LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

Sentiment Protocol: A Decentralized Protocol Leveraging Crowdsourced Wisdom: Open Review

Author: Anton Muehlemann[*†]

Reviewers: Reviewer A, Reviewer B, Reviewer C, Reviewer D

Abstract. The final version of the paper "Sentiment Protocol: A Decentralized Protocol Leveraging Crowdsourced Wisdom" can be found in Ledger Vol. 3 (2018) 48-59, DOI 10.5915/LEDGER.2018.113. There were four reviewers involved in the review process, none of whom have requested to waive their anonymity at present, and are thus listed as A, B, C, and D. After initial review by Reviewers A, B, and C (1A), it was determined that the submission required revisions. The author responded to their feedback (1B) and revised the manuscript. A third second-round reviewer, Reviewer D, was added to review the revised manuscript (2A). The author responded to their feedback (2B) and revised the manuscript. The Editors reviewed and accepted these changes, thus completing the peer-review process.

1A. Review (Initial)

Reviewer A:
Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?:

Yes

If you answered "yes" to the previous question, in one sentence, describe in your own words the novel contribution made by this paper:

A systematic approach to perform incentivized polls and elections on blockchains.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?:

Yes

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

Please assess the article's level of academic rigor:

Good (not excellent but a long way from poor)

Please assess the article's quality of presentation:

Excellent (the motivation for the work is clear, the prose is fluid and correct grammar is used, the main ideas are communicated concisely, and highly-technical details are relegated to appendixes).

How does the quality of this paper compare to other papers in this field?:

Top 20%

Please provide your free-form review for the author in this section:

I have severe problems in following the derivation and goal of the authors mathematics. Although I am mathematically well educated, the derivations and formulas are hard - partly impossible - to understand. The authors expects too much background knowledge from the reader. A better explanation would be helpful.

Additionally for me it is hard to understand, why there is so much mathematics in this paper. Even the author doesn't make subsequent proofs or derivations from his own equations. Is the goal of the mathematics to make the article look more interesting or scientific?

It is already on the protocol level very interesting and well elaborated. Why so many "stand-alone-equations"?

Further and detailed comments are in the uploaded pdf.

Reviewer B:

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?:

Yes

If you answered "yes" to the previous question, in one sentence, describe in your own words the novel contribution made by this paper:

The wisdom of crowds as quantized sentiment provides a new branch of actionable consensus/agreement

Is the research framed within its scholarly context and does the paper cite appropriate prior works?:

Yes

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

Please assess the article's level of academic rigor:

Excellent (terms are well defined, proofs/derivations are included for theoretical work, statistical tests are included for empirical studies, etc.)

Please assess the article's quality of presentation:

Excellent (the motivation for the work is clear, the prose is fluid and correct grammar is used, the main ideas are communicated concisely, and highly-technical details are relegated to appendixes).

How does the quality of this paper compare to other papers in this field?:

Top 5%

Please provide your free-form review for the author in this section:

This paper describes a new system for poll creation with protocolized incentives which improve outcome accuracy, participation, and the ease of poll moderation. The paper includes descriptions of novel contributions to extant technologies:

1. Polls exist on-chain with optionally private or transparent outcomes and submissions
2. Poll moderators can issue custom tokens against single or multiple polls
3. Widespread poll accessibility
4. Avoids the gambling designation of prediction markets; which prevent similar systems from existing in many jurisdictions

General comments:

A landscape where particular polls or groups of polls issue their own token rather than general format tokens provides for new incentive structures that currently do not exist.

The section on voting anonymization could leverage random sample elections (D.Chaum, et al.)

Example 6 is reminiscent of the current state of the NA stock market. An implementation of the sentiment protocol could provide an alternative to investing in a stock and simultaneously hedging that position in order to take a % sentiment on the stock's binary outcome in delta time with conflated external factors. The sentiment protocol is a more concise methodology for overcoming the user experience challenges of equities investment with a robust option for deal-specific carve-outs. ie. "Do you think this stock will go up or down?" with a slider for a percentage of certainty and other conditions such as "I believe f with x% certainty if yz happens."

