Received: 19 February 2021 | Revised: 3 April 2021 | Accepted: 22 April 2021

DOI: 10.1002/nem.2179

SPECIAL ISSUE PAPER

WILEY

Truthful Decentralized Blockchain Oracles

Yuxi Cai[1] | **Nafis Irtija**[2] | **Eirini Eleni Tsiropoulou**[2] | **Andreas Veneris**[1,3]

[1]Dept. of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada

[2]Dept. of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, USA

[3]Dept. of Computer Science, University of Toronto, Toronto, ON, Canada

Correspondence
Eirini Eleni Tsiropoulou, Dept. of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, USA.
Email: eirini@unm.edu

Summary

Blockchain systems rely on oracles to bridge external information to the decentralized applications residing in the systems. Astraea protocols are decentralized oracle designs utilizing majority-voting mechanism to determine the oracle outcomes and/or rewards to voters. However, the voters are indifferent between voting through a single or multiple identities, as the potential rewards by the decentralized oracles grow linearly with the voters stakes. Additionally, the majority-voting mechanism may facilitate herd behaviors among the voters, as the voters are rewarded only if they are in agreement with the majority outcomes. In this paper, a novel oracle protocol is introduced by proposing a peer prediction-based scoring scheme along with non-linear staking rules, aiming at extracting subjective data truthfully. Specifically, an incentive compatible scoring scheme is designed so that voters uniquely maximize their expected score by honest reporting. The voters are rewarded when their report achieves a relatively high score compared to the rest of the voters, as opposed to the existing schemes, where a reward is only given when they agree to the majority. Furthermore, a non-linear stake scaling rule is proposed to discourage Sybil attacks. Detailed simulation results are presented to show the operation of the proposed oracle protocol and its improvement compared to indicative mechanisms proposed in the existing literature.

1 | INTRODUCTION & MOTIVATION

Distributed Ledger Technology (DLT), or blockchain technology, has attracted the interest of the research and industrial communities due to its inherent key characteristics of decentralization, persistency, anonymity, and auditability.[1] Those characteristics provide enhanced security to transactions occurring among participating nodes, or, distributed entities. Blockchain was initially introduced as a public ledger facilitating monetary transactions.[2] However, nowadays, blockchain technology is able to support a wide variety of applications by exploiting smart contracts.[3–5] The execution of smart contracts usually needs information external to the blockchain system.[6] When the external information is necessary, verification and consensus on the information is critical to the overall security of the system. Otherwise, it will undermine the benefits of the distributed ledger and inevitably pose a security risk to the overall system.

Oracles have been introduced in the recent literature as trusted entities that provide verified external data to the blockchain systems. Among the pioneering research works in the field of oracles, Astraea protocols[7,8] constitute a series of decentralized blockchain oracle proposals. The Astraea protocols collect information from the participants (*i.e.*, voters or verifiers) by making use of staked voting mechanisms that are agnostic to the blockchain consensus mechanisms, while guaranteeing permissionless participation of the voters. The main benefit of the two proposed Astraea protocols is their low-complexity, and deterministic information collection and reward mechanisms. In order to determine the outcome of a query submitted to the oracle, as well as to distribute the rewards to the voters, both protocols adopt majority voting (*i.e.*, that is based solely on the *popularity* of a particular answer). However, such an approach may lead to several undesirable outcomes.

One of the possible undesirable voting behaviors is *lazy voting*, where a voter always votes for `True` or `False`. To discourage such a behavior, the paired-question protocol[8] has been introduced on top of the first Astraea protocol. It requires the submitters to submit antithetic proposal pairs, thus, there will be expectedly an equal amount of `True` and `False` proposals. However, there are two major drawbacks to this technique. First, the voters can enhance the expected reward by *herding*. If they believe that their true opinion belongs to the minority group, they will answer the opposite to their belief.[9] Second, the paired-question protocol cannot safeguard the system from Sybil attacks, *i.e.*, the voters create multiple fake identities. If the amount of identities controlled is significant enough, the attacker will be able to make their opinion dominant. Sybil attacks can be mitigated by increasing the cost of identity creation.[10] The first Astraea protocol introduces a cost linear to the number of identities or voting power as the stake. The paired-question protocol further increases the cost by awarding the attacker only when the antithetic questions have different majority outcomes. However, voters are still indifferent between reporting via multiple identities or a single identity.

Despite the efforts made in the previous works, the problem of jointly extracting truthful information through decentralized oracles, as well as the problem of protecting the blockchain system from Sybil attacks, has not been successfully addressed. In this work, we introduce a peer prediction-based protocol with non-linear scaling of stake to support the operation and the verified information extraction of the decentralized oracles. The main contributions of this research work that differentiate it from the previous protocols are: (1) a novel light-weight scoring rule decides the rewards to the voters such that rational voters report their private opinions truthfully; (2) the voting weight is scaled in a sub-linear manner and the reward portion follows a super-linear function regarding the submitted stake, which consequently decreases the incentive to create multiple identities. Specifically, for each vote, the oracle collects two types of reports from a voter, a binary information report and a popularity prediction report. The oracle outcome regarding the specific question is the weighted majority (based on the associated stake adjusted by a sub-linear function) of the binary information reports. Furthermore, a score is assigned to each report by accounting for the accuracy of prediction and degree of agreement within the voter groups. Subsequently, the oracle rewards only the top-scored voters, while the reward share is determined by their stake adjusted by a super-linear function.

Based on the above discussion, it should be highlighted that the proposed protocol that facilitates truthful information extraction has two fundamental benefits: (i) It incentivizes the minority-opinion voters to express their true opinion, by providing them a maximized expected reward with honest voting. Thus, the peer prediction-based mechanism is incentive compatible. (ii) The proposed non-linear stake scaling scheme incentivizes the honest voters to stake more onto a single report, as the penalty to a voter trying to bias oracle's outcome by Sybil attack is increased. This research work substantially extends our initial research findings,[11] as described above in detail.

The remainder of the paper is organized as follows. Section 2 provides a detailed literature review of existing blockchain oracles and peer prediction mechanisms. Section 3 describes the operation of a general decentralized oracle and introduces the notation that is used in the rest of the paper. Section 4 provides a detailed description of the novel decentralized oracle protocol, a proof of incentive compatibility of the scoring scheme, as well as a theoretical analysis on the expected outcome. Section 5 presents a guideline on the stake scaling rule and discusses the advantages of the proposed protocol. Section 6 concludes this paper.

2 | RELATED WORK

This section provides an overview on existing literature on blockchain oracles, decentralized voting-based oracles and peer prediction mechanisms.

2.1 | Blockchain oracles

In this paper, we refer to *decentralized oracles* as blockchain oracles with the following properties:

- Permissionless: any individual can join or leave the system without permission from existing users,
- Equi-privileged: all the users in the system have equal priority.

A number of oracles have been introduced in the existing literature. Provable,[12] formly known as Oraclize.it, derives data from a web source specified by a user, while retaining cryptographic proofs to guarantee the legitimacy of the collected information. In order to protect data from potential alteration by malicious operating systems, Town Crier[13] uses the Intel's Software Guard Extensions hardware (IntelSGX).[14] It should be noted that the aforementioned oracles operate with the support of a centralized server to deal with the query requests, thus, their decentralized nature and the equi-privilege property are violated. A validation-dispute protocol is introduced in the prediction market Augur,[5] where the token holders can report and/or challenge the outcomes of the oracle. However, for each query, a privileged reporter has priority over others to report an outcome. As a result, both of the decentralized oracle properties, *i.e.*, permissionlessness and equi-privilege, are violated. Chainlink[15] retrieves aggregated information from multiple oracles through a marketplace. It is notable that with Chainlink: (a) only registered oracles can provide data hence it is not permissionless, and (b) the system is vulnerable to denial of service attacks when a single data source is specified.

A number of oracle proposals make use of Transport Layer Security (TLS) protocols,[16] cryptographic protocols providing communication security, to prove data authenticity without sacrificing decentralization. In the TLS-N[17] protocol, servers include a proof for each TLS session in its responses to queries. This requires modification as well as additional cost of resources, both computation and storage, on the server-side. Practical Data Feed Service (PDFS)[18] authenticates, records, and verifies the TLS transactions with smart contracts to enable data transparency and consistency validations. Furthermore, DECO[19] proves data authenticity by utilizing zero-knowledge proofs. Such an approach does not rely on trusted hardware nor server side modification, while zero-knowledge proofs introduce ineligible computational overhead.

2.2 | ASTRAEA protocols

Three types of participants are involved in the original ASTRAEA[20] mechanism: *submitters*, *voters*, and *certifiers*. When submitting a proposition to the oracle, the submitter pays a bounty that will be used to reward the voters. The latter register through providing a stack, and are randomly assigned a proposition to answer. In contrast, a certifier deposit a relatively large stake, and is allowed to select a proposition to answer. Certifiers are subject to higher risk and higher potential rewards. The certifiers receive rewards from two reward pools depending on the outcome of the responded propositions. Figure 1 highlights the possible monetary flows in the system. To determine the proposition outcome, the protocol compares majority answer from the certifiers and the majority answers derived by the voters. Thus, the following two cases hold true:

- Agreement among the voters and certifiers: The majority answer is the final outcome of the oracle. The certifiers and voters that provided the same answer to the final outcome are rewarded proportionally to their stake. The certifiers and voters that provided the opposite answer to the majority answer forfeit their stake to the reward pools.

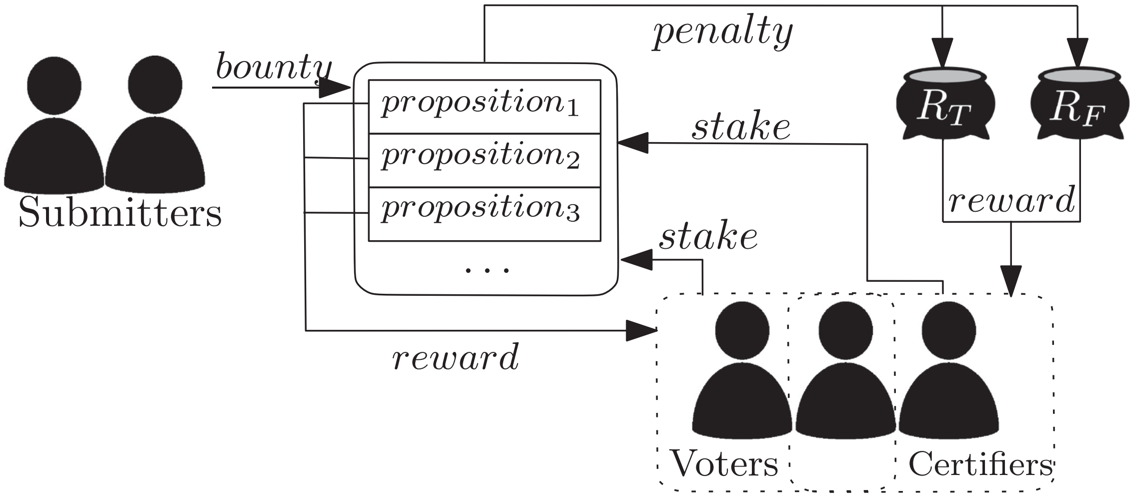

FIGURE 1 Overview of monetary flow in ASTRAEA I protocol

- Disagreement among the voters and certifiers: No answer is determined by the oracle. The certifiers lose all their stake, while the voters' stakes are refunded.

