.2022.9940336

Abstract*—Reposting plays an essential role in visibility boosting in online social networks (OSNs). In this paper, we study the problem of designing "reposting service" in an OSN to incentivize "transactions" between requesters (users who seek to enhance visibility) and suppliers (users who are willing to repost if certain incentives are given), and maximize the *welfare increase* accumulated through a given time horizon. We formulate a mathematical model for reposting which captures various factors like click through rates (CTRs), requesters' valuations and suppliers' costs. We formulate the problem of maximizing the welfare increase via judiciously assigning suppliers to requesters from two aspects: (a) "user-centric" and (b) "platform-centric". The user-centric aspect deals with the situation where requesters and suppliers would collaborate and share valuations and costs. To address the challenge of unknown CTRs, we propose an online learning protocol and achieve a sub-linear regret. The platform-centric aspect corresponds to the scenario where users keep their valuations or costs private. To address the challenges of unknown CTR, valuations and costs, we design an "explore-then-commit" online protocol which can be proved to be truthful. Lastly, we conduct extensive experiments to evaluate the efficiency and effectiveness of the proposed protocols.

I. Introduction

Online Social Networks (OSNs) have become popular venues for users to create contents and share contents with friends, followers, etc. Boosting the visibility of contents for users, i.e., making the contents reach a larger audience on the OSN has been a heated topic [1]–[4]. Take SoundCloud[1] as an example. It is a music platform powered by a global community of music artists and listeners. An artist can benefit from boosting the visibility of his music record, i.e., making the music heard by more users on SoundCloud. For example, record sales can be increased or the chance to get signed by record companies would be higher. However, one of the main obstacles to getting large visibility is that a content has to compete with many other contents [5] and the effect of "the strong get stronger" makes less popular contents buried by the most popular ones.

Reposting, a content-sharing behavior, plays an essential role in visibility boosting. For example, recently, an Atlanta-based e-sports organization Ghost Gaming posted a twitter: "Enter to win a $2,000 NZXT and a $100 gift card to the Ghost Gaming shop!".[2] They also added that users who retweet this tweet would be able to join the lucky draw which offers above valuable prizes to the winner. That twitter got 8k retweets and 7.5k likes. By contrast, the previous twitters posted by Ghost Gaming not involved in prize-giving reposting only got about 0.1k likes on average. Obviously, this twitter successfully stimulated many users to repost, so that the content could finally reach much more users and get much more likes. One can refer to [6] for more examples that show the power of reposting on Twitter.

Reposting services prevail in OSNs in a variety of forms. For example, RepostExchange[3] is a platform powered by requesters and suppliers from SoundCloud, where requesters are SoundCloud artists who will pay other users for reposting, and suppliers are SoundCloud users who will repost music if certain rewards are given. The payment and the reward are in a form called "credit", which serves as the currency on RepostExchange. Another example is Sina Weibo[4] in China, a microblogging platform similar to Twitter. The platform offers an official service named WeiTask[5], where Weibo users can opt to participate. Requesters can launch recruitment to call for reposting, and the platform would select suppliers. If selected suppliers repost the content, they would receive rewards.

Note that the above-mentioned reposting services use simple heuristic pricing mechanisms, whose efficiency and effectiveness have no theoretical guarantees. Having no theoretical guarantees may lead the system to a "socially sub-optimal" state, which hurts the ecosystem of OSNs. Thus, it is important to design theoretically efficient and effective mechanisms. To the best of our knowledge, few works have studied this problem (please refer to Section VI for a more elaboration). In this paper, we *design protocols to incentivize reposting of contents with theoretical guarantees*. We assume all the users are rational. Specifically, there is a set of requesters who are content creators and want to increase their content visibility by requesting other users to repost. There is also a set of suppliers who are willing to repost requesters' contents when certain reward is offered. A content is visible to the neighbors of the user who posts or reposts it. If a supplier s reposts requester r's contents, the visibility increase of r's content is the marginal increase of the number of users to whom the

[1] https://soundcloud.com
[2] https://twitter.com/GhostEsports/status/1161381978073116673
[3] https://repostexchange.com
[4] https://weibo.com
[5] https://weirenwu.weibo.com

978-1-6654-8234-9/22/$31.00 ©2022 IEEE

content is visible. *Click* is the actual action that users would take after viewing the post, which eventually contributes to the revenue of content creators. Without loss of generality, we use *click* to model a range of beneficial behaviors after viewing the content. For example, it can model a thumbs up to show agreement on the opinion of a post, or purchasing behavior after viewing an advertising post. We use click through rate (CTR) to denote the probability of a click after viewing the post. A requester has a personal unit valuation for each click received and a supplier has a personal unit cost for each click contributed by his reposting. The objective is to maximize the social welfare increase accumulated throughout $T \in \mathbb{N}_+$ rounds, which is the total valuations increase of requesters minus the total cost of suppliers. The decision variable is the *assignment* from suppliers to requesters.

We formulate the decision problem from two different aspects: (a) "user-centric" and (b) "platform-centric". The user-centric aspect deals with the situation where requesters and suppliers collaborate on searching for the optimal assignment among them. In this setting, requesters and suppliers would share their valuations and costs, e.g., they can set up social groups and report their valuations or costs. One challenge in searching for the optimal assignment is that the CTRs are unknown, making the social welfare increase unable to be calculated. We propose an online learning protocol to address this challenge. Our online learning protocol is built on the observation that the outcome of each assignment from suppliers to requesters reveals samples on the CTRs of requesters who are assigned at least one supplier. Note that for those requesters who are assigned no supplier at all, no sample on their CTRs is revealed. This leads to the challenge of balancing the exploration vs. exploitation tradeoff. In our online learning protocol, we employ a combinatorial multi-armed bandit algorithm to address this challenge. We also prove that our proposed protocol enjoys a sub-linear regret. The platform-centric aspect corresponds to the setting where requesters and suppliers do not collaborate and there exists a service like RepostExchange to determine the assignments. In this setting, requesters and suppliers would be asked to report (not necessarily truthfully) their unit valuations and costs to the service. Compared to the user-centric aspect, besides the challenge of unknown CTR, one extra challenge is that requesters and suppliers usually do not have the incentive to report truthfully. We design a truthful "explore-then-commit" online protocol to address these challenges. In the exploration phase, our protocol uses several rounds to estimate the CTRs. In the commit phase, our protocol uses a charging and rewarding scheme which can incentivize requesters and suppliers to report their true unit valuations and costs. We also prove the regret for this truthful online protocol. Lastly, we conduct extensive experiments on real-world datasets to evaluate the efficiency and effectiveness of the proposed protocols. The highlight of our contributions include:

- We formulate the mathematical model and the welfare increase maximization problem for the reposting service from both the user-centric and platform-centric aspects.
- We propose an online learning protocol with a probable sub-linear regret to address the challenge of unknown CTR for the user-centric aspect.
- We propose a truthful "explore-then-commit" online protocol with a sub-linear regret to address the challenge of unknown unit valuations and costs as well as unknown CTRs for the platform-centric aspect.
- We conduct experiments on the efficiency and effectiveness of the proposed protocols over real-world datasets.

The remainder of this paper is organized as follows. Section II presents the mathematical model of reposting service in OSNs and problem formulation. Section III presents the online protocol for user-centric welfare increase maximization. Section IV presents the online incentive protocol for platform-centric welfare increase maximization. Section V shows experiments on real-world OSNs. Section VI gives discussions on related literatures. Section VIII concludes.

II. Model & Problem Formulation

In this section, we first present the mathematical model of reposting service which qualities the visibility, valuations, costs, CTRs, welfare increase, etc. Then we formulate the problem of assigning suppliers to requesters from both user-centric and platform-centric aspects.

A. The Model of Reposting Service

Consider an OSN characterized by a directed and unweighted graph $\mathcal{G} \triangleq (\mathcal{U}, \mathcal{E})$, where $\mathcal{U} \triangleq \{1, \dots, N\}$ denotes a set of $N \in \mathbb{N}_+$ users and $\mathcal{E} \subseteq \mathcal{U} \times \mathcal{U}$ denotes a set of edges between users. On Twitter like OSNs, a direct edge (v, u) represents user v follows user u. On Facebook like OSNs, a friendship link between u and v can be modeled by two directed edges (u, v) and (v, u). The set of incoming neighbors of user $u \in \mathcal{U}$ is denoted by

$$\mathcal{N}_u \triangleq \{v | v \in \mathcal{U}, (v, u) \in \mathcal{E}\}.$$

Content posted by user u is visible to his incoming neighbors $\mathcal{N}_u$, and this content can also be visible to the incoming neighbors of user v, i.e., $\mathcal{N}_v$, if user v reposts it. However, the reposting behavior should not be taken for granted, since reposting is associated with some cost, e.g., the time, the social pressure of potentially annoying their incoming neighbors, etc. Let $\mathcal{R} \subseteq \mathcal{U}$ denote a set of requesters who call for reposting to boost the visibility and are willing to pay for that. Let $\mathcal{S} \subseteq \mathcal{U}$ denote a set of suppliers who are willing to repost a content if some financial reward is provided. We then give the formal definition of visibility.