Conclusion:

The sentiment protocol addresses multiple challenges in sentiment and prediction. A

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

protocolized architecture for polling can be extended into existing chains via rewards in native coins, tokens, and poll/pollster-specific new tokens. Storing poll outcomes on-chain with public access and optional privacy can provide greater public access to empirical data. A great deal of academic research could benefit from an accessible, reliable, verifiable, and transparent poll/survey ecosystem. Empirical data under peer review could have another element of public accountability if it exists on-chain. (On-chain can denote immutable chain of custody, trusted identity references, and an overall difficulty in fabrication.)

I overwhelmingly recommend that this paper become published, read, and implemented.

Reviewer C:

Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?: No.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Yes

Please assess the article's level of academic rigor:

Unsatisfactory (better than poor but a long way from excellent)

Please assess the article's quality of presentation:

Unsatisfactory (better than poor but a long way from excellent)

How does the quality of this paper compare to other papers in this field?:

Top 50%

Please provide your free-form review for the author in this section:

The author tries to be everything to everyone and thereby fails to be much of anything to anyone. At this point in time, pointing out that we can use gambling markets to get useful data about (for example) elections is not new, and discussing even a large number of variations on a theme without data, or claiming that some of the variations can overcome legal restrictions without mentioning a single legal example from a single jurisdiction, is not going to advance the state of the art.

1B. Author's Response

Response to Reviewer A:

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

Thank you for your careful reading and your feedback. I appreciate your many valuable comments. I hope that I am able to address all of them here.

Firstly, I changed all the highlighted typos - thank you for spotting them. I also reverted to the spelling “role” instead of “rôle”. According to www.oed.com both spellings are correct but the former is more common in American English.

Concerning the general comment that the manuscript is too heavy in mathematics, I can only par- tially understand this concern. This may be because my background is in mathematics. In fact, when preparing this manuscript I tried to keep the mathematical formalism to a minimum - avoiding any rigorous treatment of probability and outsourcing comments on formal definitions to footnotes. Also, most of the mathematics is limited to examples which only serve illustration purposes. In particular, the article can be read without looking at them and thus without getting into the mathematical details of each of them.

On a conceptual level, in terms of the actual Sentiment Protocol, I think some general familiarity with probability theory (not in a rigorous sense) can be beneficial to understand the connection between outcomes, sentiments and the performance evaluation function. However, I don’t think that it is necessary to follow the majority of the paper. I added a footnote after Definition 1 (p. 5) that highlights the connection between (O, S, fP E) and a probability space. Thus, the interested reader can do some further reading on that subject. I avoided using the term probability space in the paper to keep mathematical formalism to a minimum.
For the remaining comments I am using to the reviewer’s numbering.

1. Thank you for your comment on DAOs - I used your suggestions to include IoT as a use case and also emphasized that (for now) the main use case would be polling. I changed:

“Depending on the choice of parameters (cf. Section 2), the Sentiment Protocol can be used for decentralized governance, similar to projects such as Aragon, or to obtain predictions for e.g. political events or stock prices.”

to:

“Depending on the choice of parameters (cf. Section 2), the Sentiment Protocol can replace classical polling systems, can be used in financial markets to hedge positions, can activate triggers in IoT or execute policy changes in decentralized autonomous organizations.”

I also added a corresponding sentence in the abstract.

2. An element o in O, the set of outcomes, would be called an outcome.

3. The powerset (or it’s continuous analogue the Borel sigma algebra) is the biggest possible event space given a sample space. Sample space, event space and a “weighting function” are the three fundamental objects that describe a *probability space*. Wikipedia has a good overview on this topic and also lists several references for further reading.

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

4. I changed the part on “Usage of results” to be more general and included any type of event trigger rather than limiting it to DAOs. I changed:

“Apart from simply aggregating information on behalf of the pollster, the Sentiment Protocol provides the basis for running decentralized autonomous organizations. In particular, polling results can be used to directly execute other smart contracts to e.g. implement policy changes within the organization (cf. Section 4.1).”

to:

“Apart from simply aggregating information on behalf of the pollster, one could also link the ex- ecution 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 by either not earning rewards or by loosing some of their stake. Another – often controversial – use case are decentral- ized autonomous organizations.2 In this case, the polling results could be used to autonomously implement policy changes within the organization (cf. Section 4.1).”

5. The number −1 stems from the staking requirement. As the payout is proportional to the stake, the worst possible ’payout’ can only be −1, i.e. the participant loses his entire stake. A value smaller than −1 would imply that he loses more than he staked.