It should be noted that the original Astraea protocol assumes that it is equally likely for a submittor to submit a `True` or a `False` proposition. Thus, the probability of either answer as the oracle outcome is equal. This assumption prevent the draining of reward pools and formation of a lazy equilibrium. On the other hand, it complicates the analysis of the underlying system.

An improved version of the Astraea protocol,[8] the *paired-question* protocol, proposes a simplified outcome determination rule, while at the same time efficiently disincentivizing lazy voting. Specifically, the submitters are required to propose an antithetic question pair (*i.e.*, two questions with potentially opposite binary `True` or `False` answers). Opposing to both certifiers and voters in the original Astraea protocol, all responding entities are referred as voters. A voter responses to a randomly assigned question after depositing a stake. They are rewarded if the majorities outcomes of the question pair disagree with each other, and they agree with the majority in the assigned question. The improved paired-question protocol discourages lazy voting by efficiently lowering the expected payoff as shown in the extensive analysis.[8]

In Shintaku,[21] there is also one voting group, the voters. A voter deposit their stake, and are assigned randomly two propositions. The voter then submits the responses and is rewarded if the responses (i) agree with the majority, and (ii) are opposite to each other. According to the analysis,[8] the penalties for disagreements with the majority should be at least twice as large as the possible rewards in order to make the expected reward of lazy voting negligible.

In brief, existing Astraea protocols determine the oracle outcome and the voter rewards following the philosophy of the majority voting. This mechanism has some inherent drawbacks, as the voters may not reveal their true opinion and act in a dishonest way under the concern that they will end up in the minority group. Thus, the majority voting may discourage the honest reporting of the voters' opinions, if there is an a priori knowledge regarding the "popularity" of the answers. Another drawback is that they are prone to Sybil attacks, as the voters receive rewards that are linear to their submitted stakes. Thus, a voter may use multiple identities in the system to bias the final outcome of the oracle.

2.3 | Peer prediction mechanisms

Peer prediction mechanisms are designed to incentivize the participants to truthfully report private signals (*e.g.*, opinions or experiences) where an observable objective outcome is not available. The concept was first proposed as the *peer prediction method* to address honest feedback elicitation on e-commerce platforms.[22] The key insight is to induce relative truthfulness by correlating reported information, and to mitigate the problem of noisy reports by cooperating rewards. Instead of attesting on the agreement between the feedback, the *peer prediction method* maintains a distribution over all possible responses based on received reports and a prior distribution (that is, a common prior belief shared by the mechanism and all reporters). Although the peer prediction method creates a Nash equilibrium for truthful revelation, it does not discourage lazy behaviour and it requires the mechanism to be aware of the common prior belief.

The Bayesian Truth Serum (BTS) mechanism,[23] though still requiring a common prior belief among all agents, does not require it to be known by the mechanism. Therefore, the assumption is more realistic. BTS mechanism expands the report to include two components: an information report and a prediction report on the distribution of each possible information report. BTS has been shown to be incentive compatible with a large group of agents. This puts the system in jeopardy if the number of agents is small. In contrast, robust Bayesian Truth Serum (RBTS)[24] is based on the same assumption as BTS. It also requires the same two reports from the agents (*i.e.*, information and prediction reports). However, it has been shown to provide strict incentive compatibility for any number of agents, $n \geq 3$. For each agent, a RBTS score is calculated based on the information report from a peer agent and the prediction report from a reference agent. The RBTS score may not be consistent as the scores can vary depending by the choice of the peer/reference agents.

3 | MODEL SETUP

This section details the analysis model of our proposed system by introducing important definitions.

3.1 | Oracle and roles of players

The proposed decentralized oracle operates as a smart contract on a blockchain platform, like Ethereum[25] or Hyperledger.[26] We refer such a contract as as the oracle *executor* in the rest of the analysis. Two types of players participate in the operation of the decentralized oracle: the submitters and the voters. As shown in Figure 2, the submitters submit *Boolean propositions* to the executor, while the voters vote on the randomly assigned propositions. The executor is responsible to keep track of the active propositions, to receive the votes, to determine the outcome of the proposition and the corresponding scores, and to reward the voters for their engagement. A proposition can remain active for a given time duration, which is specified at its submission by the submitter, and after which it is considered closed.

3.2 | Voter belief

We denote the set of voters participating in the system as $\mathcal{V}$. A random subset of voters is selected to answer each proposition and its cardinality is much smaller compared to $|\mathcal{V}|$ to prevent the voters from choosing their preferred propositions successfully. Furthermore, we assume all voters are rational and risk-neutral expected utility maximizers. Voters share the same belief system consisting of *signals* and *states*.[27] For each proposition, voter i has a *private opinion*. The voter's private opinion is represented by a binary random variable $\mathrm{PO}_i \in \{1, 0\}$, where 1 denotes `True` and 0 stands for `False`. A state T is a random variable taking values in $\{1, \ldots, m\}$ ($m \geq 2$) and representing all possible true states of the world (*e.g.*, it is possible that voters who think Aristotle is the greatest Greek philosopher of all time make up 70% of all voters, while it is possible that only 40% of them think so). Each state is a probabilistic distribution on the possible outcomes, consisting of $\Pr(\mathrm{PO}_i = 1 | T = t)$ and $\Pr(\mathrm{PO}_i = 0 | T = t)$ ($t \in \{1, \ldots, m\}$), of the proposition.

In the rest of our analysis, we adopt the *Common Prior Assumption (CPA)*.[28] CPA consists of shared probabilistic distributions of all the states ($\Pr(T = t)$), and common initial beliefs in each state ($\Pr(o | T = t)$), where o denotes an opinion, among all voters. The common prior belief among the voters should be admissible. The common prior belief is admissible iff[24]:

1. There are two or more states, *i.e.*, $m \geq 2$;
2. All states have positive probability, *i.e.*, $\Pr(T = t) > 0$;
3. (*The assortative property*) States are distinct and sorted, *i.e.*, $\Pr(1|m) > \Pr(1|m-1) > \ldots > \Pr(1|1)$; and,
4. The signal beliefs conditional on state are fully mixed, *i.e.*, $0 < \Pr(PO = po | T = t) < 1 \ \forall t \in \{1 \ldots m\}$ and $po \in \{1, 0\}$.

It should be noted that admissibility is a weak requirement as any prior belief with two or more unique states can be mapped to an admissible prior because of conditions 2 and 4.

Moreover, we consider that the voters are Bayesian thinkers. Namely, they exploit their common prior belief and their private opinion to update their probabilistic beliefs on states. The updated beliefs are further exploited by the voters to shape their reports. Specifically, the voters' posterior belief on each state is updated based on the Bayes' rule and the corresponding received signals, as follows:

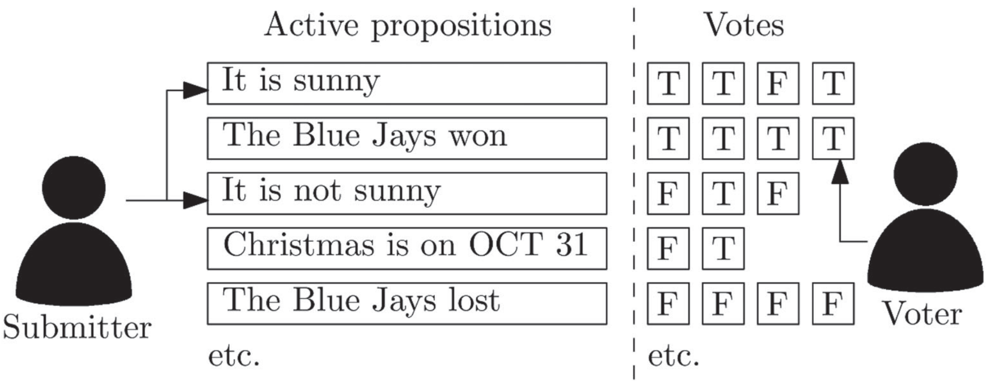

FIGURE 2 Overview of interactions among players in ASTRAEA protocols

$$\Pr(T=t|\{\mathrm{PO}_i\}) = \frac{\Pr(\mathrm{PO}_i|T=t)\cdot\Pr(T=t)}{\Pr(\mathrm{PO}_i)}$$

where

$$\Pr(\mathrm{PO}_i) = \sum_{t\in\{1\ldots m\}} \Pr(\mathrm{PO}_i|T=t)\Pr(T=t)$$

Furthermore, the private prediction PP_i is the popularity prediction on 1 (or `True`) by voter i with a private opinion PO_i as the received signal.

$$\begin{aligned}\mathrm{PP}_i &= \Pr(1|\{\mathrm{PO}_i\}) \\ &= \sum_{t\in\{1\ldots m\}} \Pr(1|T=t)\cdot\Pr(T=t|\{\mathrm{PO}_i\})\end{aligned} \tag{1}$$

3.3 | Proposition responses

A voter submits an information report $RT.\mathrm{IR}$ and a prediction report $RT.\mathrm{PR}$ to the oracle, which together define the response tuple $RT=(\mathrm{IR},\mathrm{PR})$. The information report is binary $RT.\mathrm{IR}\in\{1,0\}$, while the prediction is on the popularity of 1 (`True`) among all voters (*i.e.*, $RT.\mathrm{PR}\in[0,1]$). n voters are chosen to respond to each proposition, with $n<<|\mathcal{V}|$, as explained in Section 3.1. We denote the corresponding set of voters' responses to the specific proposition as $\mathbf{RT}=\{RT_1,RT_2,\ldots,RT_n\}$. It should be noted that the number of voters n can be variable, and we will identify a lower-bound for the value n later in our analysis in order to derive a "correct" oracle outcome with high confidence. When a voter is selected, they apply *voting strategy* function $\sigma_i((\mathrm{PO}_i,\mathrm{PP}_i))=RT_i$ to their private opionin and private prediction. An *honest* voter's strategy is $\sigma_i((\mathrm{PO}_i,\mathrm{PP}_i))=(\mathrm{PO}_i,\mathrm{PP}_i)$, while a *lazy* voter's strategy can be $\sigma_i((\mathrm{PO}_i,\mathrm{PP}_i))=(1,0.5)$ or $\sigma_i((\mathrm{PO}_i,\mathrm{PP}_i))=(0,0.5)$ for any proposition. Furthermore, we divide $\mathbf{RT}$ into subsets of responses sharing the same IR; then we have $\mathbf{RT_1}=\{RT\in\mathbf{RT}:RT.\mathrm{IR}=1\}$ and $\mathbf{RT_0}=\{RT\in\mathbf{RT}:RT.\mathrm{IR}=0\}$. Finally, we define two tuples of random variables: let $\Gamma=RT$ denote the private belief tuple of a voter selected randomly on the proposition, and $A=RT$ denote the answer tuple of a random voter.

