

RESEARCH ARTICLE

Token-Curated Registry with Citation Graph

Kensuke Ito,^{*†} Hideyuki Tanaka[‡]

Abstract. In this study, we aim to incorporate the expertise of anonymous curators into a token-curated registry (TCR), a decentralized recommender system for collecting a list of high-quality content. This registry is important, because previous studies on TCRs have not specifically focused on technical content, such as academic papers and patents, whose effective curation requires expertise in relevant fields. To measure expertise, curation in our model focuses on both the content and its citation relationships, for which curator assignment uses the Personalized PageRank (PPR) algorithm while reward computation uses a multi-task peer-prediction mechanism. Our proposed *CitedTCR* bridges the literature on network-based and token-based recommender systems and contributes to the autonomous development of an evolving citation graph for high-quality content. Moreover, we experimentally confirm the incentive for registration and curation in *CitedTCR* using the simplification of a one-to-one correspondence between users and content (nodes).

1. Introduction

For many blockchain-based decentralized applications (DApps), one of the challenges is the reliability of information originating from an *off-chain* environment. This is because the Bitcoin protocol,¹ which is the origin of DApps and has a novelty of building a reliable consensus among anonymous users (on a public peer-to-peer network), only computes information generated from an *on-chain* environment (*i.e.*, transaction records of Bitcoin). For example, consider the case of a simple DApp that provides alerts when it rains in a given location. In this case, while the DApp can ensure the on-chain state transition leads to an alert, it cannot ensure the off-chain fact (used as the trigger) that it has actually rained at the location. Therefore, most DApps rely on trusted third parties, such as the National Weather Service, for their input.² This is in contrast to the Bitcoin protocol, which functions even if the operators of each node are unknown. Consequently, DApps require an additional protocol in which anonymous users can build reliable consensus on off-chain information to maintain the novelty of the Bitcoin protocol.

A token-curated registry (TCR) is a DApp for establishing such a protocol.^{3,4} It specializes in compiling a high-quality, reliable list of off-chain content (*e.g.*, restaurants, universities, and webpages) as a recommender system.⁵ Although there are different design patterns among existing TCRs,⁶ generally their consensus building is based on a token-staking scheme in which all users can stake their tokens on a binary choice {accept, reject} as curators whenever an applicant posts new content to the list.⁷ Consensus is the selection that obtains more tokens

* bc1q3w7grsh4jd074q7m5j6nl8dt2ffmaaflvgpmuy

† Kensuke Ito (k-ito@g.ecc.u-tokyo.ac.jp) is Ph.D. student at the University of Tokyo.

‡ Hideyuki Tanaka (tanaka@iii.u-tokyo.ac.jp) is a professor at the University of Tokyo.

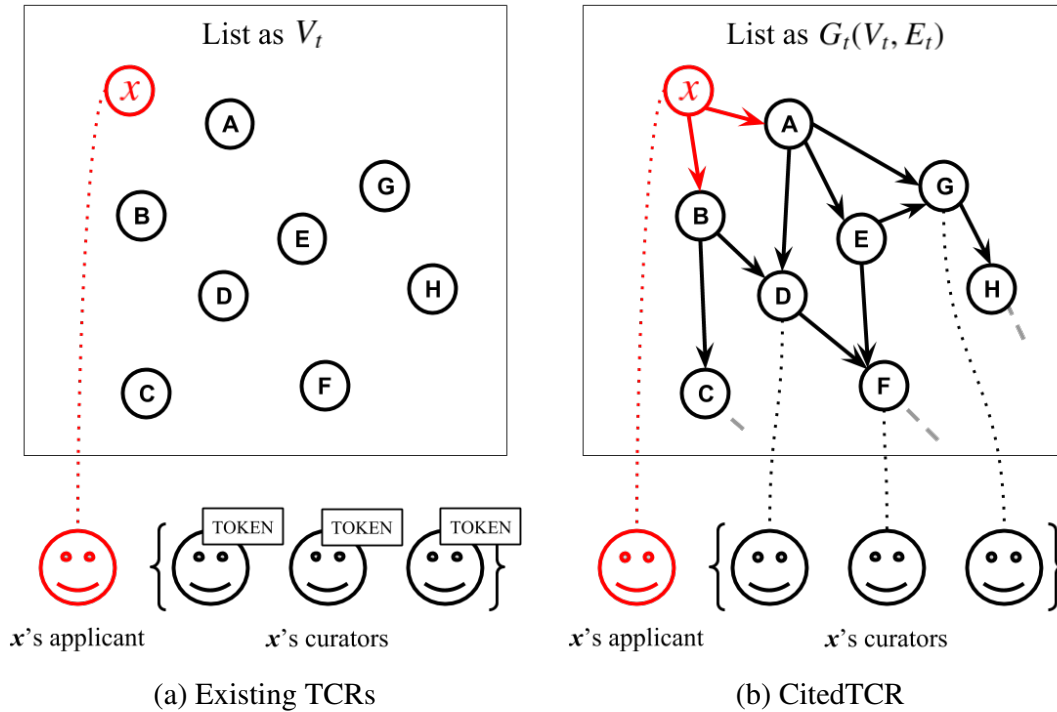


Fig. 1. TCRs in both cases select curators to decide whether to accept a newly proposed content x into the list of off-chain content $\{A, B, \dots\}$. However, while existing TCRs (a) manage an unstructured list V_t that can be curated by any user who stakes certain number of tokens, *CitedTCR* (b) manages an evolving directed acyclic graph (DAG)-structured list $G_t(V_t, E_t)$ whose curators are assigned according to citation relationships.

compared to another selection after a certain period. Moreover, all staked tokens are redistributed among curators who stake their tokens on the consensus side, *i.e.*, token staking intends to yield informative reports from anonymous curators who risk losing their tokens as well as the token price, which is assumed to fluctuate with the quality of the list. One limitation is that token staking does not reflect expertise in consensus building because, regardless of specialty, any user with a certain number of tokens can participate in the curation. Therefore, the reliability of consensus is restricted under TCRs, which primarily depend only on token staking, particularly when the off-chain content is technical (*e.g.*, academic papers and patents) and requires expertise in specific fields for effective curation.

Accordingly, in this study, we aim to incorporate the expertise of anonymous curators into TCRs using a protocol called *CitedTCR*, which leverages a citation graph for curator assignment and uses a *peer-prediction* mechanism to compute the number of reward tokens paid to the curators. Fig.1 illustrates the role of the citation graph in our protocol. Fig.1(a) shows that existing TCRs manage an evolving unstructured list (as a set) V_t , in which an applicant posts new content (as an element) x and any user can be the curator of x because of token staking. However, Fig. 1(b) shows that *CitedTCR* manages an evolving list $G_t(V_t, E_t)$ with a citation graph (*i.e.*, a DAG) structure, in which an applicant posts x , and its out-edges $(x, A), (x, B)$ point to existing nodes $\{A, B\}$ as references. Moreover, curators are stochastically assigned to a given number of users who have posted nodes (*e.g.*, $\{D, F, G\}$) that have both high similarity with x 's reference nodes $\{A, B\}$ and high centrality in G_t .⁸ *CitedTCR* assigns appropriate curators in a manner

similar to the academic peer-review process, in which researchers who have produced high-quality papers with a large number of citations are more likely to be selected as reviewers in their field of expertise. Note that this form of curator assignment serves as an incentive for applicants to register high-quality content in *CitedTCR* because users may have more opportunities to obtain reward tokens as curators if their content in G_t attracts a large number of citations.^{9,10} The citation graph serves as a proxy for the expertise of anonymous curators; therefore, the reliability of G_t from the perspective of both curation and registration is ensured.

Peer prediction is a mechanism of game theory for eliciting informative reports for tasks with no ground truth, such as the peer review of academic papers and online product reviews by consumers. In particular, peer prediction compares user reports for the same task to create a *truthful* (known as *strategy-proof* or *incentive-compatible*) environment, in which no user can obtain a higher utility by any possible strategy deviating from the user's true preferences.¹¹ *CitedTCR* uses peer prediction for reward computation, in which it is assumed that the assigned curators can obtain newly issued reward tokens if they return a binary signal {accept, reject} as a report for x and x 's citation relationships.¹² This mechanism addresses two problems in the token-staking scheme, which is even more critical under *CitedTCR*. The first problem is the risk of strategic misreports (such as collusion) among curators. Although this has been discussed for existing TCRs,^{13,14} token staking becomes more vulnerable to this risk in *CitedTCR* because *CitedTCR* assigns a fixed number of homogeneous curators with similar expertise. The second problem is the lack of incentive to participate in consensus building because of the risk of losing staked tokens (see Appendix A). Note that strengthening the weak incentive of token staking is a common topic in TCRs.¹⁵ Stronger incentives are particularly important for *CitedTCR* in which reports elicited from assigned curators are the key for reflecting expertise in G_t . Therefore, rather than token staking, we use a peer-prediction mechanism that provides maximum (new) rewards for informative reports.