Definition 1 (Visible set and visibility)**.** *If supplier $s \in \mathcal{S} \cup \{0\}$ reposts a content created by requester $r \in \mathcal{R}$, the visible set of this content is defined as*

$$\mathcal{V}_r(s) \triangleq \mathcal{N}_r \cup \mathcal{N}_s,$$

where $s = 0$ is to model that no supplier reposts this content and correspondingly, $\mathcal{N}_0 = \emptyset$. The visibility of this content is the cardinality of the visible set $|\mathcal{V}_r(s)|$.

Among the visible set of r's content, some users would click on this content. This fraction also can be interpreted as the probability of receiving a click once the content is viewed, and it is called the *click through rate (CTR)*. We assume that this probability only depends on the content creator. Let θ_r denote the CTR of a content created by requester r, which reflects the overall content quality (i.e., attractiveness) of r. Each requester r has a normalized unit valuation $v_r \in [0, 1]$ for each click of his content. Each supplier s has a normalized unit cost $c_s \in [0, 1]$ for each increased click from his incoming neighbors. For ease of presentation, we denote the CTR vector, valuation vector and cost vector respectively as

$$\boldsymbol{\theta} \triangleq (\theta_r : r \in \mathcal{R}), \quad \boldsymbol{v} \triangleq (v_r : r \in \mathcal{R}), \quad \boldsymbol{c} \triangleq (c_s : s \in \mathcal{S}).$$

We consider a total number of $T \in \mathbb{N}_+$ rounds of assigning suppliers to requesters. Each round $t \in [T] \triangleq \{1, \dots, T\}$ is associated with the same set of requesters $\mathcal{R}$ and the same set of suppliers $\mathcal{S}$. For the ease of presentation, we assume that users who are not suppliers, i.e., $\mathcal{U} \setminus \mathcal{S}$, do not repost contents when no payment is made. This assumption only simplifies the calculation of the visible set of content. If users who are not suppliers repost contents, then a content will spread faster in an OSN. In each round t, each requester creates and posts a new content. Without loss of generality, in each round, each requester is allowed to have at most one supplier repost his content, and a supplier is allowed to repost at most one requester's content. Note that this assumption is not a restriction, because the following trick of virtual requesters or virtual suppliers can handle the setting that a requester has multiple contents to repost or a supplier would like to repost multiple contents. In a situation where a requester has multiple contents, one can create multiple virtual copies of this requester while satisfying that each virtual requester has only one content to repost. In a situation where a supplier would like to repost multiple contents, one can create multiple virtual copies of this supplier while satisfying that each virtual supplier would repost only one content. Let $a_{t,r} \in \mathcal{S} \cup \{0\}$ denote the supplier assigned to requester r in round t. Here $a_{t,r} = 0$ models that no supplier is assigned to requester r in round t. Denote the requesters who are assigned a supplier in round t as

$$\mathcal{R}'_t = \{r | r \in \mathcal{R}, a_{t,r} \neq 0\}.$$

Denote the *assignment profile* in round t as

$$\boldsymbol{a}_t = (a_{t,r} : r \in \mathcal{R}).$$

We define a valid assignment profile as follows.

Definition 2 (Valid assignment profiles)**.** *An assignment profile $\boldsymbol{a}_t$ is valid if it satisfies:*

$$|\{r | a_{t,r} = s, r \in \mathcal{R}\}| \leq 1, \quad \forall s \in \mathcal{S}, \forall t \in [T]. \tag{1}$$

Condition (1) states that in a valid assignment profile, each supplier is assigned to at most one requester in each time step. We denote the set of all valid assignment profiles by

$$\mathcal{A} \triangleq \{\boldsymbol{a} | \boldsymbol{a} \text{ satisfies Eq. (1)}\}.$$

Given time step t, we use $S_{(r,a_{t,r})}$ to denote the visibility increase of requester r due to supplier $a_{t,r}$'s reposting, formally

$$S_{(r,a_{t,r})} \triangleq |\mathcal{N}_{a_{t,r}} \setminus \mathcal{N}_r|.$$

Let $V(r, a_{t,r})$ denote the corresponding expected valuation increase of requester r, formally

$$V(r, a_{t,r}) \triangleq S_{(r,a_{t,r})} \theta_r v_r.$$

Let $C(r, a_{t,r})$ denote the corresponding expected cost of supplier $a_{t,r}$ for reposting r's content, formally

$$C(r, a_{t,r}) \triangleq S_{(r,a_{t,r})} \theta_r c_{a_{t,r}}.$$

Then, $V(r, a_{t,r}) - C(r, a_{t,r})$ quantifies the overall expected benefit of the requester-supplier pair $(r, a_{t,r})$ due to assigning supplier $a_{t,r}$ to requester r. Thus, the expected welfare increase associated with assignment profile $\boldsymbol{a}_t$ is denoted by

$$W(\boldsymbol{a}_t) = \sum_{r \in \mathcal{R}} \big(V(r, a_{t,r}) - C(r, a_{t,r})\big).$$

Finally, we denote the expected *welfare increase* accumulated throughout T time steps associated with the assignment profile sequence $\boldsymbol{a}_1, \dots, \boldsymbol{a}_T$ as

$$W_T = \sum_{t \in [T]} W(\boldsymbol{a}_t),$$

which is the objective to maximize in the reposting service.

B. Problem Formulation

We formulate the assignment problem to maximize the cumulative welfare increase from both the user-centric and platform-centric aspects.

User-centric welfare increase maximization. We consider the problem that requesters and suppliers collaborate to maximize their total benefits, i.e., the welfare increase. They would share their unit valuations and unit costs, but they do not know the CTR vector $\boldsymbol{\theta}$. The challenge is to estimate the CTRs and maximize the cumulative welfare increase simultaneously. The optimal assignment profile can be stated as

$$\boldsymbol{a}^* \in \arg\max_{a \in \mathcal{A}} W(\boldsymbol{a}).$$

Our objective is to design a protocol to select a sequence of assignment profiles $(\boldsymbol{a}_t, t \in [T])$ to minimize the regret, which is defined as follows:

$$R_T \triangleq \sum_{t=1}^{T} (W(\boldsymbol{a}^*) - W(\boldsymbol{a}_t)). \tag{2}$$

Minimizing regret R_T is equivalent to maximizing welfare increase W_T, since regret R_T quantifies how well our protocols compared with the optimal assignment which has full knowledge. One can directly compare different methods using regret which can contribute to a better theoretical conclusion.

Platform-centric welfare increase maximization. We consider the problem that requesters and suppliers do not collaborate, i.e., they do not share unit valuations and unit costs. In this case, we design a reposting service for the platform to determine the assignment profiles. Compared with the user-centric setting, one additional challenge is that requesters' valuations and suppliers' costs are unknown to the platform. Our objective is to design a protocol to solicit the unit valuations and unit costs and also estimate the CTRs so that we can maximize the cumulative welfare increase or equivalently minimize the regret. Formally, the protocol is parameterized by charging scheme $\boldsymbol{p}_t \triangleq (p_{t,r} : r \in \mathcal{R})$, and reward scheme $\boldsymbol{q}_t \triangleq (q_{t,s} : s \in \mathcal{S})$, where $p_{t,r}, q_{t,s} \in [0,1], \forall t \in [T]$. More specifically, at time step t, each requester is charged $p_{t,r}$ for each of his clicks increased due to the assignment, and each supplier is rewarded $q_{t,s}$ for each increased click contributed by him. The goal is to design a charging and reward scheme to incentivize requesters and suppliers to report their unit valuations and unit costs truthfully.

III. User-centric Welfare Increase Maximization

In this section, we study the problem of user-centric welfare increase maximization, where requesters and suppliers share their unit valuations and unit cost. We start from a warm-up case where the CTR vector $\boldsymbol{\theta}$ is accessible, then design a protocol to address the challenge of unknown $\boldsymbol{\theta}$.