3.4 | Correct oracle outcome

Towards determining the "correct" answer of a decentralized oracle to a proposition, we employ the concept of the Most Probable Private Opinion (MPPO).[8] It should be noted that for subjective propositions, as the ones examined in this article, there is not always available an objective truth. Thus, the "correct" answer of the oracle should be derived by the voters' opinions in a sophisticated manner. The Most Probable Private Opinion (MPPO) is defined as the most likely private belief of a randomly selected voter on a particular proposition.

$$\mathrm{MPPO}\triangleq\begin{cases}1 & ,\Pr(\Gamma.\mathrm{IR}=1)>0.5\\ 0 & ,\Pr(\Gamma.\mathrm{IR}=1)<0.5\\ \emptyset & ,Pr(\Gamma.\mathrm{IR}=1)=0.5\end{cases}$$

The oracle answer is correct if it is equal to *MPPO*.

For the rest of the paper, we assume there is a defined correct answer. The probability of a random voter providing an answer with IR equal to MPPO is defined as follows.

$$c\triangleq Pr(A.\mathrm{IR}=\mathrm{MPPO})$$

3.5 | Prediction report means and scoring rule

Let $\bar{P}_{-i,1}$ denote the geometric mean of the voters' prediction reports *RT.PR* in $\mathbf{RT_1}$ excluding voter *i*, where the corresponding set of responses is denoted as $\mathbf{RT_{\text{-}i,1}} = \mathbf{RT_1} - \{RT_i\}$. For notation convenience, we denote $\bar{P}_{-i,1} = G(\mathbf{RT_{\text{-}i,1}})$, where *G* is the geometric mean. Similarly, $\bar{P}_{-i,0} = G(\mathbf{RT_{-i,0}})$ represents the geometric mean of the prediction reports in $\mathbf{RT_0}$, excluding voter *i*, where the corresponding set of responses is denoted as $\mathbf{RT_{-i,0}} = \mathbf{RT_0} - \{RT_i\}$. In the rest of our analysis, we use the prediction means as references to compare the chosen voter's prediction report with and to determine the corresponding reward that should be provided to the voter.

A *binary scoring rule* is a scoring rule to assign a score to an prediction $q \in [0, 1]$ according to a binary outcome $w \in \{0, 1\}$. It is *strictly proper* if the participant uniquely maximizes their expected score by reporting their prediction truthfully. The binary quadratic scoring rule (R_q)[29] is a strictly proper scoring rule, and is defined as follows:

$$R_q(q,w) = \begin{cases} 2q - q^2 & ,w = 1 \\ 1 - q^2 & ,w = 0 \end{cases}$$

3.6 | Voting weight and reward share

To discourage Sybil attack, the protocol associates the voting weight and reward share of a response submitted by voter *i* with the stake s_i. The voting weight of the response by voter *i* is denoted as $f(s_i)$, whereas the normalized weight is $\frac{f(s_i)}{\sum_{i' \in \{1,\ldots,n\}} f(s_{i'})}$. The normalized weight reflects how impactful the response is to the final oracle outcome. In the previous protocols,[7,8] voting weight is linear to the submitted stake $(f(s_i) = s_i)$. The reward share of the response by voter *i* is denoted as $g(s_i)$. Similarly, in the previous protocols,[7,8] reward share is linear to the submitted stake $(g(s_i) = s_i)$. The reward share, normalized by the total share, determines the fraction of total reward that a voter receives when eligible.

4 | COLLECTIVE TRUTH EXTRACTION PROTOCOL

This section introduces the novel decentralized oracle protocol. Firstly, we present the step by step description of the protocol, then a detailed analysis demonstrating the incentive compatibility of the protocol.

4.1 | Description

We detail the interaction between the players and the oracle system in the subsection. Similar to the paired-question protocol, any participant can submit a proposition to become a submitter as described in 4.1.1, and any participant can stake and vote on an active proposition to become a voter as described in 4.1.2. After the proposition is closed, the oracle determines the outcome as shown in 4.1.3. Depending on the oracle outcome, a score is calculated for each vote and a reward is given as shown in 4.1.4.

4.1.1 | Proposition creation

A submitter can create a query by sending a transaction to the oracle executor. The transaction contains components as specified in Table 1. A bond is deposited for quality control purposes so that the submitters are penalized if the proposition pair is badly worded or does not have opposite outcomes. The bounty serves as the potential reward to the voters. The duration determines the time period that the voters can provide answers to the two propositions. When the propositions are closed, assume that the outcome of the proposition-pair is o and o', then:

- If $o = \neg o'$, the bounty is used to reward the voter, and the bond is refunded to the submitter, or
- If $o = o'$, the bounty is refunded to the submitter, and the bond is equally distributed to the other active propositions.

4.1.2 | Voting

A voter goes through the procedures described in Figure 3 to vote on a proposition. (i) the voter i submits a monetary stake, $s_i \in [s_{min}, s_{max}]$, to the oracle executor. (ii) the voter receives a random proposition by the oracle. (iii) based on their prior belief, voter i generates private opinion PO_i and private prediction PP_i. Then, voter i generate a response tuple $RT_i = \sigma_i(\mathrm{PO}_i, \mathrm{PP}_i)$, where σ_i is their voting strategy function. (iv) the voter returns a sealed version of RT_i to the oracle. (v-vii) when the proposition becomes inactive, the voter reveals their response, and the oracle finds oracle outcomes and distributes rewards.

The random proposition assignment reduces the risk of collusion by increasing its corresponding cost. Suppose an entity controlling n voter intends to acquire the majority of responses to a proposition. We denote the number of response selected for the proposition to be m. It is notable that m is a small fraction over the complete set of voters $\mathcal{V}$. The selection of the m voters can be seen as a series of m Bernoulli trials where the selected voter is either colluded voter (with a chance of $\frac{n}{|\mathcal{V}|}$) or non-colluded (with a chance of $1 - \frac{n}{|\mathcal{V}|}$). Evidently, to control more than $\frac{m}{2}$ randomly chosen voters, the attacker must control a significant portion of total voters. For example, if an entity controls 50 voters ($n = 50$) of a total of 100 voters ($|\mathcal{V}| = 100$). For a proposition with 10 voters chosen ($m = 10$), the chance of the entity controlling at least half of the responses drops to less than 40%.

4.1.3 | Outcome Determination

After the proposition expires, the oracle determines the outcome by finding the weighted majority of the information report. Each response tuple has a weight calculated based on the associated stake and a sub-linear function. The oracle outcome o is therefore determined as follows:

$$o = \begin{cases} 1, & \sum\limits_{i \in \{i:RT_i \in \mathbf{RT}_1\}} f(s_i) > \sum\limits_{i' \in \{i':RT_{i'} \in \mathbf{RT}_0\}} f(s_{i'})) \\ 0, & \sum\limits_{i \in \{i:RT_i \in \mathbf{RT}_1\}} f(s_i) < \sum\limits_{i' \in \{i':RT_{i'} \in \mathbf{RT}_0\}} f(s_{i'})) \\ \emptyset, & otherwise \end{cases}$$

TABLE 1 Components of a proposition query

Component	Description
Proposition-pair	The query the submitter is interested in. The pair should have potentially antithetic answers.
Bond	Monetary deposit. It is refunded if the proposition-pair has opposite oracle outcomes.
Bounty	Monetary deposit. It is awarded to the voters if the proposition-pair has opposite oracle outcomes.
Duration	A period of time. During this time, the proposition-pair remains active.

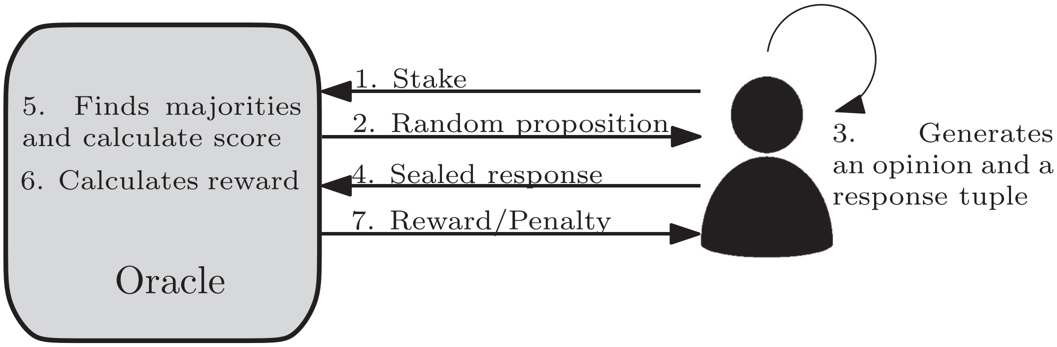

FIGURE 3 Overview of interaction between a voter and the oracle in the proposed protocol

For the sake of simplicity, we assume there always is a defined outcome for the rest of the paper. The sub-linear function f to be used is left to the actual implementation. We provide some guidelines on choosing the function in Section 5.1.

Then, the oracle checks if the paired proposition concludes to an antithetic outcome. If the same outcome is reached for the proposition pair, the submitter loses their bond, while the voters' stakes are refunded. Otherwise, the submitter regains their invested bond, and the voters' rewards are determined as presented in the following subsection.

4.1.4 | Reward determination

As the first step of reward determination, the oracle assigns a score u_i to each response $RT_i \in \mathbf{RT}$. The score composes of a *prediction score* $u_{i,\ PR}$ and an *information score* $u_{i,\ IR}$, and $u_i = u_{i,\mathrm{IR}} + u_{i,\mathrm{PR}}$.