CitedTCR is thus a hybrid of token-based and network-based recommender systems because it recommends both $V_t(G_t)$ curated by tokens and curators assigned according to G_t . In this study, as a first step of this hybrid approach, we used the Personalized PageRank (PPR) algorithm for curator assignment,¹⁶ and a peer prediction mechanism called DG13 proposed by Dasgupta and Ghosh for reward computation.¹⁷ In addition to their popularity, both PPR and DG13 have several favorable properties for *CitedTCR* as demonstrated in Sections 2 and 3. Moreover, we assume that users in this study have a one-to-one correspondence with $V_t(G_t)$. This assumption is intended to simplify the curation process into a state transition in G_t ; its details are discussed in Section 3.

The remainder of this paper is organized as follows. In Section 2, we introduce related studies and contributions from the perspective of three components: TCR, the PPR algorithm, and a peer-prediction mechanism. In Section 3, we describe the specification of *CitedTCR*, including the role of PPR and DG13. In Section 4, we examine the practical utility of our proposal using two step-wise simulations with the citation graph of academic papers. Finally, in Section 5, we conclude the paper with a summary of achievements and remaining concerns.

2. Related Work

2.1. *TCR*—Since Goldin proposed the initial design in 2017,^{3,4} TCRs have been implemented in a number of applications such as the adChain registry for webpages,¹⁸ the Ocean protocol for user reputations,¹⁹ and the Civil registry for news articles.²⁰ Because TCR is a recent development, most discussion at present focus on blog articles whose topics vary from the classification of design patterns to critical examinations of token staking.^{6,14,21} A reading list curated by the blockchain community would be helpful for summarizing this discussion.^{22,23} In addition to blog articles, TCRs have been examined in academic papers, primarily from a game-theoretic perspective. For example, Asgaonkar and Krishnamachari presented a mathematical foundation of the TCR 1.1 model to determine the sufficient conditions for each consensus at equilibrium.^{4,24} Wang and Krishnamachari introduced enhanced token staking with a new issuance of reward tokens to create an incentive to participate in consensus building.¹⁵ Moreover, Falk and Tsoukalas used an axiomatic approach to demonstrate the limitations of a token-staking scheme, in which the expected rewards are proportional to the amount of staking.¹³

As mentioned in Section 1, in this study, we aim to incorporate the expertise of anonymous curators into TCRs using a combination of citation graphs and peer prediction (*i.e.*, PPR and DG13). This approach is novel because previous studies and blog articles on TCRs have not explicitly addressed the mechanism for technical content, such as academic papers and patents, whose effective curation requires expertise in relevant fields.

2.2. *The PPR algorithm*—The PPR algorithm,¹⁶ originally named topic-sensitive PageRank, is an extension of the PageRank algorithm and computes a score of importance for each node from the viewpoint of the entire network structure.^{25,26} While the PageRank score originates from a random walk on the network, PPR allows this random walk to return to the predetermined set of nodes with a given probability,²⁷ thereby adapting the score to recommender systems (see Section 3.2 for details). In many recommender systems using PPR, *CitedTCR* is most closely related to the PaperRank algorithm proposed by Gori and Pucci,²⁸ which applies PPR to a citation graph of academic papers to generate useful paper-to-paper recommendations. Moreover, PPR is a component of several paper-to-reviewer assignment systems that attempt to recommend appropriate peer reviewers for a submitted paper.^{29,30}

From the perspective of PPR, this study provides contributions such as *CitedTCR* bridging the literature on network-based and token-based recommender systems for the first time to strengthen the reliability of the consensus. New economy movement (NEM) is a representative precedent of blockchain-based protocols that leverage a network structure for consensus building.³¹ However, NEM is not specific to TCRs and manages on-chain transaction records using a network-based score different from that of PPR.

2.3. *Peer-prediction mechanism*—Peer prediction was first introduced by Miller *et al.* as an application of the proper scoring rule and game theory.^{32–34} To model the problem of eliciting private information, reward (score) computation assumes an environment in which each user reports probabilistic but correlated signals based on the assigned tasks. As examined by Jurca and Faltings,³⁵ a common problem in the mechanism proposed by Miller *et al.* and subsequent mechanisms is that the computation has multiple Nash equilibria, including uninformative ones in which elicited reports are independent of the true signals³⁶—*e.g.*, the same signals or random signals are always reported to avoid the effort of observation. As a solution to this problem,

Dasgupta and Ghosh proposed a multi-task peer-prediction mechanism called DG13 that assigns multiple tasks to one user and computes rewards for one task using the reports produced for other tasks.¹⁷ Under the assumption of positively correlated binary signals, DG13 ensures strong truthfulness,³⁷ in which an equilibrium by informative reports has the highest rewards among other realistic equilibria (see Section 3.3 for details). *CitedTCR* uses DG13 because the abovementioned properties of multi-tasking, strong truthfulness, and binary signals are compatible with the general settings of TCRs, in which curators evaluate multiple content using binary choices.

To our knowledge, *CitedTCR* is the first proposal that uses a peer-prediction mechanism in TCRs. This proposal presents an approach that can overcome the aforementioned two problems in the token staking. In addition to DG13, recent studies on peer prediction have discussed topics relevant to TCRs. For example, Agarwal *et al.* proposed a multi-task mechanism that assigns appropriate tasks to heterogeneous users (with various propensities) based on accumulated reports.³⁸ This can contribute to TCRs with expertise as an approach different from citation graphs. Goel *et al.* assessed the robustness of a peer-prediction mechanism for the case in which an incentive for misreporting exists outside the system with an application to decentralized oracles.^{39,40} Their assessment can be applicable to TCRs with a design similar to that of decentralized oracles.

3. Model

In this section, we describe the specification of *CitedTCR* as a state transition closed on list G_t . This simplification, achieved by several assumptions, including the aforementioned one-to-one correspondence, is useful for an algorithmic expression and for the experimental simulations described in Section 4. Moreover, we present details of PPR and DG13 that clarify how these components contribute to curation in *CitedTCR*.

3.1. Setup—As depicted in Fig. 1(b), our protocol deals with an evolving DAG-structured list $G_t(V_t, E_t)$, where V_t denotes the set of registered content and $E_t \subseteq V_t \times V_t$ denotes their citation relationships. Although G_t is managed by a set of users U_t (as with other DApps), we impose the following assumption on the management of G_t .

Assumption 1 One-to-one correspondence: suppose that there is a one-to-one correspondence between U_t and V_t , *i.e.*, $f : U_t \rightarrow V_t$ is bijective.

A one-to-one correspondence indicates an environment in which a user can neither post more than one content nor share one content as a co-applicant. This setting frees our model from several complex problems in DApps, such as spamming and sybil attacks, and makes curator assignment equivalent to node selection in G_t .