A. *Offline Optimal Assignment Protocol*

Optimal assignment. We first consider the setting that the CTR vector $\boldsymbol{\theta}$ is known to users and the goal is to design a protocol that can find out the optimal assignment $\boldsymbol{a}^*$ to maximize the welfare increase. This protocol serves as a building block to study the setting that $\boldsymbol{\theta}$ is unknown to users. Note that under this setting with full knowledge there is no need to vary the assignment with time step t, and the welfare increases in different time steps are identical. Thus, we omit the subscript or superscript of t for ease of presentation.

To facilitate the protocol design, we first construct a complete weighted and undirected bipartite graph denoted by $\mathcal{B} = (\mathcal{R}, \mathcal{S}, \boldsymbol{E})$, where $\mathcal{R}$ and $\mathcal{S}$ are two disjoint node sets representing requesters and suppliers respectively, and $\boldsymbol{E} \triangleq [e_{r,s} : r \in \mathcal{R}, s \in \mathcal{S}]$ denotes the weights. The weight of edge (r,s) is set as

$$e_{r,s} = V(r,s) - C(r,s),$$

which is the expected welfare increase of requester-supplier pair (r,s). We define a matching in the bipartite graph $\mathcal{B}$ as follows.

Definition 3 (Matching)**.** *The edge set $\mathcal{M} \subseteq \mathcal{R} \times \mathcal{S}$ of graph $\mathcal{B}$ is a matching if it satisfies that $|\{s|(r,s) \in \mathcal{M}\}| \leq 1, \forall r \in \mathcal{R}$ and $|\{r|(r,s) \in \mathcal{M}\}| \leq 1, \forall s \in \mathcal{S}$.*

The following lemma states the connection between an assignment profile $\boldsymbol{a}$ and a matching in graph $\mathcal{B}$.

Lemma 1. *An assignment profile $\boldsymbol{a}$ can be mapped into a matching $\mathcal{M}(\boldsymbol{a}) = \{(r, a_r)|r \in \mathcal{R}, a_r \neq 0\}$ which satisfies*

$$W(\boldsymbol{a}) = \sum_{(r,s) \in \mathcal{M}(\boldsymbol{a})} e_{r,s}.$$

A matching $\mathcal{M}$ in $\mathcal{B}$ can be mapped into an assignment profile $\boldsymbol{a}(\mathcal{M})$ with $a_r(\mathcal{M}) = \sum_{s \in \mathcal{S}} s \mathbb{I}_{\{(r,s) \in \mathcal{M}\}}$ which satisfies

$$\sum_{(r,s) \in \mathcal{M}} e_{r,s} = W(\boldsymbol{a}(\mathcal{M})).$$

Lemma 1 implies that an effective approach to find the optimal assignment profile is to locate the maximum weighted matching and then map it to an assignment profile $\boldsymbol{a}^*$. Based on this observation, Protocol 1 outlines the protocol `OptAssign` to find the optimal assignment profile. In protocol `OptAssign`, we first construct the complete weighted bipartite graph $\mathcal{B}$. Then we find the maximum weighted matching via the algorithm `MaxWeightMatching`$(\mathcal{R}, \mathcal{S}, \boldsymbol{E})$. There are a variety of implementations of `MaxWeightMatching` and one example is the Hungarian algorithm [7]. Lastly, we map the maximum weighted matching into the optimal assignment.

Protocol 1 `OptAssign`$(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}, \boldsymbol{v}, \boldsymbol{c})$

Construct complete undirected bipartite graph $\mathcal{B} = (\mathcal{R}, \mathcal{S}, \boldsymbol{E})$
$\mathcal{M}^* \leftarrow$`MaxWeightMatching`$(\mathcal{B})$
$a_r(\mathcal{M}^*) = \sum_{s \in \mathcal{S}} s \mathbb{I}_{\{(r,s) \in \mathcal{M}^*\}}$
return assignment profile $(a_r(\mathcal{M}^*) : r \in \mathcal{R})$

Approximate protocol. Since it is computationally expensive to find a maximum weight matching for a large-scale graph using exact algorithms, e.g, $O(N^3)$ if using the Hungarian algorithm, we propose an approximation protocol to improve computational efficiency. Protocol 2 outlines `GreedyAssign`, which uses a greedy approach for approximating the maximum welfare increase. After constructing the bipartite graph, we rank all the edges in descending order by their weights and traverse them according to this order. If an edge (r,s) still exists in graph $\mathcal{B}$, we add it into the matching. Then delete the corresponding requester and supplier, as well as their adjacent edges from graph $\mathcal{B}$. We repeat until there is no edge left in $\mathcal{B}$. Finally, we get a matching and we map it to the corresponding assignment profile which is the result of the approximate protocol.

Protocol 2 `GreedyAssign`$(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}, \boldsymbol{v}, \boldsymbol{c})$

Construct complete undirected bipartite graph $\mathcal{B} = (\mathcal{R}, \mathcal{S}, \boldsymbol{E})$
Initialize $\mathcal{M} = \emptyset$
for all $(r,s) \in \boldsymbol{E}$ in descending order by weights **do**
  if (r,s) is still in $\mathcal{B}$ **then**
    add (r,s) to $\mathcal{M}$.
    delete r and s from $\mathcal{B}$.
  else
    skip this edge and continue.
  end if
end for
$a_r(\mathcal{M}) = \sum_{s \in \mathcal{S}} s \mathbb{I}_{\{(r,s) \in \mathcal{M}\}}$
return assignment profile $(a_r(\mathcal{M}) : r \in \mathcal{R})$

Theorem 1. *Protocol* `GreedyAssign` *(Protocol 2) guarantees an approximation ratio of* $1/2$*. The running time complexity of* `GreedyAssign` *is bounded by* $O(|\boldsymbol{E}| \log |\boldsymbol{E}|)$.

Remark: Theorem 1 states that `GreedyAssign` can achieve a guaranteed approximation ratio of at least $1/2$ for user-centric welfare increase with much smaller time complexity than `OptAssign` which uses an exact search algorithm like the Hungarian algorithm. The technical proof of Theorem 1 is presented in our technical report [8].

B. Online Learning Protocol

Protocol design. Now we use Protocol 1 as a building block to study the assignment problem in the user-centric aspect where the CTR vector $\boldsymbol{\theta}$ is unknown to requesters and suppliers. In this setting, one needs to balance the exploration and exploitation tradeoff. We achieve this tradeoff via the upper confidence bound (UCB) method [9]. At time step t, for each requester, an unbiased estimator of θ_r can be

$$\widehat{\theta}_r = \frac{\sum_{\tau \in [t]} \frac{N^{(\tau)}_{(r,a_{\tau,r})}}{S_{(r,a_{\tau,r})}} \mathbb{I}\{a_{\tau,r} \neq 0\}}{\sum_{\tau \in [t]} \mathbb{I}\{a_{\tau,r} \neq 0\}}, \tag{3}$$

where $N^{(\tau)}_{(r,a_{\tau,r})}$ is the observed number of increased clicks of r received due to $a_{\tau,r}$'s reposting at time step τ. Based on Eq. (3), we apply Hoeffding inequality to derive UCB index for θ_r. Then, we apply these UCB indexes to Protocol 1 for selecting the assignment profile.