The prediction score is determined by applying the quadratic scoring rule on the prediction report and the information report of a random reference response tuple $RT_{i'}$, as following:

$$u_{i,PR} = R_q(RT_i.\mathrm{PR}, RT_{i'}.\mathrm{IR})$$

The information score is calculated by subtracting 1 with the mean squared error between the prediction report and the mean prediction report of all reports sharing the same information report:

$$u_{i,IR} \triangleq \begin{cases} 1 - (\bar{P}_{-i,1} - RT_i.\mathrm{PR})^2 & \text{if } RT_i.\mathrm{IR} = 1 \\ 1 - (\bar{P}_{-i,0} - RT_i.\mathrm{PR})^2 & \text{if } RT_i.\mathrm{IR} = 0 \end{cases}$$

Each response tuple is ranked among the others based on it associated score u_i. Rewards are provided to the voters who achieved high-ranked scores, while the portion of voters being rewarded is left open to the implementation. A guideline for setting those system parameters can be found in Section 5.1. One of the key characteristics of the proposed approach is that the honest voters are rewarded regardless of whether they are in agreement or disagreement with the outcome. A super-linear function is used to distribute the bounty of the proposition to the voters proportionally to their invested stake towards encouraging their rational behavior.

The set of voters who are eligible to receive a reward is denoted by R. Each voter i, $i \in R$ receives a normalized reward share, $\frac{g(s_i)}{\sum_{i' \in R} g(s_{i'})}$, of the total reward. Recall that the bounty available to the paired-propositions is B, the actual reward for voter i, assuming they have submitted response for only one of the paired questions, is $\frac{g(s_i)}{2\sum_{i' \in R} g(s_{i'})} B$.

4.2 | Protocol analysis

This subsection provides a thorough analysis on the scoring rule of the protocol. We start with a formal proof on the Bayes-Nash incentive compatibility of the scoring rule, following with a demonstration of the robustness against dishonest strategies, and an analysis on the expected outcome.

4.2.1 | A Bayes-Nash incentive compatible scoring rule

In this section, a detailed theoretical analysis is provided to show that the expected score is maximized only when a voter reports honestly to the propositions. This is based on the assumption that they believe that the rest of the voters are honest. Given that the voters are rewarded based on their expected score, a rational voter aims at maximizing its score in order to have a higher chance to be in the high-score listed voters, thus, to receive a corresponding reward.

In the following theoretical analysis, we provide the detailed proofs to show that the proposed protocol is strictly Bayes-Nash incentive compatible given an admissible prior. Both Proposition 1 and 2 are known results from prior literature.[24,29] We present proofs in the context of this work so as to complete the underlying intuition.

Proposition 1. Strict Properness of Quadratic Scoring Rule.[29] Let $q \in [0, 1]$ be the private prediction on the outcome of a binary event $e \in \{0, 1\}$ so that $q = \Pr(e=1)$. Suppose the voter is rational, and the score is calculated based on a quadratic scoring rule. The voter uniquely maximizes their expected score by providing a prediction report $p = q$.

Proof. The expected score of having an private prediction of q and reporting any pr is $\mathbb{E}[u_p] = q(2p - p^2) + (1-q)(1-p^2)$. While reporting q has an expected score of $\mathbb{E}[u_q] = q(2q - q^2) + (1-q)(1-q^2)$. The expected loss is $\mathbb{E}[u_q] - \mathbb{E}[u_p] = (q-p)^2 \geq 0$. Hence, the expected score is maximized when $p = q$ with an expected loss of 0.

Proposition 2. Ranges of Posterior Belief.[24] It holds that $1 > \Pr(1|\{1\}) > \Pr(1) > \Pr(1|\{0\}) > 0$ for all admissible priors.

Proof. As the signal beliefs conditional on state are fully mixed, that is $0 < \Pr(po|T=t) < 1 \; \forall t \in \{1 \ldots m\}$ and $po \in \{1, 0\}$, it is easy to show that $1 > \Pr(1|\{1\}) > 0$ and $1 > \Pr(1|\{0\}) > 0$. Expanding the expression above, the following conditions must hold:

$$\Pr(1|\{1\}) = \sum_{t \in \{1 \ldots m\}} \Pr(1|T=t) \cdot \Pr(T=t|\{1\}) > \Pr(1) > \Pr(1|\{0\}) = \sum_{t \in \{1 \ldots m\}} \Pr(1|T=t) \cdot \Pr(T=t|\{0\}) \tag{2}$$

where $\Pr(1|\{1\})$ is the posterior of a voter with private opinion of 1, $\Pr(1)$ is the probability of 1 being observed by any voters, and $\Pr(1|\{1\})$ is the posterior of a voter with private opinion of 0.

Recall the assortative property of the prior belief: $\Pr(1|T=m) > \Pr(1|T=m-1) > \ldots > \Pr(1|T=1)$ and $\Pr(0|T=m) < \Pr(0|T=m-1) < \ldots < \Pr(0|T=1)$. Then, for $1 > \Pr(1|\{1\}) > \Pr(1) > \Pr(1|\{0\}) > 0$, it is sufficient that

$$\sum_{t \in \{1 \ldots t'\}} \Pr(T=t|\{1\}) < \sum_{t \in \{1 \ldots t'\}} \Pr(T=t) < \sum_{t \in \{1 \ldots t'\}} \Pr(T=t|\{0\}) \tag{3}$$

The intuition is that, assuming toward a contradiction that relationship among the signal posteriors in equation 3 are without the strict inequality, then we can strictly decrease $\Pr(1|\{1\})$ by decreasing $\Pr(T=t|\{1\})$ while increasing $\Pr(T=t-1|\{1\})$ for some $t > 1$ according to the assortative property and without breaking equation 3. Eventually, we will arrive at $\Pr(1|\{1\}) = \Pr(1|\{0\})$, reaching an contradiction. Therefore, equation 3 must hold.

Because $\Pr(T=t|\{1\}) \propto \Pr(1|T=t)\Pr(T=t)$ and $\Pr(T=t|\{0\}) \propto \Pr(0|T=t)\Pr(T=t)$, in order to for equation 3 to hold, it is sufficient to show that

$$\frac{\sum_{t \in \{1 \ldots t'\}} \Pr(1|T=t)}{\sum_{t \in \{1 \ldots m\}} \Pr(1|T=t)} < \frac{t'}{m} < \frac{\sum_{t \in \{1 \ldots t'\}} \Pr(0|T=t)}{\sum_{t \in \{1 \ldots m\}} \Pr(0|T=t)} \tag{4}$$

for all $t' \in \{1, m-1\}$.

Because of the assorative property, $\sum_{t \in \{1 \ldots m\}} \Pr(1|T=t) > m\Pr(1|T=1)$ and $\sum_{t \in \{1 \ldots m\}} \Pr(0|T=t) < m\Pr(0|T=1)$.

$$\frac{\Pr(1|T=1)}{\sum_{t \in \{1 \ldots m\}} \Pr(1|T=t)} < \frac{1}{m} < \frac{\Pr(0|T=1)}{\sum_{t \in \{1 \ldots m\}} \Pr(0|T=t)}$$

And we have

$$\frac{\sum_{t \in \{1 \ldots t'\}} \Pr(1|T=t)}{\sum_{t \in \{1 \ldots m\}} \Pr(1|T=t)} < \frac{t'\Pr(1|T=1)}{\sum_{t \in \{1 \ldots m\}} \Pr(1|T=t)} < \frac{t'}{m} < \frac{t'\Pr(0|T=1)}{\sum_{t \in \{1 \ldots m\}} \Pr(0|T=t)} < \frac{\sum_{t \in \{1 \ldots t'\}} \Pr(0|T=t)}{\sum_{t \in \{1 \ldots m\}} \Pr(0|T=t)}$$

which supports equation 4. Therefore, $1 > \Pr(1|\{1\}) > \Pr(1) > \Pr(1|\{0\}) > 0$ holds for all admissible priors.

Proposition 3. Strict Properness of Prediction Score. A voter uniquely maximizes their expected information score by truthfully reporting their private prediction if all other voters are honest.

Proof. The expected probability that a random reference response tuple containing an information report of 1, from the perspective of voter i, is $\Pr(RT_{i'}.\mathrm{IR}=1)=\Pr(1|\mathrm{PO}_i)$. According to Proposition 1, as prediction score is calculated with a quadratic scoring rule, the voter uniquely maximizes the expected score by reporting $RT_i.\mathrm{PR}=\Pr(1|\mathrm{PO}_i)$. This agrees with the strategy of a honest voter, hence honest voting maximizes the expected prediction score.

Proposition 4. Strict Properness of Information Score. A voter uniquely maximizes their expected information score by truthfully reporting their private opinion if all other voters are honest.

Proof. Given all other voters are honest and following Proposition 2, a voter would expect $(\mathrm{PP}_i-\bar{P}_{-n,\mathrm{PO}_i})^2 < (\mathrm{PP}_i-\bar{P}_{-n,\neg\mathrm{PO}_i})^2$. Therefore, reporting $RT_i.\mathrm{IR}=\mathrm{PO}_i$ yields strictly higher information score from the perspective of the voter.

Theorem 1. *The proposed scoring rule is Bayes-Nash incentive compatible.*

Proof. By Proposition 3, the pure strategy of reporting the private prediction uniquely maximizes the expected prediction score. By Proposition 4, honest reporting of the private opinion uniquely maximizes the expected information score. It follows that the strategy of honest reporting maximizes the overall expected score. Therefore, the proposed scoring rule is Bayes-Nash incentive compatible.

4.2.2 | Expected score and robustness against dishonesty

In this subsubsection, we demonstrate the expected scores of voters under specific settings. To demonstrate the incentive compatibility of the scoring rule, we use a system behavior with a 2-state voter belief system, *i.e.*, $m=2$. The initial probabilities over the two states are 0.5, *i.e.*, $\Pr(T=1)=\Pr(T=2)=0.5$. We also define the set of possible strategies that can be employed by a voter in table 2 where σ_i is the strategy function of voter i.