We further assume that only one node x proposes an additional citation graph \dot{G}_t (composed of x and the references of x) in each period, and \dot{G}_t is not delisted once it is accepted into G_t . This assumption and one-to-one correspondence make it possible to represent *CitedTCR* as a state transition $\{G_t\}_{t=0}^{\infty}$ that repeatedly determines whether to accept \dot{G}_t in each period.⁴¹ In particular, the transition from G_t to G_{t+1} has the following steps, where we summarize the meanings of each notation in Table 1 for the sake of readability:

Table 1. Notations for the state transition in *CitedTCR*

Notations	Meanings
$G_t(V_t, E_t)$	Citation graph (as an evolving DAG-structured list) in period t .
V_t	Set of nodes (as registered content) in period t .
E_t	Set of edges (as citation relationships) for V_t .
$\dot{G}_t(\{x\} \cup V_x, E_x)$	An additional citation graph proposed in period t .
x	A new node in period t .
V_x	Set of base nodes (as references) for x (<i>i.e.</i> , $V_x \subseteq V_t$).
E_x	Set of edges directed from $\{x\}$ to V_x .
\dot{C}_t	Set of curators $\{1, 2, \dots, n\}$ selected from $V_t \setminus V_x$ in period t .
\dot{R}_t	Set of reports $\{r_1^{\dot{G}_t}, r_2^{\dot{G}_t}, \dots, r_n^{\dot{G}_t}\}$ elicited from \dot{C}_t .
$r_c^{\dot{G}_t}$	A report ($\in \{0, 1\}$) that a curator $c \in \dot{C}_t$ did for \dot{G}_t .
$\dot{\Theta}_t$	Set of rewards $\{\theta_1^{\dot{G}_t}, \theta_2^{\dot{G}_t}, \dots, \theta_n^{\dot{G}_t}\}$ for \dot{C}_t .
$\theta_c^{\dot{G}_t}$	A reward for a curator $c \in \dot{C}_t$ who provided $r_c^{\dot{G}_t}$.
n	An exogenous parameter (≥ 2) for setting $ \dot{C}_t $. ⁴²
m	An exogenous parameter ($\leq n$) for the difficulty of registration.

Step 1: A new node x proposes \dot{G}_t to G_t .

Step 2: Select n of nodes as \dot{C}_t (curator assignment).

Step 3: Collect n of reports as \dot{R}_t , where $\dot{R}_t \ni r_c^{\dot{G}_t} = 0$ and 1 each designate *reject* and *accept*.

Step 4: Compute $\dot{\Theta}_t$ (reward computation).

Step 5: Update G_t to G_{t+1} , where G_{t+1} includes \dot{G}_t only if \dot{R}_t has m or more number of $r = 1$.

A pseudocode can be used to convert this state transition into Algorithms 1 and 2, in which, as commented, curator assignment (step 2) uses PPR, and reward computation (step 4) uses DG13. These algorithms include the following two properties. First, they integrate steps 2 and 3 as the $\text{CURATION}(N, C, R, G)$ function (Algorithm 2), which returns a set of reports R for the following four arguments: N , the number of reports; C , the set of nodes that are candidates for the curator; R , the initial value of the set of reports; and G , the graph containing C . This integration is intended to handle a case in which assigned curators do not provide their reports within a given period of time. In this case, $\text{CURATION}(N, C, R, G)$ continues to reselect new nodes as replacements for unresponsive curators until it collects N of reports. Second, they return not only G_{t+1} and $\dot{\Theta}_t$ but also the stock of reports R_{t+1} . This property is specific to DG13, whose reward computation leverages both the flow \dot{R}_t and stock R_t of elicited reports as one of the multi-task peer prediction mechanisms. Algorithm 1 could be simplified by adopting other intratemporal mechanisms such as token staking.

Algorithm 1 State transition in *CitedTCR*

```

1:  $G_t(V_t, E_t) \leftarrow$  list in period  $t$ 
2:  $\dot{G}_t(\{x\} \cup V_x, E_x) \leftarrow$  proposal by  $x$ 
3:  $\{m, n\} \leftarrow$  exogenous parameters
4:  $R_t \leftarrow$  stock of the reports until period  $t$  ▷ Peculiar to DG13
5:  $\dot{R}_t \leftarrow$  CURATION( $n, V_t \setminus V_x, \{\emptyset\}, G_t$ ) ▷ See Algorithm 2
6: Compute rewards  $\dot{\Theta}_t$  with  $R_t$  and  $\dot{R}_t$  ▷ Use DG13
7: return  $\dot{\Theta}_t$ 
8:  $R_{t+1} \leftarrow R_t \cup \dot{R}_t$  ▷ Peculiar to DG13
9: return  $R_{t+1}$  ▷ Peculiar to DG13
10: if  $m \geq |\{r \in \dot{R}_t | r = 1\}|$  then
11:    $G_{t+1} \leftarrow G_t$ 
12: else
13:    $G_{t+1} \leftarrow G_t \cup G_x$ 
14: end if
15: return  $G_{t+1}$ 

```

Algorithm 2 Curator assignment and report collection in *CitedTCR*

```

1: function CURATION( $N, C, R, G$ )
2:    $C' \leftarrow N$  number of curators selected from  $C$  in  $G$  ▷ Use PPR
3:    $R' \leftarrow$  reports collected from  $C'$  within a given period of time
4:    $R \leftarrow R \cup R'$ 
5:   if  $|R'| = N$  then
6:     return  $R$ 
7:   else
8:      $N \leftarrow N - |R'|$ 
9:      $C \leftarrow C \setminus C'$ 
10:    CURATION( $N, C, R, G$ )
11:   end if
12: end function

```

3.2. *PPR for curator assignment*—PPR is an algorithm that recommends relatively important nodes to a given node through iterative random walking on a network (as a Markov chain). *CitedTCR* uses PPR for a curator assignment that selects \dot{C}_t as important nodes for x . In the example presented in Fig. 1(b), the set of curators $\dot{C}_t = \{D, F, G\}$ is selected from $\{A, B, C, \dots\} \setminus \{A, B\}$ for the assessment of $\dot{G}_t(\{x, A, B\}, \{(x, A), (x, B)\})$, where \dot{C}_t is regarded as important nodes from the standpoint of x with the reference $V_x = \{A, B\}$. Nodes such as V_x are often referred to as *base nodes* in the PPR context, and are the key for computing relative importance.

To quantify the process of random walking, PPR leverages a transition matrix \mathbf{P} , which in our case is $|V_t| \times |V_t|$, and an element p_{ij} designates the probability of transition from node i to node j . In the random walk as a Markov chain, the value of p_{ij} becomes the reciprocal of node i 's out-degree.

The simplified PageRank score of V_t is the dominant eigenvector (for eigenvalue 1) of \mathbf{P} , which indicates the steady-state probability distribution as a result of iterative random walking.⁴³ Moreover, the PPR score of V_t is the dominant eigenvector of \mathbf{P}_{PPR} , which has the following modification to \mathbf{P} :¹⁶

$$\mathbf{P}_{PPR} = (1 - \alpha)\mathbf{P} + \alpha \frac{1}{|V_x|} \mathbf{B},$$

where \mathbf{B} is an additional $|V_t| \times |V_t|$ matrix whose element b_{ij} becomes 1 if j is included in base nodes V_x ; otherwise, it becomes 0 (i.e., b_{iA} and b_{iB} become 1 for all i and other elements become 0 in Fig. 1(b)). We can interpret $\mathbf{B}/|V_x|$ as another transition matrix in which all nodes in V_t must jump to one node selected from V_x uniformly at random. Thus, \mathbf{P}_{PPR} is the linear combination of the two transition matrices \mathbf{P} and $\mathbf{B}/|V_x|$ that represents biased random walking, which jumps to one of the base nodes with probability α in each step. Here, $\alpha \in [0, 1]$ is called a damping factor, and it can adjust the strength of bias as an exogenous parameter ($\alpha = 0.15$ in most cases). *CitedTCR* stochastically selects n curators in each period according to the PPR computed from \mathbf{P}_{PPR} .

Below, we discuss three properties in this application of PPR. First, similar to PaperRank,²⁸ *CitedTCR* considers G_t to be undirected when using PPR. This is important because if PPR were on a DAG structure, its score would focus on the nodes with no out-edges (i.e., the oldest content in the case of citation graph) and would thus be unreliable for recommender systems. Second, as already mentioned, *CitedTCR* excludes V_x from the candidates of \hat{C}_t . Although PPR scores high for the base nodes (reference nodes for x), we do not select them to avoid biased curation, in which assigned curators accept x simply to increase their number of citations.⁴⁴ Third, *CitedTCR* can encourage users to register high-quality content in G_t , even though the frequency with which they become curators is weighted by PPR. This is experimentally confirmed in Section 4 using the PageRank score in G_t as a proxy for quality.

