![LEDGER ledgerjournal.org]

RESEARCH ARTICLE

Sentiment Protocol: A Decentralized Protocol Leveraging Crowdsourced Wisdom

Anton Muehlemann[*][†]

Abstract. The wisdom of the crowd is a valuable asset in today's society. It is not only important in predicting elections but also plays an essential role in marketing and the financial industry. Having a trustworthy source of opinion can make forecasts more accurate and markets predictable. Until now, a fundamental problem of surveys is the lack of incentives for participants to provide accurate information. Classical solutions like small monetary rewards or the chance of winning a prize are often not very attractive for participants. More attractive solutions, such as prediction markets, face the issue of illegality and are often unavailable. In this work, we present a solution that unites the advantages from classical polling and prediction markets via a customizable incentivization framework. Apart from predicting events, this framework can also be used to govern decentralized autonomous organizations.

1. Introduction

Sentiment, in its broadest sense, is of great interest to both politics and industry. In the USA alone, the revenue of market research and public opinion polling reached $18B in 2016 and continues to grow.[1] In spite of its size, political forecasts have failed to accurately predict two major recent political events, the US presidential elections and the decision of the UK to leave the EU.[2,3] These failures may come as a surprise as, with the emergence of social networks and generally more interactive websites, it has never been easier to source the opinion of the crowd.

In the financial industry, rating agencies are highly-paid providers of sentiment on a wide range of investment vehicles. However, if one takes for example a popular US stock such as those of the company Tesla (TSLA), on Oct 13, 2017, out of 25 analysts, 35% recommended *buy*, another 35% recommended *hold* and the remaining 30% recommended *sell*,[4] therefore making the prediction no better than a simple guess.[5] Another well-documented fact is that most actively managed funds fail to beat the market, showing again that so-called expert opinions are not as valuable as they may seem.[6,7]

A common feature of the above cases is that repercussions and rewards for providing inaccurate and accurate sentiment are extremely limited, for both individuals (*i.e.*, experts) and the crowd. This is due in both cases to the difficulty of objectively evaluating their respective performances. For experts, reasons for this difficulty include the phenomenon of survival bias

───────────────
[*] A. Muehlemann (muehle@berkeley.edu) is a postdoctoral researcher at UC Berkeley, USA
[†] 1AntonM8kEGMhpaCZwh7kZHEc5cBh1qoeg

(*i.e.*, only the experts that made good predictions are cited), and also, if the expert is influential enough, the ability of turning one's prediction into a self-fulfilling prophecy. For crowds, some opinion polls offer small monetary rewards, usually in the form of a lottery, for filling out a survey. However, these rewards are paid irrespective of the quality of the provided sentiment and in particular they do not prevent anyone from intentionally providing false statements. Furthermore, specifically in web surveys, participants are able to repeatedly reply to the same surveys to increase their potential reward.

A different form and an arguably significantly more successful approach of information aggregation is speculative markets.[8–15] There, free bidding markets of outcome shares are offered and (as dictated by economic theory) those outcome share prices will become representative of the likelihood of an event coming true. Unfortunately, under most legislation (in particular in the US), speculative markets are considered gambling and thus illegal—making this option often unavailable.

In the present article, we introduce a framework that unites the legal benefits of classical polling with the predictive power of speculative markets by introducing pre-defined reward functions whose payouts are both fixed and performance-based. The structure of this article is as follows: first, in Section 2, we introduce the Sentiment Protocol. The section starts with a high-level overview and is followed by a detailed description of each component. Each stage includes basic examples that illustrate the respective concepts. In Section 3, we analyze the incentives for sentiment contributors and pollsters. We also discuss possible vulnerabilities and give suggestions on how they may be overcome. In the final section, Section 4, we discuss in further detail two use cases for the Sentiment Protocol.

2. The Sentiment Protocol

It is now widely accepted that blockchain technology, the underlying concept of Bitcoin,[16] has the ability to disrupt a wide range of industries including financial services, technology, media and telecommunications.[17,18] Arguably, one of the most revolutionary emerging concepts is that of smart contracts, which are scripts executed on a world-spanning super computer (of sorts).[19] The Sentiment Protocol leverages these new possibilities to on the one hand immutably record provided sentiment and on the other hand conduct monetary transactions, such as rewards or penalties, without the need for trusting a third party.