Since the score is directly related to the rewards a voter is able to receive (*i.e.*, the higher the score is, the more likely a voter can receive a reward based on the rewarding schemes), a rational voter will seek to maximize the expected score. We demonstrate the expected information score calculation through the following example: suppose all voters in the system are honest voters. Voter i has a private opinion of `True` (or 1) and hence a posterior private prediction of $\mathrm{PP}_i=\Pr(1|\{1\})$. They would expect the probability of another honest voter j reporting $RT_j.\mathrm{IR}=1$ with $\Pr(1|\{1\})$, and $RT_j.\mathrm{IR}=0$ with $\Pr(0|\{1\})=1-\Pr(1|\{1\})$. Therefore, the expected prediction score for voter i is:

$$\mathbb{E}[u_{i,\mathrm{PR}}]=\Pr(1|\{1\})(2RT_i.\mathrm{PR}-RT_i.\mathrm{PR}^2)+(1-\Pr(1|\{1\}))(1-RT_i.\mathrm{PR}^2)$$

And the expected information score for voter i is:

$$\mathbb{E}[u_{i,\mathrm{IR}}]=\begin{cases}1-(\Pr(1|\{1\})-\Pr(1|\{1\}))^2 & ,RT_i.\mathrm{IR}=1\\ 1-(\Pr(1|\{0\})-\Pr(1|\{1\}))^2 & ,RT_i.\mathrm{IR}=0\end{cases}$$

For the purpose of demonstration, assume that the two states belief system has the following distribution: $\Pr(1|T=1)=0.25$ and $\Pr(1|T=1)=0.75$. According to equation 1, $\mathrm{PP}_i=\Pr(1|\{1\})=0.625$. Figure 4 (a) visualizes the expected score by voter i under different combination of prediction report and information report. We present the preferred information report given the corresponding prediction report as a darkened curve. If the voter lies about their opinion and reports $RT_i.\mathrm{IR}=0$, the expected score is maximized when they also report $RT_i.\mathrm{PR}=0.5$. It is notable that the expected score is strictly the highest when the voter reports both their opinion and prediction honestly, which is when

TABLE 2 Summary of c_j values for pure strategies in response to honest voting

Name	Strategy Function	c_j
Honest	$\sigma_i((\mathrm{PO_i}, \mathrm{PP_i})) = (\mathrm{PO_i}, \mathrm{PP_i})$	>0.5
Lying	$\sigma_i((\mathrm{PO_i}, \mathrm{PP_i})) = (\neg\mathrm{PO_i}, max_{p \in [0,1]} \mathbb{E}[u_i])$	<0.5
Lazy, always `True`	$\sigma_i((\mathrm{PO_i}, \mathrm{PP_i})) = (1, 0.5)$	0.5
Lazy, always `False`	$\sigma_i((\mathrm{PO_i}, \mathrm{PP_i})) = (0, 0.5)$	0.5

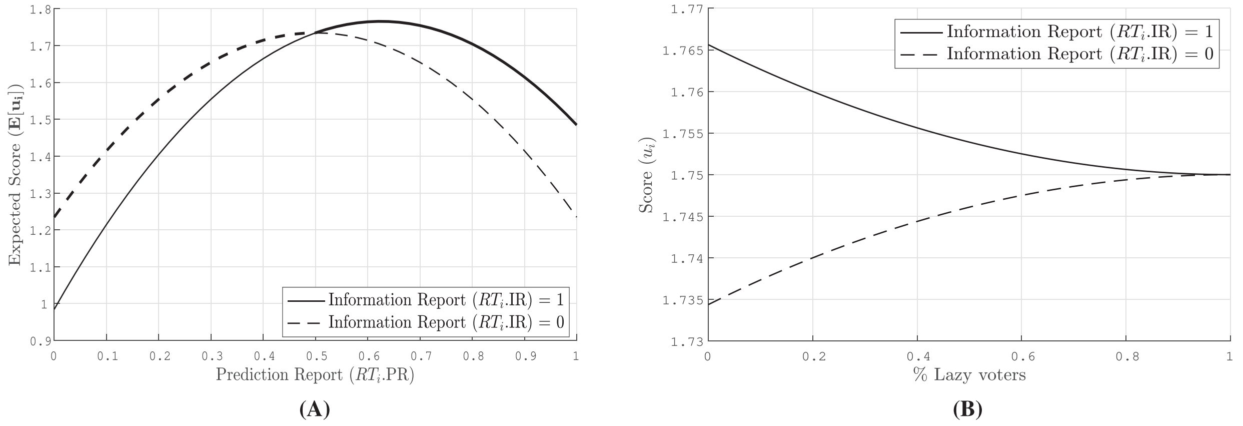

(A)

(B)

FIGURE 4 Expected score with various report values and the expected maximum score in relationship to lazy voter percentage for voter with $\mathrm{PP}_i = 0.625$

RT_i.IR$=1$ and RT_i.PR$=0.625$. Similarly, consider a more controversial proposition with $\Pr(1|T=1)=0.45$ and $\Pr(1|T=1)=0.55$, of which result is shown in Figure 5. Because of the similarity in the probabilistic state, the differentiation of either information report is not as significant. In contrary, if the majority is likely to be agreed by a greater majority ($\Pr(1|T=1)=0.10$ and $\Pr(1|T=1)=0.90$) as in Figure 6, the gain of switching from lying to honest reporting is evidently more significant.

We now relax the assumption that not all voters are honest. Suppose a portion of the voters employ the lazy strategies as shown in in table 2, and half of those voters always report RT_i.IR$=1$ while the other half always report RT_i.IR$=0$. Figure 4(b), 5(b) and 6(b) visualizes the maximized expected score by voter i under the various belief systems when there are some percentage of voters being lazy. Denote that the portion of lazy voter in the system to be μ, voter i adjusts their private prediction to $\mathrm{PP}_i(1-\mu)+0.5\mu$. As the percentage of lazy voters increases in the system, the expected score of honest reporting while the one of lying increases. However, honest reporting is strictly preferred as long as the percentage of lazy voters is less than 100%.

4.2.3 | Expected outcome

Towards determining the outcome of the proposed protocol, we adopt the weighted majority-based voting of the submitted information report. Let us consider a proposition for which n honest voters have provided an answer and a probability c that a randomly selected response agrees with MPPO. Let us also denote the probability of correct oracle outcome as P_{Corr}, which captures the probability that a majority of voters believes in MPPO.

In this subsection, we assume all voters submit the same amount of stake for the sake of simplicity. Given that generation of private opinions by any voter i and i' ($i \neq i'$) is independent, then generation of private opinion is simply a series of n Bernoulli trials, each with probability c to agree with MPPO. Hence, P_{Corr} can be calculated as follows:

$$P_{\mathrm{Corr}} = 1 - B\left(\left\lfloor\frac{n}{2}\right\rfloor, n, c\right)$$

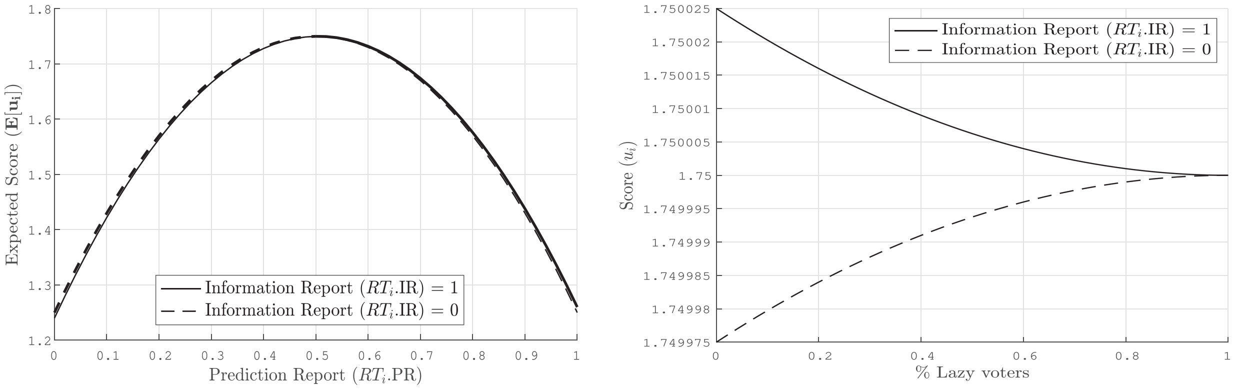

(A)

(B)

FIGURE 5 Expected score with various report values and the expected maximum score in relationship to lazy voter percentage for voter with $\text{PP}_i = 0.505$

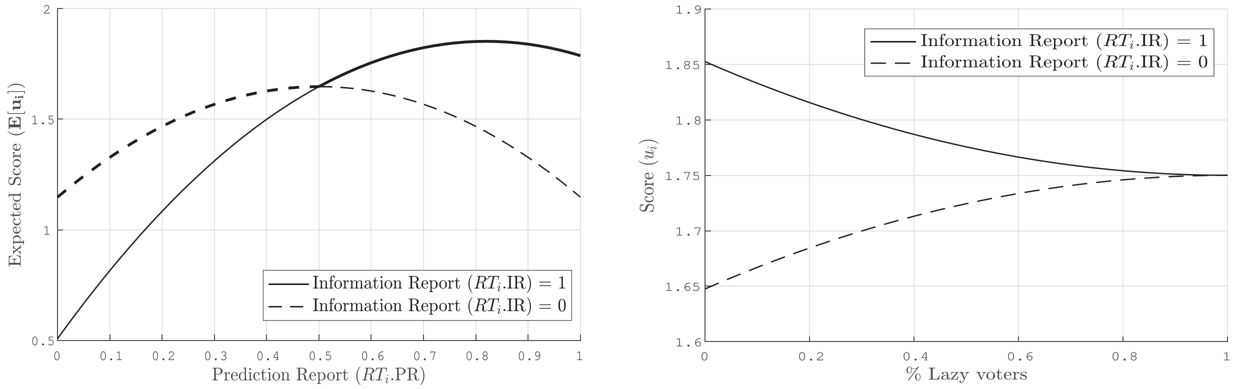

(A)

(B)

FIGURE 6 Expected score with various report values and the expected maximum score in relationship to lazy voter percentage for voter with $\text{PP}_i = 0.820$

where $B\left(\left\lfloor\frac{n}{2}\right\rfloor, n, c\right)$ is the cumulative binomial density function.

Given all honest voters, the probability of correct output associated with different c is shown in Figure 7. If only a few voters are chosen, only propositions with widely accepted answers are likely to come out correctly. On the other hand, even if a proposition is highly contentious (with c close to 0.5), the oracle will agree with MPPO with high probability provided there are enough voters. This is the same as any other oracle protocol based on majority voting while the proposed protocol results in this outcome without a rewarding scheme that relies on the most popular outcome, thereby avoiding herding effects. A reasonable minimum number of voters is 30 as a relatively controversial proposition (with $c = 0.6$) would have a chance of around 70% to come out correctly.