Protocol 3 `OnlineAssign`$(\mathcal{R}, \mathcal{S}, \boldsymbol{v}, \boldsymbol{c})$

Input: $\boldsymbol{v}, \boldsymbol{c}, (\mathcal{N}_u : \forall u \in \mathcal{U})$
for $t = 1$ to $|\mathcal{R}|$ **do**
  Select r from $\{r | r \in \mathcal{R}, N_r = 0\}$
  Select a random supplier s
  Select an assignment profile $\boldsymbol{a}_t$:
  $(a_{t,r} = s; a_{t,\bar{r}} = 0, \forall \bar{r} \in \mathcal{R} \setminus \{r\})$
  Observe number of increased clicks $N^{(t)}_{(r,s)}$
  Initialize $N_r = 1$
  Initialize $\bar{\theta}_r = N^{(t)}_{(r,s)} / S_{(r,s)}$
end for
for $t = |\mathcal{R}| + 1$ to T **do**
  Denote $\bar{\boldsymbol{\theta}}^+ = (\bar{\theta}^+_r : r \in \mathcal{R})$
  $\boldsymbol{a}_t \leftarrow$`OptAssign`$(\mathcal{R}, \mathcal{S}, \bar{\boldsymbol{\theta}}^+, \boldsymbol{v}, \boldsymbol{c})$
  Observe $N^{(t)}_{(r,a_{t,r})}$ for $\forall r \in \mathcal{R}'_t$
  for $r \in \mathcal{R}'_t$ **do**
    Update $N_r = N_r + 1$
    Update $\bar{\theta}_r = \left((N_r - 1)\bar{\theta}_r + N^{(t)}_{(r,a_{t,r})} / S_{(r,a_{t,r})} \right) / N_r$
    Update $\bar{\theta}^+_r = \bar{\theta}_r + \sqrt{3 \log t / 2N_r}$
  end for
end for

Protocol 3 outlines details of the above UCB-based method, leading to our protocol `OnlineAssign`. Initially, we use $|\mathcal{R}|$ rounds of forced exploration to obtain samples of the CTR for each requester r. In the meantime, we keep track of some statistics of the assignment history: N_r which is the number of times the requester r has been assigned a supplier till the current round; $\bar{\theta}_r$ which is the empirical mean of sampled CTRs till the current round; and the UCB index $\bar{\theta}^+_r$ which adds an adjustment term to $\bar{\theta}_r$. In the remaining rounds after the forced exploration, we take the UCB index vector $\bar{\boldsymbol{\theta}}^+ = (\bar{\theta}^+_r : r \in \mathcal{R})$ as the input CTR vector of oracle `OptAssign` (Protocol 1) to decide the best assignment profile for this round, and we update N_r, $\bar{\theta}_r$ and $\bar{\theta}^+_r$ for each requester according to the observed clicks.

Protocol analysis. First, the time complexity of Protocol 3 mainly depends on the number of calls of oracle `OnlineAssign` which is smaller than T, as well as the implementation and the time complexity of `OnlineAssign`. Besides, the following theorem states the performance guarantee of Protocol 3.

Theorem 2. *For all* $T > |\mathcal{R}|$*, protocol* `OptAssign` *(Protocol 3) achieves a sub-linear regret*

$$\mathbb{E}[R_T] \leq O(k \log(T)),$$

where k *is defined as:*

$$k \triangleq \max_{(r,s) \in \mathcal{R} \times \mathcal{S}} S^2_{(r,s)} \frac{6 \min(|\mathcal{R}|, |\mathcal{S}|)^3}{\Delta^2_{min}} \Delta_{max},$$
$$\Delta_{max} \triangleq \max_{\boldsymbol{a} \in \mathcal{A} \setminus \{\boldsymbol{a}^*\}} W(\boldsymbol{a}^*) - W(\boldsymbol{a}),$$
$$\Delta_{min} = \min_{\boldsymbol{a} \in \mathcal{A} \setminus \{\boldsymbol{a}^*\}} W(\boldsymbol{a}^*) - W(\boldsymbol{a}).$$

Remark: Theorem 2 states that the regret of protocol `OnlineAssign` is sub-linear to T, which implies the average single-round regret decreases with T and the average single-round welfare increase asymptotically approaches the optimal $W(\boldsymbol{a}^*)$. The technical proof of Theorem 2 is presented in our technical report [8].

IV. Platform-centric Welfare Increase Maximization

In this section, we consider the problem of platform-centric welfare increase maximization where unit valuations and unit costs are private information and not accessible to the platform. We first study a setting where the CTR vector $\boldsymbol{\theta}$ is known to the platform, then we generalize it to the setting where $\boldsymbol{\theta}$ is unknown.

A. Truthful Offline Protocol

We first consider the setting that the CTR vector $\boldsymbol{\theta}$ is known, but the requesters' unit valuations and suppliers' unit costs are unknown to the platform. The protocol for this setting serves as a building block for the setting where $\boldsymbol{\theta}$ is also unknown. The core idea is to design a protocol to solicit true unit valuations from requesters and true unit costs from suppliers, and then apply Protocol 1 to find the optimal assignment profile to optimize welfare increase. Note that such protocol design is independent of time step t. Thus, we omit the subscript or superscript of t for ease of presentation.

Protocol design. Before assigning suppliers to requesters, the platform asks each requester $r \in \mathcal{R}$ to report their unit valuation b_r, and asks each supplier $s \in \mathcal{S}$ to report their unit cost b_s. We define the reported valuation profile of a subset of requesters $\mathcal{X} \subseteq \mathcal{R}$ as

$$\boldsymbol{b}|_{\mathcal{X}}=(b_u : u \in \mathcal{X}),$$

and the reported cost profile of a subset of suppliers $\mathcal{Y} \subseteq \mathcal{S}$ as

$$\boldsymbol{b}|_{\mathcal{Y}}=(b_u : u \in \mathcal{Y}).$$

We then use `OptAssign` (Protocol 1) as an oracle to get the optimal assignment profile. Note that here we use the reported unit valuations $\boldsymbol{b}|_{\mathcal{R}}$ and reported unit costs $\boldsymbol{b}|_{\mathcal{S}}$ as the input unit valuations and unit costs of `OptAssign`. One can easily see that if requesters and suppliers report truthfully, i.e., $\boldsymbol{b}|_{\mathcal{R}} = \boldsymbol{v}$ and $\boldsymbol{b}|_{\mathcal{S}} = \boldsymbol{c}$, then the output assignment profile is exactly the optimal assignment profile. However, the challenge lies in designing a protocol to stimulate requesters and suppliers to report truthfully. Before diving into the protocol design, we define the following optimal assignment function to assist our presentation.

Definition 4 (Optimal assignment function and pseudo welfare increase)**.** *Suppose only a subset of requesters $\mathcal{X} \subseteq \mathcal{R}$ and a subset of suppliers $\mathcal{Y} \subseteq \mathcal{S}$ are available. Define the corresponding optimal assignment function as $\boldsymbol{A}^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}) \triangleq (A_u^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}) : u \in \mathcal{X})$, where*

$$\boldsymbol{A}^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}) = \texttt{OptAssign}(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}, \boldsymbol{b}|_{\mathcal{X}}, \boldsymbol{b}|_{\mathcal{Y}}).$$

Furthermore, we define the corresponding pseudo welfare increase with respect to $\mathcal{X}$ and $\mathcal{Y}$ as

$$W^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}) = \sum_{u \in \mathcal{X}} S_{(u, A_u^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}))} \theta_u (b_u - b_{A_u^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta})}).$$

Namely, $\boldsymbol{A}^*(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta})$ is the optimal assignment profile when only a subset of requesters $\mathcal{X} \subseteq \mathcal{R}$ and a subset of suppliers $\mathcal{Y} \subseteq \mathcal{S}$ are available and their unit valuations and unit costs are assumed to be $\boldsymbol{b}|_{\mathcal{X}}$ and $\boldsymbol{b}|_{\mathcal{Y}}$.

Based on the Vickrey-Clarke-Groves (VCG) mechanism [10]–[12], we design the charging scheme and reward scheme where each requester (resp., supplier) is charged (resp., rewarded) for the *externality*, which is the difference between the welfare increase in the absence of him and the welfare increase in the presence of him. In the presence of requester $r \in \mathcal{R}$, the expected welfare increase of other requesters $\mathcal{R} \setminus \{r\}$ and suppliers $\mathcal{S}$ can be calculated as

$$\underbrace{W^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})}_{\text{welfare increase of } \mathcal{S} \text{ and } \mathcal{R}} - \underbrace{S_{(r, A_r^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}))} \theta_r b_r}_{\text{valuation increase of } r},$$

In the absence of requester $r \in \mathcal{R}$, the expected welfare increase of other requesters $\mathcal{R} \setminus \{r\}$ and suppliers $\mathcal{S}$ can be calculated as $W^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta})$. To stimulate the requester $r \in \mathcal{R}$ to report his true unit valuations, we design a charging scheme where the platform charges r by the marginal deduction on the welfare increase of other requesters $\mathcal{R} \setminus \{r\}$ and suppliers $\mathcal{S}$, formally,

$$\begin{aligned} p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}) = {} & \underbrace{W^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta})}_{\text{welfare increase in the absence of } r} \\ & - \underbrace{(W^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}) - S_{(r, A_r^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}))} \theta_r b_r)}_{\text{welfare increase in the presence of } r}. \end{aligned} \tag{4}$$

The above can be interpreted as the loss of other participating users' welfare increase due to the existence of r. Note that this charging scheme $p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta})$ applies to all requesters $r \in \mathcal{R}$.