The novelty of the protocol is that it leverages the predictive power of speculative markets while retaining the legality of classical polling. This goal is achieved by introducing a performance-based reward function f_{PE} (cf. Section 2.4), resulting in higher payouts for better predictions, with a reward pool provided by the pollster. Since the pollster takes a distinguished role, speculative risks are taken away from the sentiment contributors.

When the performance-based reward function f_{PE} takes only non-negative values, participants can only earn rewards. Recalling that, by definition, gambling is "the act of risking money, or anything of value, on the outcome of something involving chance," it is clear that in the absence of a risk of loss, the sentiment protocol cannot be considered gambling.[20] Alternatively, if the performance-based reward function is set to also allow for negative values, participants may lose part of their invested stake, making it more similar to classical betting platforms.

Even though the protocol works exactly the same way in either case, the decision to allow or not allow negative rewards can make a big difference in practice. For example, a poll with penalties can directly pay for itself and thus provide a direct incentive to the poll creator. Furthermore, if penalties are high and comparable to rewards, the poll can closely resemble gambling with high potential rewards and high potential losses for participants. In contrast, a poll without penalties can only indirectly pay for itself through the value of the obtained information. In this case, it is likely that potential rewards are significantly smaller.[21]

Protocol overview—The Sentiment Protocol (cf. Fig. 1) has three main components:

(1) Sentiment Contribution Period,

(2) Tallying, and

(3) Performance Evaluation.

During the Sentiment Contribution Period, users (sentiment providers) can submit their sentiment to the protocol. At the end of the Sentiment Contribution Period, and after a cool-down period ΔT_0 has lapsed (which may be 0), votes are tallied. The polling results can either be used for purely informational purposes or, in the use case of decentralized governance (cf. Section 4.1), directly trigger the policy that was voted on. After another cool-down period ΔT_1 (which may again be 0), users receive a (possibly negative) performance based reward.
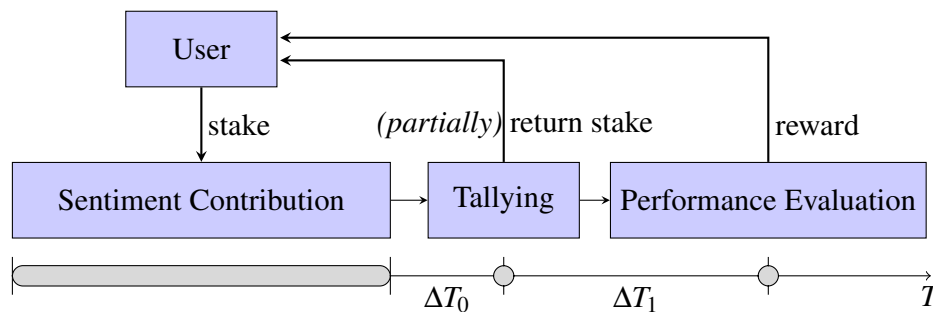

Fig. 1. Overview of the Sentiment Protocol *(with penalties)*.

Setting up a poll—To set up a poll, the creator needs to provide the following:

(1) **Topic and set of possible outcomes** $\rightarrow$ Section 2.1

(2) **Staking parameters** $\rightarrow$ Section 2.2

(3) **Information on usage of results** $\rightarrow$ Section 2.3

(4) **Performance evaluation parameters** $\rightarrow$ Section 2.4

2.1. Topic and outcome set—The creator of the poll has complete freedom over the choice of topic and also the set of possible outcomes $\mathscr{O}$. Outcomes can be discrete, such as multiple choice, or continuous, such as real numbers. The creator also has the option of using a public key encryption scheme for the submission of sentiment. By doing so, third parties are prevented from obtaining knowledge of the already-submitted sentiment.

Example 1 (Discrete Outcome set $\mathscr{O}$)**.** The poll wants to predict the outcome of the 2020 US presidential election. The outcome set is $\mathscr{O} = \{R, D\}$, where R is the Republican and D the Democratic candidate.

Example 2 (Continuous Outcome Set $\mathscr{O}$)**.** The poll aims to predict the performance of the TSLA stock from 2017-11-01 to 2018-05-01. The outcome set is $\mathscr{O} = [0, \infty)$ and $\mathscr{O} \ni o = \frac{p(\text{2018-05-01})}{p(\text{2017-11-01})}$, where p is the price of TSLA at a given date.