5 | IMPROVING RESISTANCE TO A SYBIL ADVERSARY

As a mechanism against Sybil attack, the implementation of the proposed scaling rule is critical to the viability of the proposed protocol. This section suggests a possible implementation of the scaling rule in voting-based oracles. We demonstrate its efficacy in discouraging Sybil attacks with a specific example, then analyze the advantages of the proposed system.

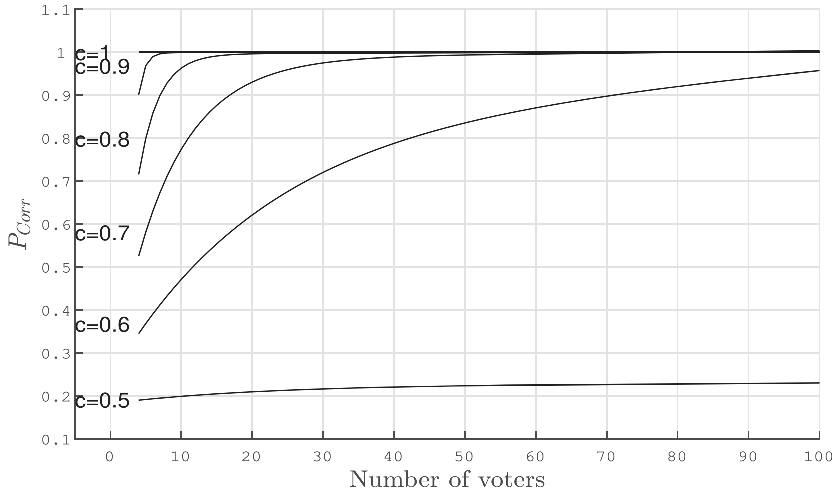

FIGURE 7 Probability of correctness as a function of n and c

5.1 | Choice of scaling functions

The scaling rule composes of two stake scaling functions as mentioned in Section 4.1.3 and 4.1.4. This distinguishes our proposal from other staked-voting systems since other staked-voting systems usually scale voting power and reward portion linearly to the associated stake.

Section 4.1.3 mentioned that the protocol scales a voter's voting weight with a sub-linear function. In the following section, we denote such a function with $f(s)$. We also denote the super-linear function applied during reward allocation in Section 4.1.4 with $g(s)$. Without loss of generosity, we let a response with minimum stake with unit voting power and unit reward share, that is $f(s_{min}) = g(s_{min}) = 1$. The paired stake scaling rule encourages voters to stake on a single identity as reward share grows super-linearly in stake, thus improving Sybil resistance. The scaling rule also prevents the formation of voting pools or a single entity from having significant voting power as voting power grows sub-linearly in stake. During implementation, one can design the pair stake scaling rule to adjust the preferred stake amount on a single identity.

For the rest of the section, we consider a group of sub-linear function with the form $f(s) = \alpha\sqrt{s} + (1-\alpha)s$ with $\alpha \in (0,1]$. We also consider a group of super-linear function with the form $g(s) = \beta s^2 + (1-\beta)s$ where $\beta \in (0,1]$. If $\alpha = 0$ and $\beta = 0$, voting power and reward allocation share of a report increase linearly in the associated stake, similar to other stake-voting protocols. By adjusting α and β, one can modify the growth rate of voting power and reward share. To perform an analysis, assume there exist an oracle system with a fixed amount of honest voters, each of whom stakes s_{min} to received proposition assignments. For the purpose of analysis, let $x = 1$, which means all voters in the system are eligible to receive a reward. For an additional voter i who stakes s_i to a single identity, i will expect their normalized voting weight and normalized reward share as in Figure 8 and 9 respectively. Additionally, we consider only a voter with strictly less than 50% of the total stake. This is because a voter with more than 50% stake can effectively manipulate the outcome in any other stake voting systems. Hence, we make an assumption that s_{max} is set to a value that an entity cannot stake an unreasonable high amount on a response, or enough entities should be selected for each proposition.

5.2 | Expected utility and Sybil resistance

This subsection showcases the improvement against Sybil attack by utilizing the paired stake scaling rule under various circumstances. Recall that the voter reward is based on the calculated score, and a voter receives a reward when:

1. there exists a weighted majority outcome,
2. the proposition pair, where the proposition belongs, has opposite oracle outcomes, and
3. their report receives a score that is ranked in the top x portion of all submitted reports.

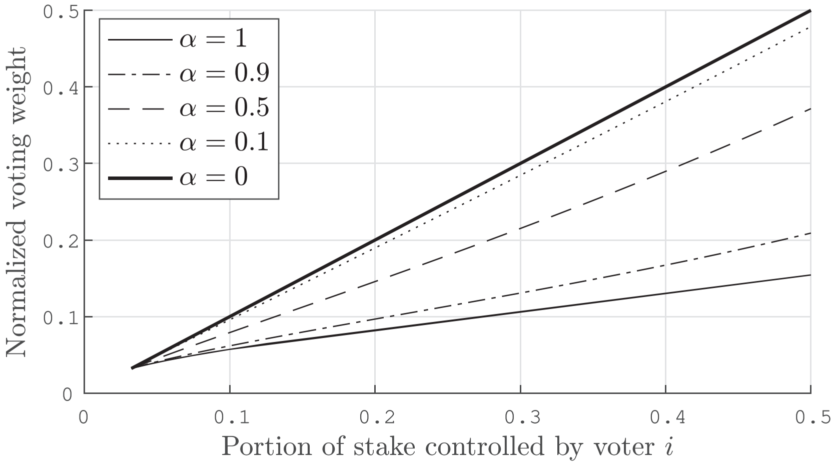

FIGURE 8 Normalized voting weight varies by portion of stake controlled with different suggested scaling functions

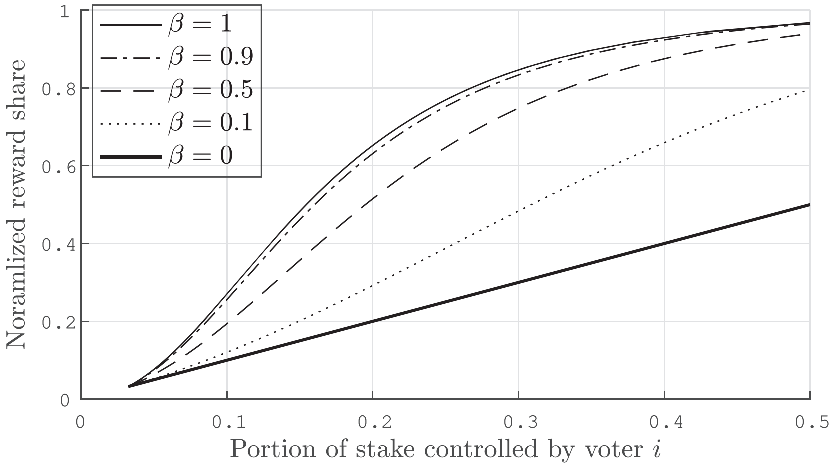

FIGURE 9 Normalized reward share varies by portion of stake controlled with different suggested scaling functions

We assume that both 1 and 2 hold in the following analysis, and that all submitted reports are equally likely to be in the top x portion. Additionally, it is worth mentioning that only a voter with minority private opinion would intend to perform such an attack in hope of biasing the oracle outcome.

We first define some common parameters for the analysis: we assume that there are 30 honest voters chosen for the interested proposition, which is a reasonable number of voters according to Section 4.2.3. Each of these 30 voters has stake $s_{min} = 1$ to the system. Additionally, there is a potentially adversary voter i with available stake of s_i. The payoff of voter i is result of (i) outcome utility, (ii) oracle reward, and (iii) voting cost. For (i), i gains a utility of u_h if the oracle outcome agrees with their private opinion ($o = \mathrm{PO}_i$), or u_l if the oracle outcome opposes their private opinion ($o = \neg\mathrm{PO}_i$). We assume that $u_h > u_l$. (ii), the total available reward for responses to the proposition is $\frac{B}{2}$. With a reward portion of $x \in (0, 1]$ and if voter i does not perform an attack, responses ranked in the first $\lfloor 31x \rfloor$ positions, according to the calculated score, are eligible for an reward. In contrast, if voter i instead vote with s_i identities each with a stack of $s_{min} = 1$, the top $\lfloor (30 + s_i)x \rfloor$ are awarded. Furthermore, if there is a tie in the score, we assume the awarded reports are randomly chosen so that the reward portion always holds. Last but not least, for (iii), we assume that submitting the first response incurs a cost of K and any additional response has a variable cost of k. K describes the cost for generating the private opinion and response tuple as well as the transaction fee for submitting the first response tuple. k includes the transaction fees for an additional report and the cost of obtaining a pseudonymous identity to vote on the specific proposition. We argue that $k > K$: to submit an additional response, the adversarial voter can either acquire an entity which as been assigned the proposition, or stake on a sufficient amount of identities before the proposition assignment. Assuming the cost of generating a private opinion is negligible, the cost of acquiring an additional identity is much higher than creating the first one.

Consider the setting where $u_h = 50$, $u_l = 0$, $\frac{B}{2} = 10$, $x = 1$, $K = 0.1$, $k = 0.3$ and $c = 0.9$. It is notable that $c = 0.9$, which means that the voter i is against a mostly agreed answer (*i.e.*, by 90% of the voters). As shown in Figure 10, despite the fact that biasing the oracle outcome successfully awards the voter more than the total available oracle reward ($u_h = 50 > 10 = \frac{B}{2}$), the voter is only incentivized to perform a Sybil attack when they have more than $\frac{21}{21+30} = 41\%$ of the total stake. This holds even with $\alpha = 0.9$ (voting power scaling slowly), and $\beta = 0.1$ (reward portion grows closely to linearly). We demonstrate an alternative scenario where the proposed protocol fails to discourage Sybil attacks in Figure 11. Keeping all other parameters constant, assume that $c = 0.55$. In other words, voter i is against the majority of a relatively controversial proposition. Most of the suggested scaling function pairs are no longer incentive compatible since the expected payoff of a Sybil attack outperforms the payoff of honest voting. This is because the probability of a successful attack against the majority in a contentious proposition is higher. Therefore, with $u_h = 50 > 10 = \frac{B}{2}$, the voter would see performing an attack beneficial. To fix this, the proposition submitter can submit a higher bounty. Figure 12 demonstrates the same controversial proposition but with $u_h = \frac{B}{2} = 10$. Because the total reward is comparable to utility from a preferred oracle outcome, voter i does not see performing an attack preferable under most of the function pairs. In the scenario described above parameter β affects the shape more noticeably than α. A higher β leads to higher expected payoff for non-attacker if the all voters are awarded.