3.3. DG13 for reward computation—As the representation of DG13 needs several more notations, we here summarize their meanings again in Table 2, along with the explanations below. DG13 and other peer-prediction mechanisms aim to elicit truthful information from the environment, in which users report the quality of a task. For example, in *CitedTCR*, n assigned curators $\hat{C}_t = \{1, 2, \dots, n\}$ provide reports $\hat{R}_t = \{r_1^{\hat{G}_t}, r_2^{\hat{G}_t}, \dots, r_n^{\hat{G}_t}\}$ on the quality of \hat{G}_t . To confirm whether a report is truthful, peer prediction assumes the stochastic signal s , which any $c \in \hat{C}_t$ can observe from \hat{G}_t and can use as input information for $r_c^{\hat{G}_t}$. DG13 focuses on binary signals $s \in \{0, 1\}$ and binary reports $r(s) \in \{0, 1\}$ (0: reject; 1: accept). We use notation $s_c^{\hat{G}_t}$ in the same manner as in reporting, i.e., curator c accepts adding \hat{G}_t to the G_t if $r_c^{\hat{G}_t}(s_c^{\hat{G}_t}) = 1$ and rejects it if $r_c^{\hat{G}_t}(s_c^{\hat{G}_t}) = 0$. This report is truthful in the $r_c^{\hat{G}_t}(0) = 0$ or $r_c^{\hat{G}_t}(1) = 1$ case and non-truthful in the $r_c^{\hat{G}_t}(0) = 1$ or $r_c^{\hat{G}_t}(1) = 0$ case. Note that $r_c^{\hat{G}_t}$ and $s_c^{\hat{G}_t}$ are sometimes denoted r_c and s_c when their task does not need to be emphasized.

We add two more assumptions that are common in the literature on peer prediction for binary signals.^{17,35,45} First, s , observed by each curator from each task, is positively correlated. Accordingly, when we randomly select another curator $\hat{c} \in \hat{C}_t$, both $Pr(s_c = 0 | s_{\hat{c}} = 0) > Pr(s_c = 0)$ and $Pr(s_c = 1 | s_{\hat{c}} = 1) > Pr(s_c = 1)$ hold for all c and \hat{c} , regardless of the tasks.⁴⁶ This requires the propensity of assigned \hat{G}_t and the peer curators of c to be somewhat homogeneous throughout each period.⁴⁷ *CitedTCR* with a citation graph ensures such an environment by curator assignment

Table 2. Notations for DG13 in *CitedTCR*

Notations	Meanings
$s_c^{\dot{G}_t}$	A signal ($\in \{0, 1\}$) that a curator $c \in \dot{C}_t$ observed from \dot{G}_t .
s_c	One of the signals that a curator $c \in \dot{C}_t$ has observed so far.
\hat{c}	Another curator who was randomly selected from \dot{C}_t , as a peer of c .
R_t	Set of all reports that <i>CitedTCR</i> has elicited until period t (<i>i.e.</i> , stock).
R_c	A subset of R_t , containing all reports that a curator c has done.
r_c	One of the reports that a curator $c \in \dot{C}_t$ has done so far.
R_c^*	A special case of R_c , whose elements are all truthful reports.

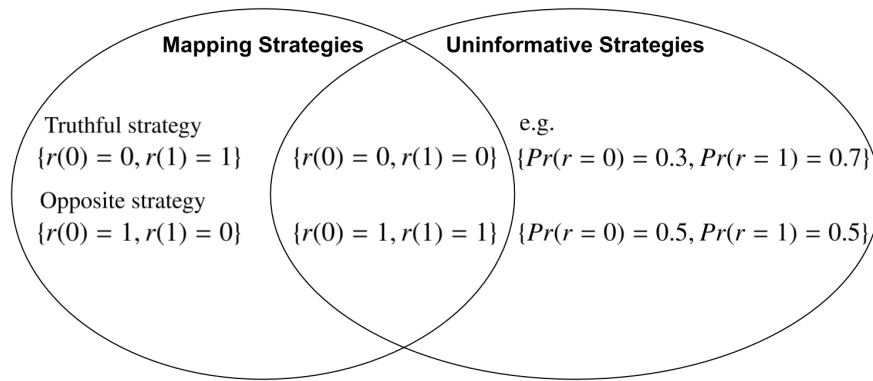


Fig. 2. Nodes (curators) can use either mapping or uninformative strategies for reporting. The strategy of always reporting 0 or always reporting 1 can be classified as both mapping and uninformative strategies.

based on PPR; this is unlike recent multi-task peer prediction,^{38,48} which becomes complex to relax this assumption. The second assumption is that each curator must select one reporting strategy from feasible choices. The set of feasible strategies in our model, presented in Fig. 2, is the union of mapping strategies and uninformative signal-independent strategies. Mapping strategies follow a mapping rule from signals to reports; however, the reports in uninformative strategies follow a given stochastic distribution independent of the observed signals. For the four possible mapping strategies under the assumption of binary signals, we specifically define a strategy that always reports truth as a *truthful* strategy, and a strategy that always reports non-truth as an *opposite* strategy.

Finally, if we let $R_c \subset R_t$ be the set of all (intertemporal) reports that c has provided for multiple \dot{G}_t s, and let R_c^* be a special case in which all elements are truthful reports (*i.e.*, c adopts a truthful strategy), the achievement of DG13 can be defined as follows:

Definition 1 Strong truthfulness: A mechanism satisfies *strong truthfulness* if $\mathbb{E} \left[\theta_c^{\dot{G}_t} \mid R_c^*, R_{\hat{c}}^* \right] \geq \mathbb{E} \left[\theta_c^{\dot{G}_t} \mid R_c, R_{\hat{c}} \right]$ holds for all $c, \hat{c}, R_c, R_{\hat{c}}$, and \dot{G}_t , where equality occurs only when both c and \hat{c} adopt the opposite strategy.⁴⁹

In other words, compared to any other strategy, the mechanism satisfying strong truthfulness

can assign strictly higher expected rewards $\mathbb{E} \left[\theta_c^{\dot{G}_t} \right]$ to the equilibrium by truthful strategies for almost all cases.

DG13, as a multi-task peer prediction mechanism, computes c 's reward $\theta_c^{\dot{G}_t}$ using not only the reports that c and randomly selected \hat{c} produced in period t (i.e., $r_c^{\dot{G}_t}, r_{\hat{c}}^{\dot{G}_t}$) but also all reports that c and \hat{c} produced until period t (i.e., $R_c, R_{\hat{c}}$). According to the original report,¹⁷ and a subsequent report for its generalization,³⁷ DG13 can be formulated as

$$\theta_c^{\dot{G}_t} = \delta_{r_c^{\dot{G}_t}, r_{\hat{c}}^{\dot{G}_t}} - \delta_{r_c \in \{R_c \setminus \dot{R}_t\}, r_{\hat{c}} \in \{R_{\hat{c}} \setminus \dot{R}_t\}},$$

where we use the following Kronecker's delta for the sake of convenience:

$$\delta_{x,y} = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{if } x \neq y \end{cases}$$

Here, $\delta_{r_c^{\dot{G}_t}, r_{\hat{c}}^{\dot{G}_t}}$ is the reward for curation in period t . It is apparent that a value of 1 is obtained when two reports for \dot{G}_t return the same signal ($r_c^{\dot{G}_t}, r_{\hat{c}}^{\dot{G}_t}$) = (0, 0) or (1, 1); otherwise, the value is 0. $\delta_{r_c \in \{R_c \setminus \dot{R}_t\}, r_{\hat{c}} \in \{R_{\hat{c}} \setminus \dot{R}_t\}}$ is a type of penalty that randomly selects two reports r_c and $r_{\hat{c}}$ produced by each curator before period t and compares them in the same manner. Assuming that c and \hat{c} always report 1 for assigned tasks irrespective of the signals, $\theta_c^t = 0$ holds because the penalty term becomes 1 even though $r_c^{\dot{G}_t}$ and $r_{\hat{c}}^{\dot{G}_t}$ always represents a reward of 1. A similar result would be derived for the case of a 50-50 uninformative strategy (i.e., $Pr(r = 0) = Pr(r = 1) = 0.5$) because the expected value of reward terms and penalty terms both become 0.5. Although $\theta_c^{\dot{G}_t}$ takes the interval $[-1, 1]$ because of the penalty, all rewards can be non-negative by adding 1 to all $\theta_c^{\dot{G}_t}$ as a basic reward.

Dasgupta and Ghosh indicated that the expected (net) reward $\mathbb{E} \left[\theta_c^{\dot{G}_t} \right]$ is maximized in the equilibrium in which all curators adopt a truthful strategy by exerting efforts on signal observation under the assumption of positively correlated signals.¹⁷

Theorem 1: DG13 satisfies strong truthfulness.