2.2. Providing sentiment by staking—To provide sentiment, each participant needs to stake a corresponding amount of assets. For simplicity, we focus on the typical use case where the user needs to stake an ERC-20 token into the polling smart contract. The set of possible sentiments $\mathscr{S} \subset \mathscr{P}(\mathscr{O})$ is a subset of the powerset of the outcome set $\mathscr{O}$.[22, 23]

During the sentiment contribution period any participant that holds the token required for staking can submit his sentiment $s \in \mathscr{S}$. The weight of the vote is directly proportional to the number of tokens submitted, *i.e.* 1 TOKEN = 1 SENTIMENT. The setup of the staking phase requires:

- **Token type**
- **Sentiment Contribution Period**
- **Limits on total submissions**

Having to commit tokens with each choice limits the amount of sentiment each participant can submit. Furthermore, in cases where the performance evaluation (cf. Section 2.4) can lead to penalties, the commitment of tokens exposes the provider to a financial risk. It is important to note that each staked token has equal weight in the sentiment contribution process. In particular, the rewards and/or penalties are directly proportional to the amount of tokens staked. This choice is due to the pseudonymous nature of the blockchain, where it is futile to limit sentiment per address. However, each pollster is free to use a custom staking token. Such a custom token could for instance be issued by the polling company and may have restrictions on transferability. Even in cases where no penalty is possible, obtaining a large amount of staking tokens puts the buyer at the volatility risk of the token price. Moreover, even if contributors are willing to accept the volatility risk, as a major stakeholder, they would have little interest in a behaviour that would undermine the polling process.

Once the Sentiment Contribution Period has ended and a time ΔT_0 (which may be 0) has lapsed, the votes are tallied. If the votes were encrypted, the pollster will have to use their private key to decrypt the submissions. If the performance evaluation does not involve penalties, the stake is returned to the sentiment provider. If penalties are possible, the stake minus the maximal possible penalty is returned.

Example 3 (Staking without penalties). Let us return to the previous example of the 2020 presidential election and assume that no penalties are possible. We set $\Delta T_0 = 24h$, $\mathscr{S} = \mathscr{O} = \{R, D\}$ and choose the following staking parameters:[24]

- Type: ETH,
- Sentiment Contribution Period: 2017-12-01 to 2017-12-10,
- Limits on total submissions: Minimum 1000 ETH (1000 SENTIMENT) and maximum 10,000 ETH (10,000 SENTIMENT).

Thus, a sentiment provider who wants to provide 100 SENTIMENT for the Democratic candidate D needs to submit 100 ETH together with his (possibly encrypted) choice D to the polling smart contract.[25] Since no penalties are possible, 100 ETH are returned to the sentiment provider on 2017-12-11.

2.3. Usage of results—Apart from simply aggregating information on behalf of the pollster, one could also link the execution of certain events to the results of the tally. For instance, one could ask individuals to estimate their energy consumption and, if it is below a certain threshold, a power plant could be idled. If individuals report inaccurate information they could be penalized

51

by either not earning rewards or by losing some of their stake. Another—often controversial—use case are decentralized autonomous organizations.[26] In this case, the polling results could be used to autonomously implement policy changes within the organization (cf. Section 4.1).

2.4. Performance evaluation—The payment of performance-based rewards is a key feature of the Sentiment Protocol. The distribution of rewards/penalties is determined by the performance evaluation function provided by the poll creator.

Definition 1 (Performance evaluation function f_{PE})**.** The performance evaluation function $f_{PE} : \mathscr{O} \times \mathscr{S} \to [-1, \infty)$, depends on the outcome $o \in \mathscr{O}$ and the submitted sentiment $s \in \mathscr{S}$ and specifies the reward/penalty per submitted token.[27]

The supremum of the performance evaluation function determines the size of the reward pool that needs to be provided to set up the poll.[28] It is given by

$$(\text{reward pool}) = \sup f_{PE} \cdot (\text{limit on total submissions}),$$

where all values are in units of tokens. If $f_{PE} \geq 0$ (as a function), the poll does not involve penalties and thus the stake is returned immediately after tallying. If f_{PE} takes negative values, $(1 + \inf f_{PE}) \cdot T$ tokens are returned after tallying, where T is the number of submitted tokens.