It is notable that adversary expect an oracle reward to all the pseudonymous identities with $x = 1$. Relaxing the reward portion, so that $x < 1$, leads to lower expected reward from an attacker in general. With a single identity,

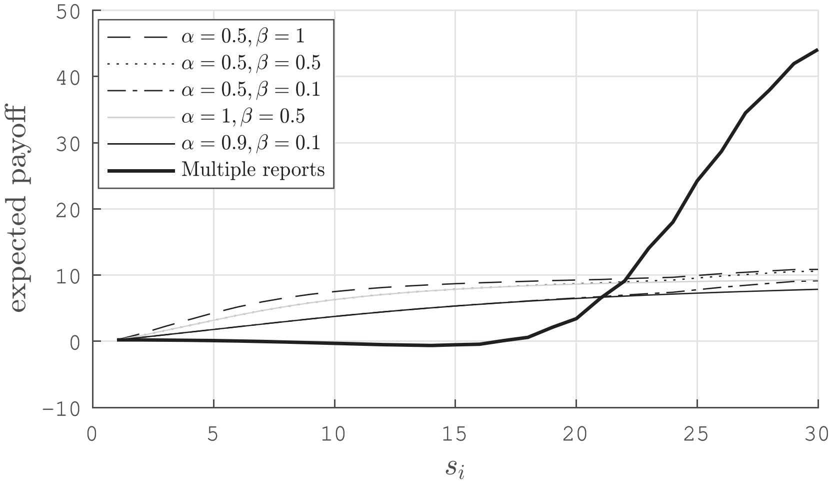

FIGURE 10 The expected payoff of voter i on a single report versus multiple reports. $c = 0.9$ and $x = 1$

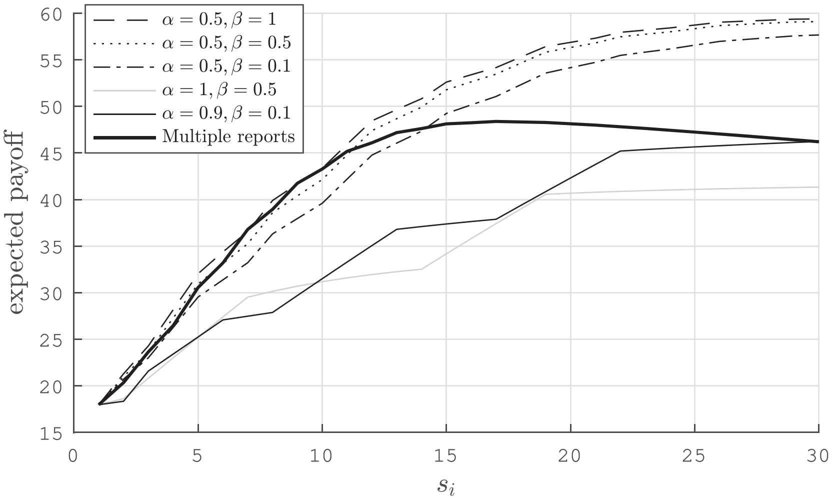

FIGURE 11 The expected payoff of voter i on a single report versus multiple reports. $c = 0.55$, $u_h = 50$ and $x = 1$

the chance of receiving a reward is x and the reward scales super-linearly in stake. On the other hand, the chance of n identities all receiving a reward is x^n. Assume that $x = 0.1$ so that the top 10% of the responses are rewarded. Most of the scaling function pairs are incentive compatible for honest voting as shown in in Figure 13. The two exceptions are when $\alpha = 0.9$ and $\alpha = 1$, the scaling function combinations fail to discourage attacks with stake less than 12 and 16 respectively. It is noticeable that α affects the shape of the payoff curve more effectively than β. Conclusively, if the reward portion is small, a larger α should be chosen to discourage attacks in a controversial proposition. In contrast, if the attacker is against a mostly agreed upon proposition as shown in Figure 14, similar level of security guarantee is achieved, even with less portion of voters being rewarded. Similar to Figure 10, the attack is only beneficial when the attacker controls more than $\frac{20}{20+30} = 40\%$ of all entities in the system.

In summary, to make the implementation incentive compatible, the total reward available to the voters should be comparable to the potential utility the potential attacker obtains from their favorite oracle outcome. Parameters, α and β, should be chosen taking the expected outcome utility, available rewards, and reward portion into consideration.

5.3 | Adversarial effects

This subsection includes additional discussion on potential adversarial behaviors. A possible adversarial behavior is to push the oracle toward ¬*MPPO* through a Sybil attack. However, this can be effectively disincentivized if the system

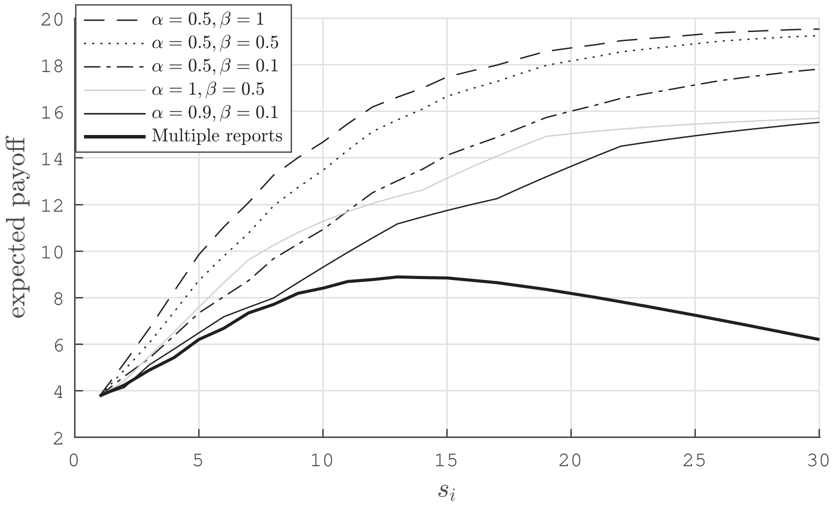

FIGURE 12 The expected payoff of voter i on a single report versus multiple reports. $c = 0.55$, $u_h = 10$ and $x = 1$

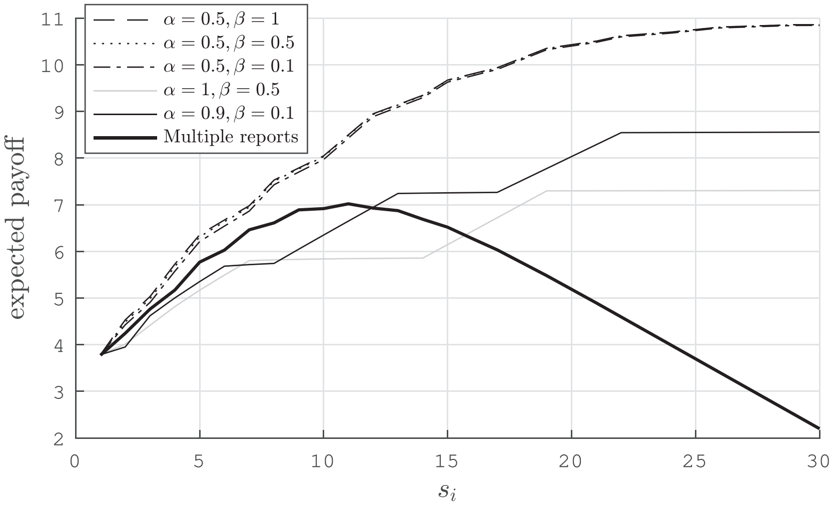

FIGURE 13 The expected payoff of voter i on a single report versus multiple reports. $c = 0.55$, $u_h = 10$ and $x = 0.1$

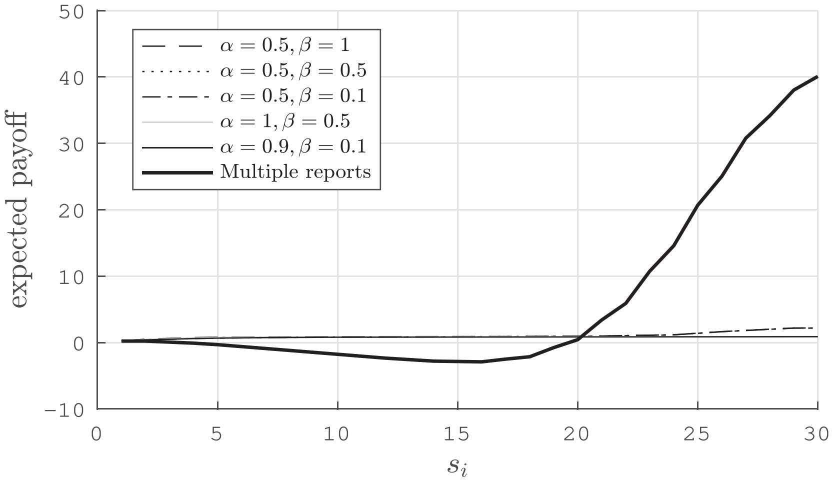

FIGURE 14 The expected payoff of voter i on a single report versus multiple reports. $c = 0.9$, $u_h = 50$ and $x = 0.1$

parameters are designed correctly as discussed in Section 5.2. Even when such an attack is reasonable, the adversarial voter has to control a considerably large portion of the total voters because of the randomization. There are various approaches to achieve randomization in a blockchain environment. For example, one can utilize the hash of the block, a random number generator such as Randao[30] on Ethereum, or a Verifiable Delay Function.[31] The chances are decreased even further by the pair-question setting, which doubles the costs of the adversary to manipulate both antithetic propositions in the pair.

Another possible adversarial behavior of an adversary is to try to achieve scores higher than any honest voters. A powerful adversary achieving such a goal can guarantee themselves a reward. However, from the point of view of an honest voter, the only possible strategy to maximize expected score is to report honestly. This mean the adversary requires additional information that is not known by honest voters. Consider the extreme scenario where the adversary knows behavior of all other voters, (*i.e.*, all the sealed response tuples). Acquiring such a knowledge incurs a significant amount of cost while the adversarial voter will have to share the reward with other top-ranking voters. Even if such an adversarial voter exists, they will not be able to secure a maximum score (for both the information score and the prediction score) because of the randomization in choosing the reference response. Furthermore, it is notable that an adversary controlling multiple identities can effectively bias the calculation of prediction means. An improvement is to, instead of using the mean of prediction reports, use the median so that it is more robust against outliers and manipulations.