Similarly, we calculate the reward for suppliers as follows. In the presence of the supplier $s \in \mathcal{S}$, the welfare increase to requesters $\mathcal{R}$ and other suppliers $\mathcal{S} \setminus \{s\}$ can be calculated as

$$\underbrace{W^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})}_{\text{welfare increase of } \mathcal{R} \text{ and } \mathcal{S}} + \underbrace{S_{(A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}), s)} \theta_{A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})} b_s}_{\text{cost of supplier } s}.$$

where $A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})$ denotes the requester who is assigned to supplier s under the assignment profile $\boldsymbol{A}^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})$, and we set $A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}) = 0$ by default when s is not assigned to any requester. In the absence of supplier $s \in \mathcal{S}$, the welfare increase to requesters $\mathcal{R}$ and other suppliers $\mathcal{S} \setminus \{s\}$ can be calculated as $W^*(\mathcal{R}, \mathcal{S} \setminus \{s\}, \boldsymbol{\theta})$. To stimulate the supplier $s \in \mathcal{S}$ to report his true unit costs, we design a reward scheme where the platform rewards s by the marginal contribution to the welfare increase of requesters $\mathcal{R}$ and other suppliers $\mathcal{S} \setminus \{s\}$, formally,

$$\begin{aligned} & q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}) \\ & = \underbrace{W^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}) + S_{(A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}), s)} \theta_{A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})} b_s}_{\text{welfare increase in the presence of } s} \\ & \quad - \underbrace{W^*(\mathcal{R}, \mathcal{S} \setminus \{s\}, \boldsymbol{\theta})}_{\text{welfare increase in the absence of } s}. \end{aligned} \tag{5}$$

The above can be interpreted as the gain of other participating users' welfare increase due to the existence of s. Note that this reward scheme $q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta})$ applies to all suppliers $s \in \mathcal{S}$.

Protocol analysis. We analyze the properties of our protocol specified by the charging scheme $p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}), \forall r \in \mathcal{R}$ and the reward scheme $q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}), \forall s \in \mathcal{S}$. We first define the utility for requesters and suppliers.

Definition 5 (Utility)**.** *Given the reported profile $(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}})$, the CTR vector $\boldsymbol{\theta}$, and the assignment profile $\boldsymbol{A}^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta})$ (i.e., the result of $\texttt{OptAssign}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}, \boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}})$), then the utility of a requester $r \in \mathcal{R}$ is defined as:*

$$U_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta}) \triangleq v_r N_{(r, A_r^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}))} - p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}),$$

and the utility of a supplier $s \in \mathcal{S}$ is also defined as his marginal gain:

$$U_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta}) \triangleq q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta}) - c_s N_{(A_s^{*-1}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}), s)},$$

where $N_{(r,s)}$ is the increased number of clicks of requester r contributed by s's reposting, i.e., the number of clicks from users in $\mathcal{N}_s \setminus \mathcal{N}_r$.

In the following, we introduce some conceptions of mechanism design from the perspective of reposting service.

Definition 6 (Efficient)**.** *A protocol is efficient if the selected assignment profile $\boldsymbol{a}$ maximizes the welfare increase, i.e., $\boldsymbol{a} \in \arg\max_{\boldsymbol{a}\in\mathcal{A}} W(\boldsymbol{a})$.*

The efficient property states that by using this protocol, the assignment profile achieves the maximum welfare increase.

Definition 7 (Dominant strategy incentive compatible (DSIC) [13])**.** *A protocol is DSIC (or truthful) if it satisfies the following conditions. For $\forall r \in \mathcal{R}, \forall b_r, \forall \boldsymbol{b}|_{\mathcal{R}\setminus\{r\}}, \forall \boldsymbol{b}|_{\mathcal{S}}$ we have*

$$\mathbb{E}[U_r((v_r, \boldsymbol{b}|_{\mathcal{R}\setminus\{r\}}), \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta})] \geq \mathbb{E}[U_r((b_r, \boldsymbol{b}|_{\mathcal{R}\setminus\{r\}}), \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta})],$$

and for $\forall s \in \mathcal{S}, \forall b_s, \forall \boldsymbol{b}|_{\mathcal{R}}, \forall \boldsymbol{b}|_{\mathcal{S}\setminus\{s\}}$ we have

$$\mathbb{E}[U_s(\boldsymbol{b}|_{\mathcal{R}}, (c_s, \boldsymbol{b}|_{\mathcal{S}\setminus\{s\}}); \boldsymbol{\theta})] \geq \mathbb{E}[U_s(\boldsymbol{b}|_{\mathcal{R}}, (b_s, \boldsymbol{b}|_{\mathcal{S}\setminus\{s\}}); \boldsymbol{\theta})].$$

The DISC property states that reporting unit valuation (resp., cost) truthfully is a weakly-dominant strategy for each requester (resp., supplier). Given that users are all rational, DSIC implies that all requesters and suppliers will report truthfully.

Definition 8 (Ex-interim individually rational (EIIR))**.** *A protocol is EIIR if for $\forall r \in \mathcal{R}$ we have*

$$\mathbb{E}[U_r((v_r, \boldsymbol{b}|_{\mathcal{R}\setminus\{r\}}), \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta})] \geq 0,$$

and for $\forall s \in \mathcal{S}$ we have

$$\mathbb{E}[U_s(\boldsymbol{b}|_{\mathcal{R}}, (c_s, \boldsymbol{b}|_{\mathcal{S}\setminus\{s\}}); \boldsymbol{\theta})] \geq 0.$$

The EIIR property states that participating in and reporting unit valuation/cost truthfully will not lead to a negative expected utility. This guarantees that requesters and suppliers have the incentive to participate in the reposting service.

Theorem 3. *The proposed protocol is efficient, DSIC and EIIR.*

Remark: Theorem 3 states that our proposed protocol can guarantee that requesters and suppliers would achieve the optimal and non-negative utilities by reporting their unit valuations/costs truthfully. As a result, the protocol can guarantee the maximum welfare increase even without given unit valuations and unit costs at first. We postpone the proof to Section VII.

B. Truthful Online Protocol

Now we study the most challenging setting where the CTRs $\boldsymbol{\theta}$, requesters' unit valuations $\boldsymbol{v}$ and suppliers' unit costs $\boldsymbol{c}$ are all unknown to the platform.

Protocol design. Since $\boldsymbol{\theta}$ is unknown, the charging scheme $p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta})$ and the reward scheme $q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \boldsymbol{\theta})$ which depend on $\boldsymbol{\theta}$ are also unable to calculate. To address this challenge, we propose an "explore-then-commit" online protocol `StrategicOnlineAssign`, which is outlined in Protocol 4. This protocol has two phases: the exploration phase and the commit phase. The exploration phase runs for $\max\{\gamma, \frac{|\mathcal{R}|\gamma}{|\mathcal{S}|}\}$ rounds so that each requester is selected for at least γ rounds. Exploration rounds do not involve charges and rewards. After the exploration phase, we use $\bar{\boldsymbol{\theta}}$, the empirical mean of CTR samples, to estimate $\boldsymbol{\theta}$. Then, the protocol goes into the commit phase. In this phase, we use $\bar{\boldsymbol{\theta}}$ to estimate the optimal assignment profile, i.e., $\boldsymbol{A}^*(\mathcal{R}, \mathcal{S}, \bar{\boldsymbol{\theta}})$, and this assignment profile is fixed for the remaining rounds. We also use $\bar{\boldsymbol{\theta}}$ to calculate the charging scheme $p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \bar{\boldsymbol{\theta}})$ and the reward scheme $q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \bar{\boldsymbol{\theta}})$ in each round in the commit phase.

Protocol 4 `StrategicOnlineAssign`$(\mathcal{R}, \mathcal{S})$.