Once the performance evaluation time is reached (ΔT_1 after tallying), the rewards are determined and paid out to the sentiment providers. Any remaining balance in the reward pool is returned to the pollster. See also Appendix A for an extension of the protocol that allows multiple performance evaluations.

Example 4 (Constant positive f_{PE})**.** In this trivial case $f_{PE} \equiv c$ for some positive constant $c \in \mathbb{R}^+$. Thus, the reward is independent of both the provided sentiment and the outcome. Assuming the same staking parameters as in Example 3, the reward pool is $10{,}000 \cdot c$ and, since the reward is independent of the outcome or sentiment, there is no need to wait ($\Delta T_1 = 0$) and all sentiment providers receive $(1 + c)T$ tokens on 2017-12-11, where T is the number of submitted tokens. This trivial case corresponds to the current practice of most polling companies.

Example 5 (Positive f_{PE} for discrete $\mathscr{O}$)**.** We return to Example 3. To incentivize sentiment providers to make good predictions we define

$$f_{PE}(D,D) = f_{PE}(R,R) = c \text{ and } f_{PE}(D,R) = f_{PE}(R,D) = 0.$$

Thus, only people that voted for the winning candidate get rewarded. Since the results of the elections will not be finalized until mid November 2020, $\Delta T = 1071$ days. As in Example 4, the reward pool is $10{,}000 \cdot c$. Compared to Example 4, participants will only get rewarded for correct predictions and thus are incentivized to think more carefully about their forecasts. Furthermore, with the same parameters as in Example 4, the poll creator will have to pay fewer rewards. Alternatively, the creator could increase the rewards and make participation more appealing.

Example 6 (f_{PE} for continuous $\mathscr{O}$ with penalties)**.** We return to Example 2 and use the same staking parameters as in Example 3. The sentiment set is $\mathscr{S} = \{\text{buy}, \text{sell}\}$ and we wish to choose f_{PE} such that providers of buy ratings get rewarded for a positive development of the TSLA stock price ($o > 1$) and penalised for a negative development.[29] Similarly, providers of sell ratings shall get rewarded for a negative development of the TSLA stock price ($o < 1$) and penalized

for a positive development. In the interest of fairness, if the stock price *e.g.* doubles ($o = 2$), the buy rater should get the same reward as a sell rater if the stock price halves ($o = 1/2$). A natural choice for f_{PE} that satisfies these requirements and is also bounded is

$$f_{PE}(o,s) = \frac{2c}{\pi} \cdot \text{sgn}(o-1) \cdot \arctan\left(\max\left\{o-1, o^{-1}-1\right\}\right) \cdot g(s),$$

where $g(\text{buy}) = 1$ and $g(\text{sell}) = -1$. Note that arctan has the appealing properties of being approximately linear around 0, strictly monotone increasing but still bounded by $\pi/2$. Thus for small price changes

$$f_{PE}(o,\text{buy}) \approx \begin{cases} c(o-1), & o \geq 1 \\ -c(o^{-1}-1), & o < 1 \end{cases}$$

and similarly $f_{PE}(o,\text{sell}) \approx -c(o-1)$ for $o \geq 1$ and $f_{PE}(o,\text{sell}) \approx c(o^{-1}-1)$ for $o < 1$. Owing to monotonicity, higher (sentiment-aligned) performance results in higher rewards. As in the previous examples, the reward pool is $10{,}000 \cdot c$. The maximal possible penalty is $-c$ and thus only $(1-c)T$ tokens are returned after tallying. If the value of c is close or equal to 1, participants can use this poll to bet for or against the stock. Since c is large, possible returns and losses are comparable to their invested stakes. This risk of loss is likely to result in more accurate forecasts. Apart from the gained insights, the poll creator can also benefit monetarily from the poll if the majority of participants makes inaccurate predictions.

3. Incentives and Vulnerability Analysis

In this section, we analyse the incentives for pollsters and sentiment providers to use the Sentiment Protocol. We distinguish between public polls, polls among experts, and gambling. We also explain its advantages over current (centralized) polling mechanisms and discuss possible vulnerabilities.

3.1. Benefits of using blockchain technology—Firstly, it is important to note that opinion polling is not a truly decentralized mechanism and in particular the Sentiment Protocol does not claim to be fully decentralized. Nonetheless, using blockchain technology offers important advantages over classical, completely centralized, solutions. For the Sentiment Protocol, the two most important properties of blockchain technology are:

    (1) the ability to easily transact and store value without the need of trusting a third party, and

    (2) the ability to immutably store data.