5.4 | Advantages of the proposed protocol

Comparing to the existing Astraea protocols, a notable improvement of the proposed protocol is the discouragement of herding behaviors. With a belief model presented in the previous section, a rational voter will always vote for the popular answer to improve their expected reward under the existing Astraea protocols. Reporting a minority answer leads to a lower chance to be in agreement with the majority voters. However, the proposed protocol decouples voter rewards from popularity of the submitted answer, honest voters are incentivized to report their private opinion regardless of the expected popularity. One may argue that the previous protocols encourage a faster convergence to a majority answer, and hence more efficient. Nevertheless, the popularity is critical in many circumstances, including elicitation of feedback, governance decision making and opinion polls. The proposed protocol is a *much* better option as the responses provides an honest measure of how supported the outcome is, and potentially benefit future decision-making.

Moreover, the proposed protocol creates an incentive for the voters to stake on a single identity rather than multiple ones. Efficiently, with the right choice of parameters, Sybil attack is strictly unfavorable even for an adversarial voter owning a significant portion of total stake. However, an entity with a higher stake is getting more reward than a voter with the same amount of stake in a linear staking system. One potential issue is that this may lead to uneven distribution of wealth, or a phenomenon of ”the rich getting richer”. However, it can be limited by setting the maximum allowed stake, s_{max}.

6 | CONCLUSION

This paper introduces a truth-inducing decentralized oracle protocol. The proposed protocol rewards the voters, who provide their subjective data, based on the score determined with a incentive compatible scoring rule and a non-linear scaling system to discourage herding and Sybil attacks. We provide a peer prediction-based scoring rule along with a detailed proof of its Bayes-Nash incentive compatibility. The paper also includes a guideline to the selection of scaling functions as well as other the system parameters to prevent Sybil attacks. A detailed evaluation is provided to demonstrate the benefits of the proposed protocol. Part of our current and future work is the addition of a reputation system to allow the distribution of awards based on previous performance, and to incentivize sustainable participation.

ACKNOWLEDGEMENT

We would like to thank Mr. Georgios Fragkos for his initial contribution at the preliminary research work presented in the 2nd Conference on Blockchain Research & Applications for Innovative Networks and Services, BRAINS 2020.[11]

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in Professional Academic Website at http://ece-research.unm.edu/tsiropoulou/.

REFERENCES

1. Zheng Z, Xie S, Dai H, Chen X, Wang H. An overview of blockchain technology: Architecture, consensus, and future trends. In: 2017 IEEE International Congress on Big Data (BigData Congress). Honolulu, HI, USA; 2017:557-564.
2. Nakamoto S. Bitcoin: A peer-to-peer electronic cash system. Available from: http://www.bitcoin.org/bitcoin.pdf; 2009.
3. Mussenbrock C, Karpischek S. Etherisc whitepaper. Available from: https://etherisc.com/files/etherisc_whitepaper_1.01_en.pdf; 2017.
4. IBM. Blockchain for supply chain. Available from: https://www.ibm.com/blockchain/supply-chain/
5. Peterson J, Krug J, Zoltu M, Williams AK, Alexander S. Augur: a decentralized oracle and prediction market platform; 2018.
6. Greenspan G. Why many smart contract use cases are simply impossible. Available from: https://www.coindesk.com/three-smart-contract-misconceptions; 2016.
7. Adler J, Berryhill R, Veneris A, Poulos Z, Veira N, Kastania A. Astraea: A decentralized blockchain oracle. arXiv preprint arXiv: 1808.00528; 2018.
8. Merlini M, Veira N, Berryhill R, Veneris A. On public decentralized ledger oracles via a paired-question protocol. In: 2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). Seoul, Korea (South); 2019:337-344.
9. Çelen B, Kariv S. Distinguishing informational cascades from herd behavior in the laboratory. *Am Econ Rev*. 2014;94(3):484-498.
10. Douceur J. The Sybil Attack. In: Druschel P., Kaashoek F., Rowstron A., eds. *IPTPS '01 Revised Papers from the First International Workshop on Peer-to-Peer Systems*. Berlin, Heidelberg: Springer; 2002:251-260.
11. Cai Y, Fragkos G, Tsiropoulou EE, Veneris A. A truth-inducing sybil resistant decentralized blockchain oracle. In: 2020 2nd conference on blockchain research & applications for innovative networks and services (brains) IEEE. Paris, France; 2020:128-135.
12. Provable. The provable™ blockchain oracle for modern dapps. Available from: https://provable.xyz/
13. Zhang F, Cecchetti E, Croman K, Juels A, Shi E. Town crier: An authenticated data feed for smart contracts. In: Proceedings of the 2016 aCM sIGSAC conference on computer and communications security ACM; 2016:270-282.
14. Costan V, Devadas S. Intel SGX explained. *IACR Cryptology ePrint Archive*. 2016;2016(086):1-118.
15. Ellis S, Juels A, Nazarov S. Chainlink a decentralized oracle network. Available from: https://link.smartcontract.com/whitepaper; 2017.
16. Dierks T, Rescorla E. The Transport Layer Security (TLS) Protocol, Version 1.2. Available from: https://tools.ietf.org/html/rfc5246; 2008.
17. Ritzdorf H, Wust K, Gervais A, Felley G, Juels A. Tls-n: Non-repudiation over tls enabling ubiquitous content signing. In: Network and distributed system security symposium (NDSS). San Diego, California; 2018.
18. Guarnizo J, Szalachowski P. PDFS: Practical Data Feed Service for Smart Contracts. In: 24th edition of esorics; 2019.
19. Zhang F, Maram SKD, Malvai H, Goldfeder S, Juels A. DECO: Liberating Web Data Using Decentralized Oracles for TLS. In: 24th edition of esorics; 2019.
20. Adler J, Berryhill R, Veneris A, Poulos Z, Veira N, Kastania A. ASTRAEA: A Decentralized Blockchain Oracle. In: Proceedings of the 2018 IEEE conference on blockchain. Halifax, NS, Canada; 2018:1145-1152.
21. Kamiya R. Shintaku: An end-to-end-decentralized general-purpose blockchain oracle system. Available from: https://gitlab.com/shintaku-group/paper/raw/master/shintaku.pdf; 2018.
22. Miller N, Resnick P, Zeckhause R. Eliciting informative feedback: The peer prediction method. *Management Science*. 2005;51(9): 1359-1373.
23. Prelec D. A bayesian truth serum for subjective data. *Science*. 2004;306(5695):462-466.
24. Witkowski J, Parkes DC. A robust bayesian truth serum for small populations. In: Proceedings of the twenty-sixth AAAI conference on artificial intelligence AAAI. Québec City, Québec, Canada; 2014:1492-1498.

25. Buterin V. Ethereum: A next-generation smart contract and decentralized application platform. https://github.com/ethereum/wiki/wiki/White-Paper; 2014.
26. Cachin C. Architecture of the hyperledger blockchain fabric. In: Workshop on Distributed Cryptocurrencies and Consensus Ledgers. Vol. 310. Chicago, Illinois, USA; 2016.
27. Pearl J. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. 2nd ed. San Mateo, CA, USA: Morgan Kaufmann; 1988.
28. Harsanyi JC. Games with incomplete information played by 'bayesian' players, part iii. the basic probability distribution of the game. *Manag Sci*. 1968;14(7):486-502.
29. Selten R. Axiomatic characterization of the quadratic scoring rule. *Exp Econ*. 1998;1:43-62.
30. randao.org. Randao: Verifiable random number generation. Available fromhttps://randao.org/whitepaper/Randao_v0.85_en.pdf; 2017.
31. Boneh D, Bonneau J, Bünz B, Fisch B. Verifiable delay functions. *IACR Cryptol ePrint Arch*. 2018;2018:601.

AUTHOR BIOGRAPHIES

Yuxi Cai received a B.A.Sc. degree in Electrical and Computer Engineering from the University of Toronto in 2019. She is currently pursuing the M.A.Sc. degree with the Department of Electrical and Computer Engineering, University of Toronto. Her current research interests include decentralized blockchain oracles, and economic incentive design.

Nafis Irtija is a Ph.D. student and research assistant in the Department of Electrical and Computer Engineering, University of New Mexico. He received his bachelor's and master's in Electrical and Electrical Engineering from the University of Dhaka, Bangladesh. His main research interests include distributed decision making, reinforcement learning, game theory, optimization, contract theory, and prospect theory.

Eirini Eleni Tsiropoulou is currently an Assistant Professor at the Department of Electrical and Computer Engineering, University of New Mexico. Her main research interests lie in the area of cyber-physical social systems and wireless heterogeneous networks, with emphasis on network modeling and optimization, resource orchestration in interdependent systems, reinforcement learning, game theory, network economics, and Internet of Things. Five of her papers received the Best Paper Award at IEEE WCNC in 2012, ADHOCNETS in 2015, IEEE/IFIP WMNC 2019, and INFOCOM 2019 by the IEEE ComSoc Technical Committee on Communications Systems Integration and Modeling, and IEEE/ACM BRAINS 2020. She was selected by the IEEE Communication Society - N2Women - as one of the top ten Rising Stars of 2017 in the communications and networking field. She received the NSF CRII Award in 2019 and the Early Career Award by the IEEE Communications Society Internet Technical Committee in 2019.

Andreas Veneris received a Diploma in Computer Engineering and Informatics from the University of Patras in 1991, an M.S. degree in Computer Science from the University of Southern California, Los Angeles in 1992 and a Ph.D. degree in Computer Science from the University of Illinois at Urbana Champaign in 1998. In 1998 he was a visiting faculty at the University of Illinois until 1999 when he joined the Department of Electrical and Computer Engineering and the Department of Computer Science at the University of Toronto where today he is a Professor. Since 2018 he is a Connaught Scholar for his contributions to blockchain technology. His research interests include CAD for debugging, verification, synthesis and test of digital circuits/systems, crypto-economics, decentralized blockchain technology, and combinatorics. He has received several teaching awards, a best paper award and a Ten Year Best Paper Retrospective Award. He is the author of one book and he holds several patents. He is a member of IEEE, ACM, AMS, AAAS, Technical Chamber of Greece, Professionals Engineers of Ontario and The Planetary Society.

How to cite this article: Cai Y, Irtija N, Tsiropoulou EE, Veneris A. Truthful Decentralized Blockchain Oracles. *Int J Network Mgmt*. 2021;e2179. https://doi.org/10.1002/nem.2179