ISSN 2379-5980 (online)
associated article DOI 10.5915/LEDGER.2019.166

Decentralized Common Knowledge Oracles: Open Review

Austin K. Williams,[*†] Jack Peterson[‡]

Reviewers: Reviewer A, Reviewer B

Abstract. The final version of the paper "Decentralized Common Knowledge Oracles" can be found in Ledger Vol. 4 (2019) 157-190, DOI 10.5915/LEDGER.2019.166. There were two reviewers involved in the review process, neither of whom have requested to waive their anonymity at present, and are thus listed as A and B. After initial review by Reviewers A and B, the submission was returned to the authors with feedback for revision (1A). The authors responded (1B) and resubmitted their work. It was once again sent to Reviewers A and B, who indicated that the revisions made were sufficient to address their concerns, thus ending the peer review process. Author responses are bulleted for clarity.

1A. Review

Reviewer A

The paper "Decentralized common knowledge oracles" describes and mathematically proofs the design of an oracle algorithm with an economic motivation of returning true results. As well the mathematical and game-theoretical introduction as the stepwise derivation of the optimal oracle algorithm are not easy but generally understandable by a mathematically sound reader. All equations, symbols and sets are explained and - where ever necessary - colloquially expressed repeatedly. Starting with a single round oracle algorithm up to an multiple round algorithm with the possibility of disputes and without the need of a-posteriori-knowledge all single steps are explained with figures, pseudo-code algorithms and examples. As a reviewer with sound mathematical knowledge but without deep background in game theory I can understand the correctness of the assumptions without being fully able to confirm the overall correctness of all (game) theories. Generally I can recommend the publication of this paper without changes. It seems to be a bit lengthy and does not directly invite for reading 24 pages of content. But any reduction or compression risks to also reduce the possibility of understanding the well described mathematical considerations. The only improvement - if really necessary - considers the theorems 5.1 and 5.2. Athough being proven in the annex, it is equally hard to understand as well the theorems as the proofs. If the figure 4

[*] 3NubPpzpE3DQm9g6XxkueB7GEjuDZDewCF
[†] A. K. Williams (austin.williams@onewayfunction.com) is a researcher at OpenZeppelin.
[‡] J. Peterson (jack@tinybike.net) is a researcher at INDY Labs, USA.

and the table 5 were a little bit more explained, especially regarding the terms in the brackets (fig. 4) and the two terms in each cell of table 5, this could be helpful for the understanding. Nevertheless, I already yet recommend to accept the paper.

Reviewer B

The paper introduces a new (or rather three new) mechanisms for solving the problem of honest decentralized oracles. This problem is a huge limiting factor for interactions between blockchains and out-of-blockchain (=real world) events. Thus, any contribution that helps to overcome this issue is immensely important for a wide-spread adoption of Blockchain technology. The authors promise to deliver such a solution.

The idea presented by the authors to provide honest decentralized oracles is clever: a distinction between what the authors call reporters and queriers of oracles. Repoters can still cause an oracle to potentially lie about an event but have the economic incentive not to, because the queriers have the power to punish dishonest reporters by devaluing any paid out tokens to zero.

The curious reader would expect the description of this idea as the focal point of the paper to be mentioned in the abstract, the introduction and the conclusion. However, this fundamental idea is not well-carved out and presented to the reader. Instead, it is still very vague in the abstract and introduction and the reader has to wait until section 4. It is highly critical to carve out and describe this idea of a distinction between reporters and queriers. Otherwise readers will stop reading the paper after the abstract or the introduction at the latest.