Input: $(\mathcal{N}_u : \forall u \in \mathcal{U})$
Solicit unit valuations $\boldsymbol{b}|_{\mathcal{R}}$ and unit costs $\boldsymbol{b}|_{\mathcal{S}}$
Initialize $N_r = 0, \bar{\theta}_r = 0, \forall r \in \mathcal{R}$
for $t = 1$ to $\max\{\gamma, \frac{|\mathcal{R}|\gamma}{|\mathcal{S}|}\}$ **do**
  Select an assignment $\boldsymbol{a}_t$, s.t. $|\mathcal{M}(\boldsymbol{a}_t)| = \min\{|\mathcal{R}|, |\mathcal{S}|\}$
  for $r \in \mathcal{R}'_t$ **do**
    Observe $N^{(t)}_{(r,a_{t,r})}$
    Update $N_r = N_r + 1$
    Update $\bar{\theta}_r = \left((N_r - 1)\bar{\theta}_r + N^{(t)}_{(r,a_{t,r})}/S_{(r,a_{t,r})}\right)/N_r$
  end for
end for
Denote $\bar{\boldsymbol{\theta}} = (\bar{\theta}_r : r \in \mathcal{R})$
Denote $\hat{\boldsymbol{a}}^* \leftarrow$`OptAssign`$(\mathcal{R}, \mathcal{S}, \bar{\boldsymbol{\theta}}, \boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}})$
for $t = \max\{\gamma, \frac{|\mathcal{R}|\gamma}{|\mathcal{S}|}\} + 1$ to T **do**
  Select assignment $\boldsymbol{a}_t = \boldsymbol{A}^*(\mathcal{R}, \mathcal{S}, \bar{\boldsymbol{\theta}})$
  for $r \in \mathcal{R}'_t$ **do**
    r is charged $p_{t,r} = p_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \bar{\boldsymbol{\theta}})$
    $s = a_{t,r}$ is rewarded $q_{t,s} = q_s(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}, \bar{\boldsymbol{\theta}})$
  end for
end for

We prove that protocol 4 can stimulate requesters/suppliers to report their unit valuations/costs truthfully. We also prove that protocol `StrategicOnlineAssign` also has a sublinear regret. This sublinear regret guides us to select γ. Due to page limit, we present them in our technical report [8].

V. Experiments on Real-world Datasets

In this section, we conduct experiments on two real-world datasets Twitter and Google+ to evaluate the performance of our proposed protocols.

A. Experimental Settings

The real-world datasets we use for evaluation are described as follows.

- **Twitter Social Network [14]:** It is a sub-network of the Twitter user-user following network which contains 23,370 nodes and 33,101 directed edges. A node represents a user and a directed edge indicates that a user follows another user.
- **Google+ Social Network [14]:** It is a sub-network of the Google+ user-user following network which contains

23,628 nodes and 39,242 directed edges. A node represents a user and a directed edge denotes that one user has the other user in his circles.

In the above two datasets, there are no requesters or suppliers. Thus, we sample two subsets of users uniformly at random as requesters and suppliers in each dataset. Specifically, we sample 0.25% users as requesters and 0.25% users as suppliers. Besides, the datasets do not contain CTRs $\boldsymbol{\theta}$. The CTR of each requester is independently sampled from $[0, 1]$ uniformly at random. We synthesize the unit valuation of a requester with degree d as

$$\frac{(1 + d/d_{max})^{\lambda}}{2^{\lambda}} \in [0, 1], \tag{6}$$

and synthesize the unit cost of a supplier with degree d as

$$1 - \frac{(1 + d/d_{max})^{\lambda}}{2^{\lambda}} \in [0, 1], \tag{7}$$

where d_{max} is the largest degree in the network and $\lambda > 0$ is a parameter of the function to control the relationship between degree and unit valuation/cost. Specifically, Eq. (6) models that a requester with a larger degree tends to have a larger unit valuation v_r and Eq. (7) models that a supplier with a smaller degree tends to have a larger unit cost c_s. We use $\lambda = 0.8$ in the experiments unless otherwise specified.

We compare our protocols with other baselines. We use `Optimal` to refer to the optimal result that can be achieved under the setting with full knowledge, i.e., CTRs, unit valuations and unit costs are all accessible. We also have the following three heuristic methods: `DegradedAssignI`, `DegradedAssignII` and `DegradedAssignIII`. They use an framework similar to `OptAssign` to determine the assignment in each round. The only difference is the step in Line 1 of Protocol 1, which is to compute the weight $e_{r,s}$ of each edge (r, s) when constructing the bipartite graph. More specifically, `DegradedAssignI` use $e_{r,s} = S_{(r,s)}(v_r - c_s)$, `DegradedAssignII` use $e_{r,s} = S_{(r,s)}\theta_r$, and `DegradedAssignIII` use $e_{r,s} = S_{(r,s)}$.

B. Evaluate `OptAssign` and `GreedyAssign`

We compare two protocols, i.e., `OptAssign` which uses the Hungarian algorithm to find the maximum weighted matching in Protocol 1, and `GreedyAssign` which uses the greedy algorithm stated in Protocol 2 to find the maximum weighted matching, concerning welfare increase and running time. We evaluate these two protocols on datasets in the setting with full knowledge. From Fig. 1(a), one can observe that for both datasets, the welfare increases achieved by `GreedyAssign` is only slightly less than that achieved by `OptAssign`. Fig. 1(b) shows the running time of `OptAssign` and `GreedyAssign`. From Fig. 1(b), one can observe that `OptAssign` takes much more running time than `GreedyAssign` does. These results show that in the user-centric setting with known CTRs (full knowledge), `GreedyAssign` can reduce the computational complexity significantly with only a slight drop in the welfare increase.

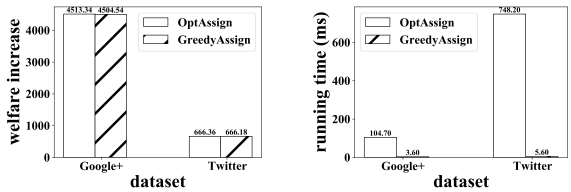

(a) welfare increase (b) running time

Fig. 1: Compare `OptAssign` and `GreedyAssign`.

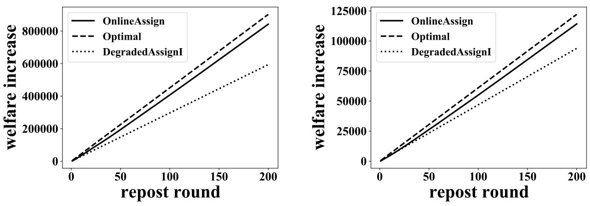

(a) Google+ (b) Twitter

Fig. 2: Compare welfare increase of different methods.

C. Evaluate `OnlineAssign`

We now evaluate `OnlineAssign` (Protocol 3) which is designed for the user-centric aspect with unknown CTRs. We run $T = 200$ rounds. We assume that for requester r, the observed samples of CTR across different rounds are independent and identically distributed Gaussian distribution with mean θ_r and variance 1. For comparison, we also show the cumulative welfare increase of heuristic method `DegradedAssignI`, which can be interpreted as a situation where the CTRs of all the requesters are assumed to be 1 seeing as CTRs are unknown. Fig. 2 shows the cumulative welfare increase of different methods. For both datasets, one can observe that the slope of the `OnlineAssign` curve increases with time and converges to the slope of `Optimal`. This observation verifies the logarithmic growth of regret stated in Theorem 2. In other words, protocol `OnlineAssign` can achieve a near-optimal single-round welfare increase in the later rounds, even though the exact values of CTRs are unknown at the beginning. Besides, in each dataset, after round 20 or so, the curve of `OnlineAssign` is always above `DegradedAssignI` and steeper than `DegradedAssignI`, which shows `OnlineAssign` performs much better than the heuristic method when the time horizon is large.

D. Evaluate Truthful Offline Protocol

We evaluate our truthful protocol proposed in Subsection IV-A for the platform-centric aspect where the platform knows CTRs but does not know the unit valuations and unit costs. Recall that the protocol charges requesters according to Eq. (4) and rewards suppliers according to Eq. (5). We compare users' utilities achieved by truthful reporting and untruthful reporting to verify the DSIC property. For untruthful reporting, the

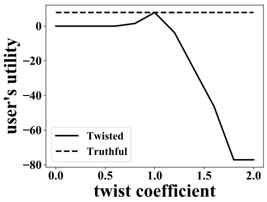

(a) Google+. Requester No.17404

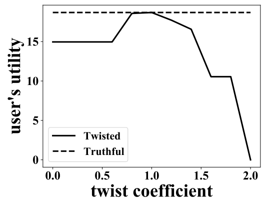

(b) Google+. Supplier No.1276

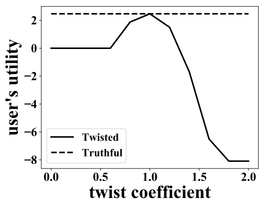

(c) Twitter. Requester No.2809

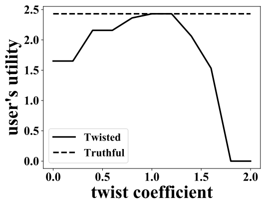

(d) Twitter. Supplier No.12925

Fig. 3: User's utility under different reporting strategies.

reports of requesters and suppliers are twisted as $b_r = \alpha v_r$ and $b_s = \beta c_s$ respectively, where $\alpha, \beta \in \mathbb{R}_+$ are twist coefficients. Specifically, we sample one requester and one supplier uniformly at random from each dataset. The values of twist coefficients α and β are ranged from 0 to 2, with step size 0.2. Fig. 3 shows the utilities of sampled requesters and suppliers when they use different reporting strategies. For both datasets and all the requesters and suppliers, the utilities achieved by untruthful reporting (i.e., twist coefficients are not 1) are lower than or equal to the utilities achieved by truthful reporting. It implies that reporting truthfully is a weakly-dominant strategy which maximizes one's utility.