See Appendix B for the proof of this theorem. Note that DG13 in *CitedTCR* must collectively compute rewards for previous reports after c and \hat{c} both finish reporting three times. Three is the number that satisfies the minimum requirements for establishing multi-task peer prediction without loss of generality : (i) two users, (ii) three total tasks, and (iii) two or more tasks per user, including at least one common task.³⁷ Although each node curates many \dot{G}_t s during $\{G_t\}_{t=0}^{\infty}$ (as long as it has high quality), *CitedTCR* with iterative reward computation cannot satisfy (iii) when either c or \hat{c} produces a report for the first time. Because of this, we postpone reward computation until both c and \hat{c} are sure to meet all minimum requirements by three reports,⁵⁰ and thus, DG13 can elicit truthful reports from curators.

4. Experimental Studies

Although Section 3 describes the utility of PPR and DG13, our study must assess how their combination contributes to the construction of the reliable list G_t . In this section, we perform this assessment experimentally using two step-wise simulations that are both based on the




Fig. 3. Our experiments use a DAG structure with 1,421 time-ordered nodes, where green represents the citation relationships of the first 421 nodes, while red represents the citation relationships of the last 1,000 nodes. We consider the state transition $\{G_t\}_{t=0}^{1000}$ by letting the green (subgraph) be G_0 .

DAG-structured dataset formatted from the arXiv high-energy physics theory (HEP-TH) citation network. In particular, the simulation first uses only PPR to examine the strength of the incentive for registering high-quality content. It then incorporates DG13 to confirm the incentive for eliciting informative reports. All materials used for this experiment are available in the Github repository.⁵¹

4.1. Dataset—The arXiv HEP-TH citation network is a dataset provided by Stanford Network Analysis Project (SNAP), which contains the citation relationships of academic papers in the HEP-TH category submitted from January 1993 to April 2003.⁵² We selected one component with 1,421 papers since January 2000, and constructed a DAG structure as depicted in Fig. 3 (powered by Cytoscape).⁵³ Here, the green component represents the citation relationships of the first 421 nodes, while the red component represents the citation relationships of the last 1,000 nodes (*i.e.*, the green part is a subgraph of the DAG structure). Our experiments consider the green component the initial state G_0 and consider the state transition $\{G_t\}_{t=0}^{1000}$ by sequentially adding the nodes and edges in the red component to G_t .

4.2. Incentive for registering high-quality content—Thus far, we have assumed that *CitedTCR* tends to select curators more frequently from nodes that are regarded as important in G_t , which serves as an incentive for users to register high-quality content. However, this assumption is not obvious because the curator assignment in each period is weighted by the PPR algorithm, which excludes even base nodes from the candidate list. To determine the true strength of the incentive for registering high-quality content, our first experiment computes the correlation between the frequency distribution for 1,421 nodes to be selected as a curator because of sequential assignments up to G_{1000} , and the (not simplified) PageRank score for 1,421 nodes in G_{1000} .⁵⁴

Here, the former designates the number of opportunities in which each node can earn rewards as a curator for the state transition $\{G_t\}_{t=0}^{1000}$, while the latter designates the importance of each node from the viewpoint of the entire DAG in G_{1000} . We specifically computed Spearman's rank correlation coefficient of these values 10 times for each 20 cases with a different number of assigned curators: $n = \{1, 2, \dots, 20\}$.^{55,56}

Fig. 4 summarizes the trend of 200 derived correlation coefficients in a box plot that depicts the median value as orange lines, 25/75 percentile as boxes, pseudo-maximum/minimum value as bars, and outliers as circles. This figure reveals that all correlation coefficients are within the range of 0.4 to 0.7, which can be regarded as moderately correlated. Moreover, they begin to converge between 0.65 and 0.7 when n exceeds 10. These results indicates that *CitedTCR* can retain sufficient incentive to register high-quality content, especially when it assigns more than 10 curators to \hat{G}_t , even though curator assignment relies on the PPR algorithm without base nodes.

4.3. Incentive for eliciting informative reports—After the simulation of curator assignment, the second experiment adds the DG13 mechanism to the first experiment to compute the expected reward $\mathbb{E}[\theta_c^{\hat{G}_t}]$ stemming from $r_c^{\hat{G}_t}$ and $r_c^{\hat{G}_t}$. To simulate the settings of DG13, in which the user reports the received signal $s \in \{0, 1\}$ according to a given strategy, we stochastically allocate the strategy and $s \in \{0, 1\}$ in advance to all 1,421 nodes. In this experiment, the nodes are assumed to use either the truthful strategy or the aforementioned 50-50 uninformative strategy. The allocation of the two strategies is subject to the exogenous randomness parameter $\varepsilon = \{0.0, 0.1, \dots, 1.0\}$, where the expected number of nodes with the uninformative strategy is $\varepsilon \cdot 1,421$, and the expected number of nodes with the truthful strategy is $(1 - \varepsilon) \cdot 1,421$. Similarly, $s \in \{0, 1\}$ is allocated to 1,421 nodes by another exogenous parameter $Pr(s = 0) = \{0.0, 0.1, \dots, 1.0\}$. We computed $\mathbb{E}[\theta_c^{\hat{G}_t}]$ by averaging the total reward generated in $\{G_t\}_{t=0}^{1000}$ for each of the 121 environments comprising different allocations of these two exogenous parameters $\{0.0, 0.1, \dots, 1.0\} \times \{0.0, 0.1, \dots, 1.0\}$, in which $n = 10$ and $m = 0$ are fixed in any environment (*i.e.*, \hat{G}_t is always accepted into G_t regardless of the reports).

Fig. 5 summarizes our results in 11 graphs with different $Pr(s = 0)$, which depicts the trend that $\mathbb{E}[\theta_c^{\hat{G}_t}]$ in the same $Pr(s = 0)$ is maximized when all curators use the truthful strategy (*i.e.*, $\varepsilon = 0.0$), even though the amount of maximized $\mathbb{E}[\theta_c^{\hat{G}_t}]$ decreases as $Pr(s = 0)$ deviates from 0.5 and becomes indifferent with respect to ε if $Pr(s = 0) = 0.0$ or 1.0 .⁵⁷ This result is consistent with the *strong truthfulness* discussed in Section 3.3 and indicates that *CitedTCR* retains incentive to elicit informative reports from assigned curators through DG13-based rewards.

5. Conclusion

In this study, we proposed *CitedTCR*, which incorporates the expertise of anonymous curators into existing TCRs by constructing a reliable citation graph, which is a common proxy for measuring the quality of technical content (*e.g.*, academic papers, patents). To achieve this enhancement on a public peer-to-peer network, we leveraged the PPR algorithm and DG13 mechanism, where the former assigns appropriate curators and the latter elicits informative reports from the assigned curators. As a hybrid of network-based and token-based recommender systems, the combination of previous methods can lead to an incentive design that provides more reward tokens to users as they register high-quality content and continue producing informative reports. Although




Fig. 4. The first experiment computes Spearman’s rank correlation coefficients between the frequency distribution of curator assignment up to G_{1000} and the PageRank score to the DAG in G_{1000} . The box plot for all 200 coefficients (10 times for each $n = \{1, 2, \dots, 20\}$) represents the moderate positive correlation, which increases as n increases and converges between 0.65 and 0.7. This result supports our assumption that *CitedTCR* tends to select curators more frequently from nodes that are considered important in G_t .