The Sentiment Protocol uses (1) for both the staking and the rewards/penalties. In a centralized system, users are hesitant to commit a significant value for period of time ($\geq \Delta T_0$) for the chance of receiving a comparatively small reward. Furthermore, by additionally using (2), the submitted sentiment is immutably stored and the user can easily prove that he is entitled to receive a reward for an outcome—which may be far ahead in the future (cf. Example 5 and Appendix A).

3.2. Public polls—The goal of a company conducting public polls is to receive reliable sentiment from a large number of users. The users on the other hand wish to be rewarded for providing accurate information. In the following we describe a possible setup of the Sentiment Protocol to align their interests.

To ensure the diversity of sentiment providers, the polling company issues their own token called POLL which can be arbitrarily transferred to and from polling smart contracts but cannot

be transferred between users. However, all earned rewards are freely transferable. Further, we assume that at most times there are several polls to chose from and that each user has a starting balance of 100 POLL. The performance evaluation function f_{PE} is chosen positive (*i.e.* no penalties) and the cool down period ΔT_0 is one week. Due to this long cool-down period, POLL holders are incentivized to participate in polls where they feel most confident in predicting the result. Unless they have absolute certainty that their prediction is correct, they are also incentivized to diversify their tokens on several polls. If however, they (think that they) are absolutely certain about the outcome, they may be inclined to stake all their POLL tokens on a single event. Of course, such behaviour is also in the interest of the polling company.

3.3. Experts' opinions—The goal of the polling company is to receive high quality predictions. However, in contrast to public polls, the users (experts) may wish to demonstrate their commitment to a forecast by putting their own capital at risk and thus earn credibility.

Since the users have an interest in demonstrating credibility, the polling company chooses a performance evaluation function with penalties. Since each submission of sentiment may result in a loss of value, the company can simply choose ETH as a staking token. A sentiment provider that feels more certain about a prediction will choose to stake more tokens than a provider who is not as certain. As in the previous case, this behaviour is in the interest of the polling company.

3.4. Gambling—If penalties and rewards are comparable to the invested stake, the sentiment protocol can resemble classical gambling platforms. In particular, the protocol inherits the well-known properties of prediction markets such as the possibility of high rewards for participants, very accurate predictions, and a long-term monetary profit for the poll creator.

3.5. Limitations and possible vulnerabilities—An important limitation of blockchains is their inability to source external information. Thus, in order to determine rewards, users need to rely on oracles. This problem is not specific to the Sentiment Protocol but poses a general problem in developing truly decentralized applications. Another important limitation of smart contract platforms is their high cost for computations. Thus, the cost for determining performance based rewards may be significant.

If these general limitations of blockchains can be overcome (or are irrelevant as in Example 4) and if the incentive structure in the Sentiment Protocol is set up correctly, it is impossible for the pollster or the user to illegitimately receive funds.

Owing to currently having these limitations, the polling company may act in a much more centralized manner and may choose to conduct the rewards calculations off-chain with data provided by regular (centralized) data-feeds. Of course, this special role allows the pollster to illegitimately keep rewards. However, due to the public availability of the provided sentiment and the performance evaluation function, cheating by the pollster can easily be proven and the polling company will quickly lose credibility. Thus, cheating by the company becomes unprofitable in the long run.

To consider possible cheating from the users' perspective it is convenient to distinguish between performance evaluation function with and without penalties. If no penalties are possible then a user has no risk of loss. However, even if the pollster does not use a custom token as detailed in Section 3.2, the commitment of huge amounts of tokens is on the one hand very costly and on the other hand only profitable if the predictions are right in the end. Furthermore, if the token is native to the polling company, amassing large amounts makes the user a major stakeholder of the company and thus any misbehaviour counterproductive.

If penalties are possible and if potential rewards/penalties are sufficiently large, it may be advisable to choose a 'zero-sum' performance evaluation function, *i.e.* there should not exist a combination of sentiments and staking choices such that the sentiment provider earns a profit independent of the outcome. Of course, as in the case without penalties, it is at the discretion of the pollster to allow this possibility.

4. Use Cases

In this section we discuss two specific use cases. We will not discuss the two cases in full details but rather assume familiarity with Section 2.