6. c is a positive constant in the real numbers. I added an insert to make this more explicit. ″In this trivial case fP E ≡ c for some positive constant c ∈ R+.″

7. The examples are used to provide different scenarios on how the model can be applied. As such they are not all the same (intentionally). ΔT1 is the waiting time between the tallying event and the performance evaluation. In Example 4, the performance evaluation does not depend on anything (particularly not on performance) and thus there is no need to delay the payout. That is why ΔT1 = 0 (see also Fig. 1 in the manuscript).

8. First of all, indeed it should be o instead of x. Thank you for spotting this. I have changed this. As mentioned before, the examples are just one of infinitely many different choices one can make and are purely for illustration purposes. It is clear that the performance evaluation function needs to be bounded below by 1 (see 5). If, furthermore, one wanted wanted to have some symmetry, he upper bound would need to be +1. Now, if one wanted to reward better performance with more rewards then one needs to find a monotone increasing and bounded function. If one furthermore requires linearity around 0, i.e. rewards are directly proportional to performance for small changes, then one naturally arrives at arctan. Again, it is not integral to the paper to understand each example. Everything in the paper can in fact be done and understood without looking at any of the examples. I tried to give a variety of examples with some being trivial, such as Example 4 (a performance function that does not depend on performance), and some more elaborate ones such as Example 6.

9. x determines the decrease in supply. For bitcoin, this x is set to 1/2 every four years (i.e. block rewards half every four years). Such a mechanism is necessary if one wants to have a limited supply (such as 21,000,000 bitcoins). If x > 1, there would be a hyperinflation.

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

10. Such a behavior can be specified in the token smart contract. Once vested the tokens will not be restricted anymore and I did not intend to color them. However, again, this is just an example and the actual implementation will depend on the intention of the creator of the poll/series of polls.

11. I extended the caption of the figure, explicitly stating which color correspond to which sentiment. Concerning numbering of equations: they are only numbered if referred to in the text.

Response to Reviewer B:

Thank you for your careful reading and your feedback. I am particularly happy about how you try to push the developed ideas even further. Your ideas are new and interesting to me. When developing the sentiment protocol it was exactly my intention to have others use this general framework for their own use cases. It is great to see how this seems to have already begun.

Response to Reviewer C:

Thank you for your careful reading and your feedback. Unfortunately, it seems that some of the main messages of this paper did not come through. As explained in Section 5, the sentiment protocol aims to take a step away from the gambling nature of classical prediction markets. This is achieved by introducing an asymmetry between the pollster and the participants. Just as in classical polling, the pollster needs to incentivize participation by putting up a reward pool - similar to e.g. amazon gift cards or fixed rewards. But instead of randomly distributing the rewards to the participants, the rewards will depend on the outcome (which is a desirable feature of prediction markets).

Concerning your comment on the lack of evidence that such an approach would indeed be legal, it is of course difficult to find court cases that deal with matters that have never been considered illegal in the first place. Admittedly, I am not a legal expert and also the paper is a technical paper that leaves any particular implementation to the user. However, I like to argue that if such a system (a distribution of rewards without any downsides) would be deemed illegal, then all polling that offers any kind of incentive for participation would also need to be deemed illegal. Even without a legal background, I doubt that this is the case. If it was the case, the number of companies to sue would be endless and would include Google, iPoll and many others.

2A. Review, Second Round

Reviewer A:

Although I still have the opinion that in terms of readability and understanding, maths should be used there where it serves proofs and understanding and here it is also used for compression and illustration purposes, my understanding and the readability of the paper has been improved to the last version.

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

Reviewer C:

I've re-reviewed the paper. The author didn't address my concerns, so my previous review stands.

Reviewer D:
Does this paper represent a novel contribution to cryptocurrency or blockchain scholarship?:

Not sure. The author may have some novel work in the ability to use this protocol as a legal alternative to betting markets.

Is the research framed within its scholarly context and does the paper cite appropriate prior works?