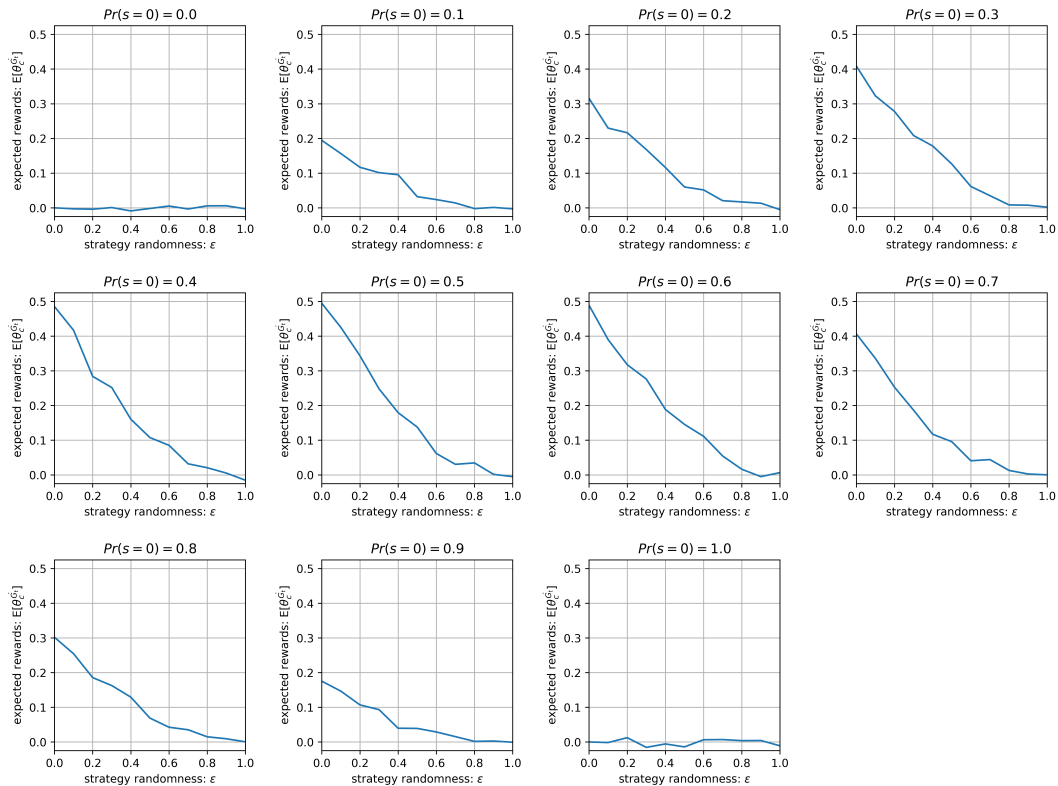


Fig. 5. The second experiment computes $\mathbb{E} [\theta_c^{G_t}]$, which varies depending on strategy randomness ϵ and signal distribution $Pr(s = 0)$ in G_t . 11 graphs for 121 $\{0.0, 0.1, \dots, 1.0\} \times \{0.0, 0.1, \dots, 1.0\}$ cases reveal that $\mathbb{E} [\theta_c^{G_t}]$ is maximized when all curators select the truthful strategy (i.e., $\epsilon = 0.0$), except the $Pr(s = 0) = 0.0$ or 1.0 case. This result is consistent with *strong truthfulness*, which can elicit informative reports from assigned curators in *CitedTCR*.

this incentive design has a different approach than existing TCRs that involve token staking, *CitedTCR* has sufficient utility, which was confirmed theoretically and experimentally. This study can contribute to the emerging discussion on TCRs through its use of a citation graph and peer-prediction mechanism.

However, for practical implementation of this proposal, two remaining issues must be addressed in future work. One involves relaxing the strong assumption of a one-to-one correspondence between users and nodes. Despite the importance of being spam- and sybil-proof for the robustness of peer-to-peer systems, *CitedTCR* without one-to-one correspondence is vulnerable to such attacks because the role of the applicant and its curators can easily overlap if users can create many sybil accounts or post many contents to G_t . To overcome these attacks, an environment may be required in which curators are selected not from V_t , but from U_t , and U_t has no incentive to create sybil accounts when posting multiple content. The indices or algorithms for addressing similar issues have been proposed in the relevant fields of *CitedTCR* such as *SocialRank* in network-based recommender systems,⁵⁸ *h-index* in citation analysis,⁵⁹ and *Proof of Stake* in blockchain.^{60,61} It is therefore a topic for future research to assess the availability of such existing studies in *CitedTCR*.

The second remaining task is to design a valuable reward token. Although this study assumes that users act to maximize the amount of reward tokens, the power of tokens as an incentive is subject to their value, which is determined based on their utility, scarcity, and sustainability. *CitedTCR* therefore requires additional mechanisms to ensure the value of reward tokens as in the Bitcoin protocol, where block-reward halving fixes total supply, and difficulty adjustment stabilizes hash rate. A potential approach is to charge every applicant a token-based registration fee whose price is elastic and based on the frequency with which \hat{G}_t is proposed in a given period.⁶² This approach is worth considering as a registration fee gives the reward token a utility and can serve to prevent spam attacks.

Notes and References

¹ Nakamoto, S. “Bitcoin: A Peer-to-Peer Electronic Cash System.” (2008) (accessed 3 December 2019) <https://bitcoin.org/bitcoin.pdf>.

² This situation is often referred to as *Oracle problem*.

³ Goldin, M. “Token-Curated Registries 1.0.” *Medium* (accessed 2 April 2019) <https://medium.com/@ilovebagels/token-curated-registries-1-0-61a232f8dac7>.

⁴ Goldin, M. “Token Curated Registries 1.1, 2.0 TCRs, New Theory, and Dev Updates.” *Medium* (accessed 3 April 2019) <https://medium.com/@ilovebagels/token-curated-registries-1-1-2-0-tcrs-new-theory-and-dev-updates-34c9f079f33d>.

⁵ See Section 2 for some application examples.

⁶ Lockyer, M. “Token Curated Registry (TCR) Design Patterns.” *Medium* (accessed 3 April 2019) <https://hackernoon.com/token-curated-registry-tcr-design-patterns-4de6d18efa15>.

⁷ Wang and Krishnamachari referred to the TCR for binary choices as *objective TCR*.¹⁵

⁸ As described later, *CitedTCR* uses the Personalized PageRank (PPR) algorithm to measure both similarity and centrality.

⁹ We will confirm the strength of this incentive in Section 4.

¹⁰ As a similar incentive, TCRs using the token staking often require applicants to stake a certain amount of their token on {accept} choice.

¹¹ Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. V. *Algorithmic Game Theory*. Cambridge: Cambridge University Press (2007).

¹² As described in Section 3, x is listed including its citation relationship if the number of {accept} reports exceeds a given threshold.

¹³ Hemenway Falk, B., Tsoukalas, G. “Token-Weighted Crowdsourcing.” *SSRN* (accessed 3 December 2019) <https://dx.doi.org/10.2139/ssrn.3295811>.

¹⁴ Bulkin, A. “Curate This: Token Curated Registries That Don’t Work.” *Medium* (accessed 15 April 2019) <https://blog.coinfund.io/curate-this-token-curated-registries-that-dont-work-d76370b77150>.

¹⁵ Wang, Y. L., Krishnamachari, B. “Enhancing Engagement in Token-Curated Registries via an Inflationary Mechanism.” *arXiv* (2018) (accessed 3 December 2019) <https://www.arXiv.org/abs/1811.09680>.

¹⁶ Haveliwala, T. H. “Topic-Sensitive Pagerank.” In *Proceedings of the 11th International Conference on World Wide Web* ACM 517–526 (2002) <https://doi.org/10.1145/511446.511513>.

¹⁷ Dasgupta, A., Ghosh, A. “Crowdsourced Judgement Elicitation with Endogenous Proficiency.” In *Proceedings of the 22nd International Conference on World Wide Web* ACM 319–330 (2013) <https://doi.org/10.1145/2488388.2488417>.

¹⁸ <https://publisher.adchain.com/>.

¹⁹ <https://oceanprotocol.com/>.

²⁰ <https://registry.civil.co/registry/approved>.

²¹ Balasanov, S. “TCR Design Flaws: Why Blockchain Needs Reputation.” *Medium* (accessed 15 April 2019) <https://blog.relevant.community/tcr-design-flaws-why-blockchain-needs-reputation-c5771d97b210>.

²² Token Curated Registry. “The Token Curated Registry Reading List.” *Medium* (accessed 15 April 2019) <https://medium.com/@tokencuratedregistry/the-token-curated-registry-whitepaper-bd2fb29299d6>.

²³ No Author. “Token Curated Registries — Messari - Crypto News, Pricing, and Research.” *Messari* (2018) (accessed 15 April 2019) <https://messari.io/resource/token-curated-registries>.

²⁴ Asgaonkar, A., Krishnamachari, B. “Token Curated Registries-A Game Theoretic Approach.” *arXiv* (2018) (accessed 3 December 2019) <https://www.arxiv.org/abs/1809.01756>.

²⁵ Brin, S., Page, L. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” *Computer Networks and ISDN systems* **30.1-7** 107–117 (1998) [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X).

²⁶ Page, L., Brin, S., Motwani, R., Winograd, T. “The PageRank Citation Ranking: Bringing Order to the Web.” *Stanford InfoLab* (1999) (accessed 3 December 2019) <http://ilpubs.stanford.edu:8090/422/>.