4.1. Governance—Any token that is able to access smart contract functionality (*e.g.* ERC-20 tokens) can use the Sentiment Protocol for decentralized governance. Since policy changes are not objectively verifiable as right or wrong, it is natural to choose a constant performance evaluation function $f_{PE} \equiv c$ (cf. Example 4). Thus, each token holder that participates in the governance process gets a fixed reward for his engagement. The required reward pool can be provided by the token issuing company. To ensure an indefinite supply of reward pools R_i for the ith voting event, the company can *e.g.* use a distribution according to a geometric series. That is, if the company plans on using 100,000 tokens for incentivizing governance participation (via the Sentiment Protocol), they can set $R_{i+1} = x \cdot R_i$,[30] where $x < 1$ and

$$R_1 = 100,000 \cdot (1-x). \tag{1}$$

Since the performance evaluation function is independent of the outcome, neither oracles nor complex computations are necessary and thus the (governance) Sentiment Protocol can be entirely implemented on-chain. In particular, this allows running a decentralized autonomous organization, where user are allowed to vote on policy changes and decisions are automatically executed according to the voting results.

4.2. Community driven rating agency—In this scenario, the polling company wishes to regularly receive sentiment on the performance expectations of a wide range of stocks. To ensure broad reach, the polling company issues 1,000,000 custom tokens of which 900,000 are distributed to the public and 100,000 are held by the company to fund reward pools. The transferability of originally distributed tokens between users is limited (*e.g.* only 10% per quarter). However, earned rewards can be transferred freely. To ensure the indefinite supply of rewards pools, the company chooses $x = 0.99 = 99\%$ (cf. (1)) and thus the reward pool for the first sentiment round is $R_1 = 1000$ tokens and each subsequent pool is 99% of the size of the previous pool. In each round the users are asked to provide their sentiment on the performance of 10 different stocks within the next three months ($\Delta T_1 = 3$ months) by choosing between $\Uparrow$, $\Leftrightarrow$ and $\Downarrow$ for each stock.

The performance evaluation function is defined by $f_{PE}(o, \Uparrow) = c \min\{1, \max\{0, o-1\}\}$, $f_{PE}(o, \Downarrow) = c \min\{1, \max\{0, 1/o - 1\}\}$ and

$$f_{PE}(o, \Leftrightarrow) = \begin{cases} 5c(1.1 - 1/o) & o \in [0.\overline{90}, 1], \\ 5c(1.1 - o) & o \in [1, 1.1], \\ 0 & \text{otherwise.} \end{cases}$$
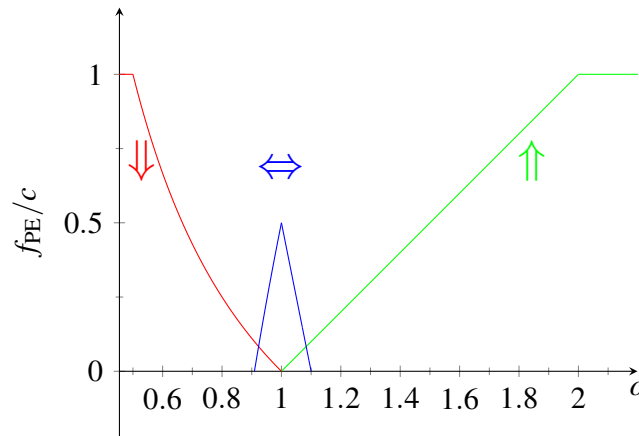
Fig. 2. The performance evaluation function f_{PE} for the three sentiments $\Downarrow$, $\Leftrightarrow$ and $\Uparrow$. The red curve indicates the function $o \mapsto f_{PE}(o, \Downarrow)$, the blue curve the function $o \mapsto f_{PE}(o, \Leftrightarrow)$ and the green curve the function $o \mapsto f_{PE}(o, \Uparrow)$. Lines not shown correspond to $f_{PE} = 0$.

Thus, a $\Uparrow$-rater gets rewarded only for positive stock performance and his reward is linear in the performance with a cap at 100%. A $\Leftrightarrow$-rater receives a maximal payout if the stock price does not change and does not get rewarded if the performance is outside the interval $(0.\overline{90}, 1.1)$. See Fig. 2 for a plot of this function. We note that the performance evaluation function only takes non-negative values and thus there is no risk of loss for participants. In particular, participation in the poll does not involve gambling.