To show the performance associated with welfare increase, we compare our proposed protocol with heuristic method `DegradedAssignII`, which can be interpreted as a situation where the unit valuations of all the requesters are assumed to be 1 and unit costs of all the suppliers are assumed to be 0 seeing as unit valuations and unit costs CTRs are unknown. Fig. 4 shows the single-round welfare increase achieved by different methods. For both datasets, one can observe that the welfare increases achieved by our proposed protocol are the same as the welfare increases achieved by `OptAssign` (results in Fig. 1(a)). This observation verifies the efficient property of our protocol. Besides, the welfare increases achieved by our proposed protocol are higher than the heuristic method.

E. Evaluate `StrategicOnlineAssign`

We evaluate the protocol `StrategicOnlineAssign` proposed in Subsection IV-B for the platform-centric aspect with unknown CTRs, i.e., the setting where the platform does not know the CTRs, unit valuations and unit costs. We apply protocol `StrategicOnlineAssign` on both datasets for $T = 200$ rounds. We vary γ, the number of rounds in the exploration phase, to study its impact on cumulative welfare

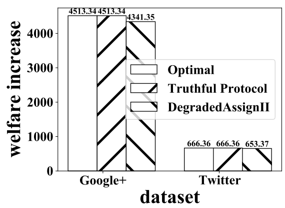

Fig. 4: Compare welfare increase of different methods.

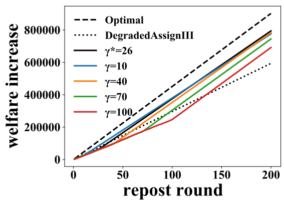

(a) Google+

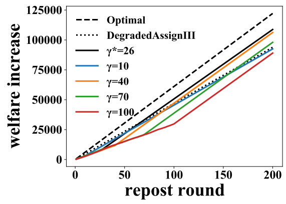

(b) Twitter

Fig. 5: Compare cumulative welfare increase of different methods and different values of γ.

increase. We also compute the value γ^* that minimizes the regret upper bound of Protocol 4 (please refer to our technical report [8] for the regret upper bound) and ensures the tightest upper bound given $T = 200$. For both datasets, the solution are $\gamma^* = 26$. For comparison, we compare it with heuristic method `DegradedAssignIII` which can be interpreted as a situation where the CTRs of all the requesters are assumed to be 1, the unit valuations of all the requesters are assumed to be 1 and unit costs of all the suppliers are assumed to be 0, seeing as all of these are unknown to the platform. Fig. 5 shows the cumulative welfare increases when we use `DegradedAssignIII` and `StrategicOnlineAssign` with different values of γ. We can observe that for both datasets, the protocol using $\gamma = \gamma^*$ achieves a larger cumulative welfare increase in the end than using other γ values. Besides, compared with `DegradedAssignIII`, `StrategicOnlineAssign` ($\gamma = \gamma^*$) achieves a larger cumulative welfare increase in the end in both datasets. We can also observe that, for both datasets, the slope of `StrategicOnlineAssign` ($\gamma = \gamma^*$) in round 200 is steeper than `DegradedAssignIII` and similar to `Optimal`, which implies a near-optimal single-round welfare increase in the later rounds.

VI. Related Work

Application perspective. From an application perspective, there are lots of previous works studying the topic of social network visibility from different aspects. Firstly, some works study visibility from a measurement perspective [15], [16], which revealed insights into the pattern of visibility of different users. Secondly, many works focus on how to boost visibility in OSNs, i.e., to reach a large and diverse population of users, and the methods can be categorized into two classes. The first class of methods resorts to new connections establishment, i.e., selecting appropriate users to connect so that the visibility can be efficiently increased [17]–[19]. The second class is by way of optimizing the broadcasting of contents, e.g., to select the appropriate time for a user to post or share his content [20]–[24] and select users who are influential to broadcast the content (influence maximization problem) [25]–[27]. Our work also falls in the second class. Lastly, there are also some works devoted to users who run businesses in OSNs associated with their visibility. Some of them consider the setting where multiple product sellers compete for making their products get larger visibility [1]–[4]. Besides, some analytical models were proposed to analyze the competition and equilibrium among multiple product sellers [1], [2], [4] and the fairness among product sellers [3]. The main differences between our work and mentioned previous work are as follows: (a) we use a quite different model of visibility which captures factors such as CTRs, unit valuations and unit costs, but in previous works these factors are either not captured or assumed to be known; and (b) the visibility is boosted via reposting behavior, which is different from the strategies used in the above-mentioned works.

Methodology perspective. From a methodology perspective, our work is closely related to combinatorial multi-armed bandit (CMAB), mechanism design and online mechanism design. Firstly, CMAB [28], [29] is a variant of the classic MAB model [30], where multiple arms (a.k.a. a super arm) can be pulled in each round in contrast. In protocol `OnlineAssign`, we customize the framework of CMAB to our setting by taking an assignment profile as a super arm. The second aspect is mechanism design [31]–[33] which aims to incentivize players to act truthfully. Vickrey-Clarke-Groves (VCG) mechanism [10]–[12] is one of the most well-known auction mechanisms. Our truthful protocols in Subsection IV-A and Subsection IV-B are variants of the VCG mechanism. The last aspect, the MAB mechanism, is at the intersection of the above two fields. The traditional MAB model simply assumes that all arms are static choices. However, in many applications [34], [35], the arms can represent rational and selfish individuals. Thus, mechanism design has been applied in MAB context to deal with the interplay between online learning and the strategic players, leading to MAB mechanisms [36]. A typical application of MAB mechanism is sponsored search auction (SSA) [37]. However, the solution to SSA can not be applied to our problem for the following major reasons: (a) SSA simply uses a greedy assignment rule while we use a more complicated matching-based assignment rule; and (b) SSA involved one-side private information while our problem has two-side private information (unit valuations and unit costs).

VII. Selected Proof

A. Proof Sketch of Theorem 3

Proof. For requester $r \in \mathcal{R}$, his utility under an assignment profile $\boldsymbol{a}$ is

$$U_r(\boldsymbol{b}|_{\mathcal{R}}, \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{\theta}) = N_{(r,a_r)} v_r - S_{(r,a_r)} \theta_r b_r$$
$$+ \sum_{r \in \mathcal{R}} S_{(r,a_r)} \theta_r (b_r - b_{a_r}) - W^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta}). \quad (8)$$

One can observe that $W^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta})$ is independent of r. Thus, maximizing r's expected utility is equivalent to expecting the assignment profile to maximize

$$\mathbb{E}\left[N_{(r,a_r)} v_r - S_{(r,a_r)} \theta_r b_r + \sum\nolimits_{r \in \mathcal{R}} S_{(r,a_r)} \theta_r (b_r - b_{a_r})\right]$$
$$= S_{(r,a_r)} \theta_r v_r - S_{(r,a_r)} \theta_r b_{a_r} + \sum\nolimits_{\bar{r} \in \mathcal{R} \setminus \{r\}} S_{(\bar{r},a_{\bar{r}})} \theta_r (b_{\bar{r}} - b_{a_{\bar{r}}}). \quad (9)$$

We define the pseudo welfare increase function of a given assignment profile $\boldsymbol{a}$ for $\forall \mathcal{X} \subseteq \mathcal{R}, \forall \mathcal{Y} \subseteq \mathcal{S}$ as

$$W(\mathcal{X}, \mathcal{Y}, \boldsymbol{\theta}, \boldsymbol{b}|_{\mathcal{X}}, \boldsymbol{b}|_{\mathcal{Y}}; \boldsymbol{a}) = \sum\nolimits_{u \in \mathcal{X}} S_{(u,a_u)} \theta_u (b_u - b_{a_u}).$$

One can observe that the value of Eq. (9) is exactly the value of $W(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}, (v_r, \boldsymbol{b}|_{\mathcal{R} \setminus \{r\}}), \boldsymbol{b}|_{\mathcal{S}}; \boldsymbol{a})$. If requester reports $b_r = v_r$, then the platform would select the assignment profile $\boldsymbol{a}^*$ returned by $\texttt{OptAssign}(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}, (v_r, \boldsymbol{b}|_{\mathcal{R} \setminus \{r\}}), \boldsymbol{b}|_{\mathcal{S}})$, which also maximizes Eq. (9). The analysis of the suppliers is similar. Thus, we prove DSIC.