²⁷ For this property, PPR is often referred to as the random walk with restart (RWR) algorithm.

²⁸ Gori, M., Pucci, A. “Research Paper Recommender Systems: A Random-Walk Based Approach.” In *2006 IEEE/WIC/ACM International Conference on Web Intelligence (WI 2006 Main Conference Proceedings)(WI'06)* IEEE 778–781 (2006) <https://doi.org/10.1109/WI.2006.149>.

²⁹ Liu, X., Suel, T., Memon, N. “A Robust Model for Paper Reviewer Assignment.” In *Proceedings of the 8th ACM Conference on Recommender Systems* ACM 25–32 (2014) <https://doi.org/10.1145/2645710.2645749>.

³⁰ Küçüküntç, O., Saule, E., Kaya, K., Çatalyürek, Ü. V. “Recommendation on Academic Networks Using Direction Aware Citation Analysis.” *arXiv* (2012) (accessed 3 December 2019) <https://arxiv.org/abs/1205.1143>.

³¹ No Author. “NEM Technical Reference.” *nem.io* (2018) (accessed 3 December 2019) https://nem.io/wp-content/themes/nem/files/NEM_techRef.pdf.

³² Miller, N., Resnick, P., Zeckhauser, R. “Eliciting Informative Feedback: The Peer-Prediction Method.” *Management Science* **51.9** 1359–1373 (2005) <https://doi.org/10.1287/mnsc.1050.0379>.

³³ Gneiting, T., Raftery, A. E. “Strictly Proper Scoring Rules, Prediction, and Estimation.” *Journal of the American Statistical Association* **102.477** 359–378 (2007) <https://doi.org/10.1198/01621450600001437>.

³⁴ See Faltings’ and Radanovic’s *Game Theory for Data Science: Eliciting Truthful Information* for more comprehensive review on peer-prediction method and other information elicitation models.⁶³

³⁵ Jurca, R., Faltings, B. “Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms.” In *International Workshop on Internet and Network Economics* Springer 268–277 (2005) https://doi.org/10.1007/11600930_26.

³⁶ Uninformative equilibria are designated as a *blind agreements* in original DG13.¹⁷

³⁷ Shnayder, V., Agarwal, A., Frongillo, R., Parkes, D. C. “Informed Truthfulness in Multi-Task Peer Prediction.” In *Proceedings of the 2016 ACM Conference on Economics and Computation* ACM 179–196 (2016) <https://doi.org/10.1145/2940716.2940790>.

³⁸ Agarwal, A., Mandal, D., Parkes, D. C., Shah, N. “Peer Prediction with Heterogeneous Users.” In *Proceedings of the 2017 ACM Conference on Economics and Computation* ACM 81–98 (2017) <https://doi.org/10.1145/3033274.3085127>.

³⁹ Goel, N., Filos-Ratsikas, A., Faltings, B. “Decentralized Oracles via Peer-Prediction in the Presence of Lying Incentives.” (2019) Presented at the 14th Tinbergen Institute Conference “Bayesian Crowd” at Erasmus University Rotterdam, 24 June 2019, https://lia.epfl.ch/~goel/upload/doc/papers/2019/goel_outside_incentives.pdf.

⁴⁰ Decentralized oracle is a broader concept than TCR, which includes every DApp responsible for consensus-building on off-chain contents, *i.e.*, TCR can be interpreted as one of the decentralized oracle systems. The term decentralized oracle is often used in the context of prediction market, and representative platforms (*e.g.*, Augur and Gnosis) use the token staking for their consensus building as with the case of TCRs.^{64,65}

⁴¹ Therefore, when managing *CitedTCR*, we need to prepare in advance an initial state G_0 with a sufficient number of nodes and edges.

⁴² Here, the condition $n \geq 2$ is important for DG13 mechanism as we will see in Section 3.3. This condition provides $|V_x|$ with an upper limit number which must satisfy $|V_t \setminus V_x| \geq n$ for all t .

⁴³ Although original paper uses Simplified PageRank as an introduction of model description,²⁶ PageRank is the dominant eigenvector of the matrix $\mathbf{P}_{PR} = (1 - \alpha)\mathbf{P} + \alpha(1/|V_t|)\mathbf{1}$, where $\mathbf{1}$ is $|V_t| \times |V_t|$ matrix whose elements are all 1. Namely, \mathbf{P}_{PR} quantifies the random walking which, with probability α , jumps to one of all existing nodes uniformly at random (random-surfer model). This is to make PageRank work even in the directed network including dead-end loop or the node with no out-edges.

⁴⁴ Note that even this modification cannot completely eliminate the biased curation, as long as the curation affects the future structure of G_t . Analyzing the strength of this bias is one of our future tasks.

⁴⁵ Witkowski, J., Parkes, D. C. “Peer Prediction Without a Common Prior.” In *Proceedings of the 13th ACM Conference on Electronic Commerce* ACM 964–981 (2012) <https://doi.org/10.1145/2229012.2229085>.

⁴⁶ Accordingly, $Pr(s_c = 1 | s_{\hat{c}} = 0) < Pr(s_c = 1)$ and $Pr(s_c = 0 | s_{\hat{c}} = 1) < Pr(s_c = 0)$ hold, simultaneously.

⁴⁷ The homogeneity required for positively correlated signals is not as strong in binary signals as in multiple signals.

⁴⁸ Mandal, D., Leifer, M., Parkes, D. C., Pickard, G., Shnayder, V. “Peer Prediction With Heterogeneous Tasks.” *arXiv* (2016) (accessed 3 December 2019) <https://arxiv.org/abs/1612.00928>.

⁴⁹ The original definition generalizes both truthful strategy and opposite strategy as a *permutation strategy* to encompass the case of multiple (non-binary) signals.³⁷

⁵⁰ Two reports cannot satisfy (ii) if c and \hat{c} share two tasks.

⁵¹ https://github.com/knskito/materials_CitedTCR.

⁵² <https://snap.stanford.edu/data/cit-HepTh.html>.

⁵³ Shannon, P., *et al.* “Cytoscape: A Software Environment for Integrated Models of Biomolecular Interaction Networks.” *Genome Research* **13.11** 2498–2504 (2003) <https://doi.org/10.1101/gr.1239303>.

⁵⁴ We set $\alpha = 0.15$ in both the PageRank and the PPR algorithms.

⁵⁵ We cannot use Pearson correlation coefficient because both frequency distribution and PageRank scores follow not normal distribution but power-law distribution.

⁵⁶ Correlation coefficients are different in each of 10 computations because curators are assigned stochastically according to PPR algorithm, contrary to the constant PageRank score.

⁵⁷ This is because truthful strategy becomes indifferent with an uninformative strategy that always returns $r = 0$ (or 1), if $Pr(s = 0) = 0.0$ (or 1.0) holds. Note that $Pr(s = 0) = 0.0$ and 1.0 are outside the scope of DG13 as we cannot put the assumption of positively correlated signals on the environments.

⁵⁸ Tsai, M.-H., Aggarwal, C., Huang, T. “Ranking in Heterogeneous Social Media.” In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining* ACM 613–622 (2014) <https://doi.org/10.1145/2556195.2556254>.

⁵⁹ Hirsch, J. E. “An Index to Quantify an Individual’s Scientific Research Output.” *Proceedings of the National Academy of Sciences* **102.46** 16569–16572 (2005) <https://doi.org/10.1073/pnas.0507655102>.

⁶⁰ Gui, G., Hortacsu, A., Tudon, J. “A Memo on the Proof-of-Stake Mechanism.” *arXiv* (2018) (accessed 3 December 2019) <https://arxiv.org/abs/1807.09626>.

⁶¹ Saleh, F. “Blockchain Without Waste: Proof-of-Stake.” *SSRN* (2018) (accessed 3 December 2019) <https://dx.doi.org/10.2139/ssrn.3183935>.

⁶² This concept corresponds to the difficulty adjustment in Bitcoin protocol.

⁶³ Faltings, B., Radanovic, G. *Game Theory for Data Science: Eliciting Truthful Information*. San Rafael: Morgan & Claypool Publishers (2017).

⁶⁴ Peterson, J., Krug, J., Zoltu, M., Williams, A. K., Alexander, S. “Augur: A Decentralized Oracle and Prediction Market Platform.” *arXiv* (2015) (accessed 3 December 2019) <https://arxiv.org/abs/1501.01042>.