Since the reward pool is 1000 tokens and 900,000 tokens are held by the sentiment providers we set $c = 1000/(10 \cdot 900,000) = 0.000\overline{1}$. By allowing the rating of 10 stocks simultaneously, users are incentivized to prioritize stocks (by staking more tokens) that they feel most confident in making a good prediction.

The obtained sentiment data is a valuable asset to the company and could for instance be sold to third parties or used to create a community driven portfolio. In the latter case, the community itself would have the chance to invest in such a portfolio and thus double their incentive to participate in the rating process.

5.   Conclusions

In this paper, we have shown how the Sentiment Protocol can leverage blockchain technology to align the incentives of pollsters and sentiment contributors. We discussed several use cases including decentralized governance and prediction markets. By introducing a fixed performance evaluation function we are able to reward predictions finely graduated. We have shown how such an approach offers a clear advantage over classical polling solutions, where rewards are either non-existing or small and independent of the quality of the contributed sentiment, and also over classical prediction markets, which are considered illegal in most jurisdictions.

Acknowledgments

Notes and References

[1] No Author. "Revenue of market research and public opinion polling." *Statista* (accessed 17 October 2017) `https://www.statista.com/forecasts/409757/united-states-market-research-and-public-opinion-polling-revenue-forecast-naics-54191`.

[2] No Author. "Nationwide opinion polling for the United States presidential election, 2016." *Wikipedia* (accessed 17 October 2017) `https://en.wikipedia.org/wiki/Nationwide_opinion_polling_for_the_United_States_presidential_election,_2016#Polls_conducted_in_2016`.

[3] No Author. "Opinion polling for the United Kingdom European Union membership referendum, 2016." *Wikipedia* (accessed 17 October 2017) `https://en.wikipedia.org/wiki/Opinion_polling_for_the_United_Kingdom_European_Union_membership_referendum`.

[4] Parker, J. "Analysts Remain Uncertain on Tesla Stock amid Production Woes." *Market Realist* (2017) (accessed 17 October 2017) `https://marketrealist.com/2017/10/analysts-remain-uncertain-on-tesla-stock-amid-production-woes/`.

[5] Ferri, R. "Any Monkey Can Beat the Market." *Forbes* (2012) (accessed 13 August 2018) `https://www.forbes.com/sites/rickferri/2012/12/20/any-monkey-can-beat-the-market/`.

[6] Lowenstein, R. "Why Buffett's Million-Dollar Bet Against Hedge Funds Was a Slam Dunk." *Fortune* (11 May 2016) (accessed 17 October 2017) `http://fortune.com/2016/05/11/warren-buffett-hedge-fund-bet/`.

[7] Berk, J. B., Green, R. C. "Mutual fund flows and performance in rational markets." *Journal of political economy*. **112.6** 1269–1295 (2004).

[8] Hirshleifer, J. "The private and social value of information and the reward to inventive activity." *The American economic review*. **61.4** 561–574 (1971).

[9] Berg, J. E., Rietz, T. A. "Prediction markets as decision support systems." *Information systems frontiers*. **5.1** 79–93 (2003).

[10] Berg, J. E., Nelson, F. D., Rietz, T. A. "Prediction market accuracy in the long run." *International Journal of Forecasting*. **24.2** 285–300 (2008).

[11] Epley, N., Keysar, B., Van Boven, L., Gilovich, T. "Perspective taking as egocentric anchoring and adjustment." *Journal of personality and social psychology*. **87.3** 327 (2004).

[12] Hanson, R. "Shall we vote on values, but bet on beliefs?" *Journal of Political Philosophy*. **21.2** 151–178 (2013).

[13] Buterin, V. "An Introduction to Futarchy." *Ethereum Blog* (2014) (accessed 30 October 2017) `https://blog.ethereum.org/2014/08/21/introduction-futarchy/`.

[14] Peterson, J., Krug, J. "Augur: a decentralized, open-source platform for prediction markets." *arXiv preprint* (2015) `https://arxiv.org/abs/1501.01042`.

[15] Koppelmann, M., George, S. "Gnosis whitepaper." *Gnosis* (2017) (accessed 22 October 2017) `https://gnosis.pm/assets/pdf/gnosis-whitepaper.pdf`.

[16] Nakamoto, S. "Bitcoin: A peer-to-peer electronic cash system." (accessed 22 October 2017) `https://bitcoin.org/bitcoin.pdf`.