Yes

Please assess the article's level of academic rigor:

Good (not excellent but a long way from poor)

Please assess the article's quality of presentation:

Good (not excellent but a long way from poor)

How does the quality of this paper compare to other papers in this field?:

Top 50%

Comments:

This paper has a problem with novelty, or at the least with the way it presents itself. The problem as I see it is that the idea to use a cryptocurrency gambling market to harness the wisdom of the crowd isn't novel (because as the author points out, others have already done this). The novelty in this paper is the ability to do things in a way that more closely resembles classical polling, where the costs are taken on by the pollster, which would arguably make it more legal. I think to properly address this, the author would have to reorganize the paper, and more convincingly argue the case that it would be legal (This is not to say that the argument couldn't be made, merely that at present it isn't being made. There is a significant difference between making a claim and making an argument, and at present only the former is present.). I would suggest taking section five and starting the paper there, with Futarchy, and making the case for the author's Sentiment Protocol as a potential solution to the problem of illegality. From there the author could describe other potential benefits of the protocol, including the possibility of using the same protocol for other scenarios. Because it lacks a focus on its novelty, the paper goes in a number of directions at once, and so a reader must go hunting

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

through the paper to look for any novelty. In my opinion it ought to be presented much more clearly from the beginning, and more forcefully argued, at that.

2B. Author's Response

Response to Reviewer A:

Thank you. I am glad to hear that my changes improved the readability and, again, I am very grateful for your previous suggestions.

Response to Reviewer C:

Thank you for your updated review. From my understanding, your concerns were similar to the ones raised by Reviewer D and I reorganized the paper accordingly (see above). I hope that this addresses your concerns.

Response to Reviewer D:

Thank you for your constructive criticism.

I followed your suggestion and reorganized the paper to put more emphasis on the fact that the sentiment protocol unites the legal benefits of classical polling with the predictive advantages of speculative markets. These changes are reflected in several places.

Abstract:

[...] 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.

Introduction

[...] increase their potential reward.

A different form and an arguably significantly more successful approach of information aggregation are speculative markets. There, free bidding markets of outcome shares are offered and, dictated by economic theory, share prices will become representative of the likelihood of an event coming true. Unfortunately, under most legislations (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

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

functions whose payouts are both fixed and performance based. The structure of this article is as follows [...]

Sentiment Protocol (Section 3)
[..] 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 (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. Whenever the performance based reward function fPE only takes 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. If the performance based reward function can also attain negative values, participants can lose part of their invested stake, making it more similar to classical prediction markets.

Use Cases (Section 4)
[...] 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.

Conclusions (now Section 5)
[...] 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.
By incorporating the comparison with (and benefits over) prediction markets early on, I felt like (former) Section 5 was now superfluous and removed it entirely.

3A. Review, Third Round

Reviewer D:

This is an improvement over the last draft, and does a better job of highlighting the paper's contributions overall. I think it could be further improved by gearing the examples provided more toward the two Use Cases outlined in section 4.

While the author takes pains to repeat that the version of the protocol in which rewards can at worst be zero (i.e. lack of penalties) does render the protocol legal in jurisdictions in which betting markets are forbidden, I feel as though it could be more fully explored, especially in regard to incentives.

I would like to see a little more explanation of why a pollster would run the protocol for the given examples—for example, in a case where provider loss *is* possible, the poll would be able to pay for itself, and therefore the insights it provides to the pollster would be more likely

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

to be profitable (if accurate). But in a case in which the rewards are provided by the pollster, then direct profit is much less likely a motive for the pollster (which, of course, makes sense for the election example, but may need language explicitly saying so). The Use Cases section does a good job of this, but the "examples" come first, and lead to reader questions that may weaken the paper's position.

In the case of Example 6, for example, in which the providers could lose funds for a wrong guess on what would happen to Tesla stock, what incentive do they have to invest in the poll and not simply in Tesla stock itself? In Use Case 4.2, where the potential for loss on the participants' side is negated, however, participation makes more sense.

3B. Author's Response

Response to Reviewer D:

Thank you for your feedback and valuable suggestions. I implemented your suggestions in several places.

2. The Sentiment Protocol (added a paragraph)

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

Example 4 (added the word "positive" in the title to make the distinction more clear)

Example 5 (added the word "positive" in the title and added short incentive analysis in the end)

[...] 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 (added "with penalties" in the title and added short incentive analysis in the end)

> [...] 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.

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113

3. Incentives and vulnerability analysis (added gambling)

> [..] We distinguish between public polls, polls among experts and *gambling*.[..]
>
> 3.4 Gambling (new paragraph; since prediction markets have been discussed at length in the references, I kept this very short)
>
> 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.

I hope that these changes sufficiently address Reviewer D's comments.

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

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

ISSN 2379-5980 (online)
Associated article DOI
10.5915/LEDGER.2018.113