⁶⁵ No Author. “Gnosis Whitepaper.” *Gnosis.io* (2017) (accessed 3 December 2019) <https://gnosis.io/pdf/gnosis-whitepaper.pdf>.

⁶⁶ If we assume the cost of curation as c , the expected rewards in this example become $\mathbb{E}(k) = p(k - c) - (1 - p)(q + c)$. This extension shifts the condition for $\mathbb{E}(k) = 0$, from $\frac{p/(1-p)}{n^*/(n-n^*)} = 1$ to $\frac{p/(1-p)}{n^*/(n-n^*)} = \frac{q+c}{q-\frac{n^*}{n-n^*}c}$, where the right-hand side of the new condition must be greater than one.

⁶⁷ Furthermore, if we designate $Pr(s_c = 0, s_{\hat{c}} = 0) - Pr(s_c = 0)Pr(s_{\hat{c}} = 0) = P_{00}$, $Pr(s_c = 0, s_{\hat{c}} = 1) - Pr(s_c = 0)Pr(s_{\hat{c}} = 1) = P_{01}$, $Pr(s_c = 1, s_{\hat{c}} = 0) - Pr(s_c = 1)Pr(s_{\hat{c}} = 0) = P_{10}$, $Pr(s_c = 1, s_{\hat{c}} = 1) - Pr(s_c = 1)Pr(s_{\hat{c}} = 1) = P_{11}$, they have the following relations: $P_{00} = P_{11}$, $P_{01} = P_{10}$, $P_{00} + P_{01} + P_{10} + P_{11} = 0$.

Appendix A: Expected Rewards in a Simple Token-Staking Scheme

Consider a simple token staking example in which n curators stake a fixed q number of tokens on one of the options. Let k be the amount of (net) rewards that curators can obtain when their selections become the consensus, and let p be the curators’ subjective probability of the realization of this event. Then, the expected reward in this example is $\mathbb{E}(k) = pk - (1 - p)q$.

Specifically, k is the redistribution of the total staked tokens nq among the curators who have staked on the consensus with the exception of one’s own stake, q . Accordingly, if we let n^* be the number of curators who have staked on the consensus, $k = \frac{n}{n^*}q - q = \frac{n-n^*}{n^*}q$. By substituting this into the equation of $\mathbb{E}(k)$, we can derive the following condition:

$$\mathbb{E}(k) \begin{cases} > \\ = \\ < \end{cases} 0, \text{ if } \frac{p/(1-p)}{n^*/(n-n^*)} \begin{cases} > \\ = \\ < \end{cases} 1,$$

where $\frac{p/(1-p)}{n^*/(n-n^*)}$ represents the odds ratio between the expected and actual value of the probability of one's choice becoming the consensus; *i.e.*, the expected reward in the model takes a positive value only when we estimate the odds to be higher than their actual value and is zero as long as our estimation is precise (as a result of the zero-sum game). Furthermore, the expected reward under precise odds estimation is negative if we take the cost of curation into account.⁶⁶

These results reveal that the token-staking scheme does not have sufficient incentive to engage curators in consensus building. Providing new reward tokens to curators in proportion to the score of the peer-prediction mechanism is one possible approach to this problem.

Appendix B: Proof of the Strong Truthfulness of the DG13 Mechanism

This proof uses notations that are compatible with Section 3.3. The expected value of the reward term $\delta_{r_c^{\hat{G}_t}, r_{\hat{c}}^{\hat{G}_t}}$ depends not only on the results of $r_c^{\hat{G}_t}$ and $r_{\hat{c}}^{\hat{G}_t}$, but also on the probability distribution of input signals that each node observes in period t , as follows:

$$\mathbb{E} \left[\delta_{r_c^{\hat{G}_t}, r_{\hat{c}}^{\hat{G}_t}} \right] = \sum_{s_c=0}^1 \sum_{s_{\hat{c}}=0}^1 Pr(s_c, s_{\hat{c}}) \cdot \delta_{r_c(s_c), r_{\hat{c}}(s_{\hat{c}})},$$

where $Pr(s_c, s_{\hat{c}})$ is the joint probability distribution of the signals that c and \hat{c} can receive from \hat{G}_t . Note that the right-hand side does not require superscript \hat{G}_t because of the assumption of positively correlated signals.

As described in Section 3.3, the penalty term is the result of the comparison between two randomly picked reports that c and \hat{c} produce prior to period t . We can write the expected value of the penalty in a similar form to the reward term as follows:

$$\mathbb{E} \left[\delta_{r_c \in \{R_c \setminus \dot{R}_t\}, r_{\hat{c}} \in \{R_{\hat{c}} \setminus \dot{R}_t\}} \right] = \sum_{s_c=0}^1 \sum_{s_{\hat{c}}=0}^1 Pr(s_c) Pr(s_{\hat{c}}) \cdot \delta_{r_c(s_c), r_{\hat{c}}(s_{\hat{c}})}.$$

This uses product distribution $Pr(s_c)Pr(s_{\hat{c}})$ rather than joint distribution $Pr(s_c, s_{\hat{c}})$ because the penalty term covers all intertemporal reports included in $R_c \setminus \dot{R}_t$ and $R_{\hat{c}} \setminus \dot{R}_t$.

Consequently, $\mathbb{E}(\theta_c^{\hat{G}_t})$ can be expressed as

$$\mathbb{E} \left[\theta_c^{\hat{G}_t} \right] = \sum_{s_c=0}^1 \sum_{s_{\hat{c}}=0}^1 [Pr(s_c, s_{\hat{c}}) - Pr(s_c)Pr(s_{\hat{c}})] \cdot \delta_{r_c(s_c), r_{\hat{c}}(s_{\hat{c}})}.$$

The terms in square brackets correspond to the correlation of s_c and $s_{\hat{c}}$. If one assumes that $Pr(s_c, s_{\hat{c}}) - Pr(s_c)Pr(s_{\hat{c}}) > 0$, then both $Pr(s_c|s_{\hat{c}}) > Pr(s_c)$ and $Pr(s_{\hat{c}}|s_c) > Pr(s_{\hat{c}})$ hold because $Pr(s_c, s_{\hat{c}}) = Pr(s_c|s_{\hat{c}})Pr(s_{\hat{c}}) = Pr(s_{\hat{c}}|s_c)Pr(s_c)$, *i.e.*, s_c and $s_{\hat{c}}$ are positively correlated in this case.

Because DG13 assumes positively correlated binary signals, the following condition holds in

the expanded form of $\mathbb{E}(\theta_c^{\hat{G}_t})$:

$$\begin{aligned} \mathbb{E} \left[\theta_c^{\hat{G}_t} \right] &= [Pr(s_c = 0, s_{\hat{c}} = 0) - Pr(s_c = 0)Pr(s_{\hat{c}} = 0)]_{>0} \cdot \delta_{r_c(0), r_{\hat{c}}(0)} \\ &+ [Pr(s_c = 0, s_{\hat{c}} = 1) - Pr(s_c = 0)Pr(s_{\hat{c}} = 1)]_{<0} \cdot \delta_{r_c(0), r_{\hat{c}}(1)} \\ &+ [Pr(s_c = 1, s_{\hat{c}} = 0) - Pr(s_c = 1)Pr(s_{\hat{c}} = 0)]_{<0} \cdot \delta_{r_c(1), r_{\hat{c}}(0)} \\ &+ [Pr(s_c = 1, s_{\hat{c}} = 1) - Pr(s_c = 1)Pr(s_{\hat{c}} = 1)]_{>0} \cdot \delta_{r_c(1), r_{\hat{c}}(1)}, \end{aligned}$$

where $[x]_{>0}$ and $[x]_{<0}$ indicate that x is positive and negative, respectively.⁶⁷

It is apparent that $\mathbb{E}(\theta_c^{\hat{G}_t})$ is maximized only when both c and \hat{c} provide truthful reports ($r(0) = 0, r(1) = 1$) or opposite reports ($r(0) = 1, r(1) = 0$). Any other pattern, such as nodes using asymmetric strategies or always reporting the same signal, produces less expected values. Under the assumption of using one reporting strategy, this outcome indicates that $\mathbb{E}(\theta_c^{\hat{G}_t})$ is maximized only when both x and \hat{c} adopt either a truthful or opposite strategy. Thus, DG13 satisfies strong truthfulness.



Articles in this journal are licensed under a Creative Commons Attribution 4.0 License.

Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.