[17] Schatsky, D., Muraskin, C. "Beyond Bitcoin: Blockchain is Coming to Disrupt Your Industry." *Deloitte* (2015) (accessed 2 July 2018) `https://www2.deloitte.com/insights/us/en/focus/signals-for-strategists/trends-blockchain-bitcoin-security-transparency.html`.

57

[18] Wattenhofer, R. *The science of the blockchain*. CreateSpace Independent Publishing Platform (2016).

[19] Buterin, V. "Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform." *Ethereum* (accessed 17 October 2017) `http://ethereum.org/ethereum.html`.

[20] No Author. "Definition of Gambling." *dictionary.com* (accessed 28 March 2018) `http://www.dictionary.com/browse/gambling`.

[21] Compare this situation to classical polls that only raffle a gift card among participants.

[22] Recall, that the powerset $\mathscr{P}(\mathscr{O})$ of a set $\mathscr{O}$ is the set of all subsets of $\mathscr{O}$.

[23] For the technical reader: if $\mathscr{O}$ is continuous we are using the Borel sigma algebra $\sigma(\mathscr{O})$ instead of $\mathscr{P}(\mathscr{O})$.

[24] For the technical reader: we identify $\{\{R\},\{D\}\} \cong \{R,D\}$.

[25] For the technical reader: the user needs to encrypt the value with a nonce and also submit the encrypted nonce.

[26] Cuende, L., Izquierdo, J. "Aragon white paper." *Aragon* (2017) (accessed 17 October 2017) `https://wiki.aragon.one/documentation/whitepaper/`.

[27] The technical reader may have noticed that the triple $(\mathscr{O},\mathscr{S},f_{\mathrm{PE}})$ is reminiscent of a probability space. Of course, this is not a coincide.

[28] If $|\mathscr{O}| < \infty$, $\sup f_{\mathrm{PE}} = \max f_{\mathrm{PE}}$.

[29] We can identify the *sell* rating with $[0,1) \in \mathscr{P}(\mathscr{O})$ and the *buy* rating with $(1,\infty) \in \mathscr{P}(\mathscr{O})$.

[30] Indeed, $\sum_{i=1}^{\infty} R_i = 100{,}000\,(1-x)\sum_{i=0}^{\infty} x^i = 100{,}000\,(1-x)\frac{1}{1-x} = 100{,}000$.

Appendix A: Sentiment Protocol with multiple performance evaluations

To incentivize users to provide sentiment on long term performance (*e.g.* stock prices), we extend the protocol to allow several performance evaluations (cf. Fig. 3). Let us illustrate this by assuming that the performance is evaluated every three months, *i.e.* $\Delta T_i = 3$ months. To this end, we can *e.g.* define the performance evaluation function f_{PE_i} at the ith performance evaluation as

$$f_{\text{PE}_i} = 2^{-i} f_{\text{PE}},$$

where f_{PE} is the original (single event) performance evaluation function. Since $\sum_{i=1}^{\infty} 2^{-i} = 1$, the reward pool does not need to be increased. Of course, the pollster is free to make different choices for f_{PE_i}. The only restriction is that the reward pool needs to be big enough to cover *all* performance evaluations.
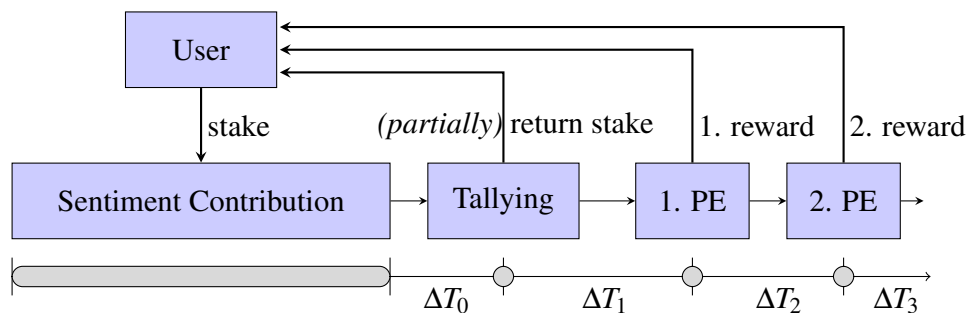

Fig. 3. Overview of Sentiment Protocol *(with penalties)* with long term rewards.