Then, the protocol is obviously efficient due to DSIC and the optimality of the selected assignment profile.

Finally, we prove EIIR. The expected utility of r is

$$\mathbb{E}[U_r(\boldsymbol{v}, \boldsymbol{c}; \boldsymbol{\theta})] = W^*(\mathcal{R}, \mathcal{S}, \boldsymbol{\theta}) - W^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta}). \quad (10)$$

Since $\boldsymbol{A}^*(\mathcal{R} \setminus \{r\}, \mathcal{S}, \boldsymbol{\theta})$ is also available in the real world, Eq. (10) is non-negative. The analysis of suppliers is similar. Thus we can prove the protocol is EIIR.

VIII. Conclusion

In this paper, we design protocols for the reposting service to incentivize "transactions" between requesters and suppliers. We present a mathematical model for the reposting service and formulate the welfare increase maximization problem from both user-centric and platform-centric aspects. For the user-centric aspect, we proposed an online protocol `OnlineAssign` which uses `OptAssign` as an oracle to determine the assignment profiles sequentially to maximize the cumulative welfare increase with a provable sub-linear regret. For the platform-centric aspect, we design an truthful online protocol `StrategicOnlineAssign` which uses a VCG-based charging and reward scheme and the "explore-then-commit" mode. We prove that the protocol can solicit the true unit valuations/costs and we also prove its regret. Extensive experiments on real-world datasets validate the effectiveness of our proposed protocols.

Acknowledgment

We would like to thank anonymous reviewers and our shepherd Prof. Qiang Liu for their helpful comments that helped us improve this work. The work of John C.S. Lui was supported in part by the RGC's GRF 14200321. The work of Hong Xie was supported in part by Chongqing Talents: Exceptional Young Talents Project (cstc2021ycjhbgzxm0195). Hong Xie is the corresponding author.

REFERENCES

[1] E. Altman, P. Kumar, S. Venkatramanan, and A. Kumar, "Competition over timeline in social networks," in *2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2013)*. IEEE, 2013, pp. 1352–1357.

[2] R. Dhounchak, V. Kavitha, and E. Altman, "A viral timeline branching process to study a social network," in *2017 29th International Teletraffic Congress (ITC 29)*, vol. 3. IEEE, 2017, pp. 19–24.

[3] E. Hargreaves, C. Agosti, D. Menasché, G. Neglia, A. Reiffers-Masson, and E. Altman, "Fairness in online social network timelines: Measurements, models and mechanism design," *Performance Evaluation*, vol. 129, pp. 15–39, 2019.

[4] A. Reiffers-Masson, E. Hargreaves, E. Altman, W. Caarls, and D. S. Menasché, "Timelines are publisher-driven caches: Analyzing and shaping timeline networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, no. 3, pp. 26–29, 2017.

[5] S. Asur, B. A. Huberman, G. Szabo, and C. Wang, "Trends in social media: Persistence and decay," in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 5, no. 1, 2011, pp. 434–437.

[6] A. Jungherr, "Twitter use in election campaigns: A systematic literature review," *Journal of information technology & politics*, vol. 13, no. 1, pp. 72–91, 2016.

[7] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[8] S. Zheng, X. Liu, H. Xie, and J. C. Lui, "Reposting service in online social networks: Modeling and online incentive protocols," https://www.dropbox.com/sh/r3uy26olpwp0ib9/AAADrTsJAKT3k7Dws3cnaOMwa, 2022, online.

[9] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.

[10] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of finance*, vol. 16, no. 1, pp. 8–37, 1961.

[11] E. H. Clarke, "Multipart pricing of public goods," *Public choice*, pp. 17–33, 1971.

[12] T. Groves, "Incentives in teams," *Econometrica: Journal of the Econometric Society*, pp. 617–631, 1973.

[13] R. B. Myerson, "Optimal auction design," *Mathematics of operations research*, vol. 6, no. 1, pp. 58–73, 1981.

[14] J. J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks." in *NIPS*, vol. 2012. Citeseer, 2012, pp. 548–56.

[15] A. Bessi, F. Zollo, M. Del Vicario, M. Puliga, A. Scala, G. Caldarelli, B. Uzzi, and W. Quattrociocchi, "Users polarization on facebook and youtube," *PloS one*, vol. 11, no. 8, p. e0159641, 2016.

[16] T. Bucher, "Want to be on the top? algorithmic power and the threat of invisibility on facebook," *New media & society*, vol. 14, no. 7, pp. 1164–1180, 2012.

[17] K. Gatmiry and M. Gomez-Rodriguez, "The network visibility problem," *ACM Transactions on Information Systems (TOIS)*, vol. 40, no. 2, pp. 1–42, 2021.

[18] S. Zheng, H. Xie, and J. C. Lui, "Social visibility optimization in osns with anonymity guarantees: Modeling, algorithms and applications," in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2021, pp. 2063–2068.

[19] ——, "Pricing social visibility service in online social networks: modeling and algorithms," in *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2021, pp. 209–213.

[20] M. R. Karimi, E. Tavakoli, M. Farajtabar, L. Song, and M. Gomez Rodriguez, "Smart broadcasting: Do you want to be seen?" in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 1635–1644.

[21] N. Spasojevic, Z. Li, A. Rao, and P. Bhattacharyya, "When-to-post on social networks," in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 2127–2136.

[22] U. Upadhyay, A. De, and M. Gomez Rodriguez, "Deep reinforcement learning of marked temporal point processes," *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[23] A. Zarezade, A. De, U. Upadhyay, H. R. Rabiee, and M. Gomez-Rodriguez, "Steering social activity: A stochastic optimal control point of view." *J. Mach. Learn. Res.*, vol. 18, pp. 205–1, 2017.

[24] A. Zarezade, U. Upadhyay, H. R. Rabiee, and M. Gomez-Rodriguez, "Redqueen: An online algorithm for smart broadcasting in social networks," in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 2017, pp. 51–60.

[25] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proc. ACM KDD*, 2003.

[26] S. Aral and P. S. Dhillon, "Social influence maximization under empirical influence models," *Nature human behaviour*, vol. 2, no. 6, p. 375, 2018.

[27] Y. Lin and J. C. Lui, "Analyzing competitive influence maximization problems with partial information: An approximation algorithmic framework," *Performance Evaluation*, vol. 91, pp. 187–204, 2015.

[28] W. Chen, Y. Wang, and Y. Yuan, "Combinatorial multi-armed bandit: General framework and applications," in *International Conference on Machine Learning*. PMLR, 2013, pp. 151–159.

[29] Y. Gai, B. Krishnamachari, and R. Jain, "Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations," *IEEE/ACM Transactions on Networking*, vol. 20, no. 5, pp. 1466–1478, 2012.

[30] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *arXiv preprint arXiv:1204.5721*, 2012.

[31] Y. Narahari, *Game theory and mechanism design*. World Scientific, 2014, vol. 4.

[32] N. Nisan and A. Ronen, "Computationally feasible vcg mechanisms," *Journal of Artificial Intelligence Research*, vol. 29, pp. 19–47, 2007.

[33] T. Roughgarden, "Algorithmic game theory," *Communications of the ACM*, vol. 53, no. 7, pp. 78–86, 2010.

[34] H. Ma, D. Zhao, and P. Yuan, "Opportunities in mobile crowd sensing," *IEEE Communications Magazine*, vol. 52, no. 8, pp. 29–35, 2014.

[35] L. Anderegg and S. Eidenbenz, "Ad hoc-vcg: a truthful and cost-efficient routing algorithm for mobile ad hoc networks with selfish agents," in *Proceedings of the 9th annual international conference on Mobile computing and networking*, 2003, pp. 245–259.

[36] C. Liu, Q. Cai, and Y. Zhang, "Multi-armed bandit mechanism with private histories," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 1607–1609.

[37] M. Babaioff, Y. Sharma, and A. Slivkins, "Characterizing truthful multi-armed bandit mechanisms: Extended abstract," in *Proceedings of the 10th ACM Conference on Electronic Commerce*, 2009, pp. 79–88.