The economics and game theory within the described mechanisms are sound and valid. However, there are two major limitations to these mechanisms that will likely reduce the real world applicability significantly. (1) It is implicitly assumed that reporters cannot be queriers or at least these two groups do not overlap strongly. The reviewer cannot see why these groups should not be identical and a malicious agent or group of agents control the majority of reporters and queriers. In that case, the punishment option will not help to incentivize honesty. (2) Effects of the on-chain token payout for reporters are not considered in the "economic soundness condition", although these are even discussed later as "open interest" on a betting market (it might be that the authors implicitly assume that the benefits of winning a bet due to having made an oracle misreport an event always comes in tokens that can be rendered valueless by queriers – but this is neither stated nor realistic). Such effects could still render an attack economically viable. It is the old problem that arises when derivative markets become larger than their underlying market like in sports betting for low-class leagues (this problem is the major limitation to Augur, the company the authors work for).

The basic idea of the paper is very interesting and helpful for the debate. The reviewer thus suggests that paper can potentially be published, but only after a strong revision. In this revision (1) the distinction between reporter and querier needs to be presented as the core of the paper and (2) the limitations of the model are clearly spelled out and either discussed and/or debunked. At the same time, the authors could think of only presenting one mechanism instead of three (and effectively omit sections 7-10) and rather present the other mechanisms as an add-on in a follow up paper. This would give the paper a clearer focus and stronger focus but is not a must.

Two small things the reviewer stumbled upon: (1) why is the absolute value of R and T used in the pay definition? Can there be a negative number of tokens? If yes, how? (2) On page 10 it might need to be "weakly dominant" instead of "weakly dominate".

1B. Author Response

Reviewer A

It seems to be a bit lengthy and does not directly invite for reading 24 pages of content. But any reduction or compression risks to also reduce the possibility of understanding the well described mathematical considerations.

- We agree. It was a tough choice between "reductive understanding" and brevity. We opted for the former since this is the first formal paper on this approach and we want future researchers to fully understand the reasons for each component of the final mechanism. We hope that the revisions we made to the introduction will help this be a little more inviting.

The only improvement - if really necessary - considers the theorems 5.1 and 5.2. Although being proven in the annex, it is equally hard to understand as well the theorems as the proofs. If the figure 4 and the table 5 were a little bit more explained, especially regarding the terms in the brackets (fig. 4) and the two terms in each cell of table 5, this could be helpful for the understanding. Nevertheless, I already yet recommend to accept the paper.

- We appreciate this feedback. The figures mentioned are game trees, decision trees, and game representations in normal form. These are very common in game theory literature but are less common outside the field of game theory. We revised the wording in the sections referencing those figures, the figure descriptions, and the introduction to the theorems in order to help non-game-theorist readers have a better sense of what they are looking at. There is a certain minimum game theoretical background needed to understand the meaning of the theorems and to follow the proofs. While we think that providing that background is outside the scope of the paper, we agree that non-game-theorist readers would benefit from some better descriptions of the figures. We revised them accordingly.

Reviewer B

The idea presented by the authors to provide honest decentralized oracles is clever: a distinction between what the authors call reporters and queriers of oracles. Reporters can still cause an oracle to potentially lie about an event but have the economic incentive not to, because the queriers have the power to punish dishonest reporters by devaluing any paid out tokens to zero.

The curious reader would expect the description of this idea as the focal point of the paper to be mentioned in the abstract, the introduction and the conclusion. However, this fundamental idea is not well-carved out and presented to the reader. Instead, it is still very vague in the abstract and introduction and the reader has to wait until section 4. It is highly critical to carve out and describe this idea of a distinction between reporters and queriers. Otherwise readers will stop reading the paper after the abstract or the introduction at the latest.

- We appreciate this recommendation. We heavily revised the introduction section to put this querier/reporter distinction front and center, and also to present our overall approach to the reader early on. We similarly revised the abstract.

The economics and game theory within the described mechanisms are sound and valid. However, there are two major limitations to these mechanisms that will likely reduce the real world applicability significantly. (1) It is implicitly assumed that reporters cannot be queriers or at least these two groups do not overlap strongly. The reviewer cannot see why these groups should not be identical and a malicious agent or group of agents control the majority of reporters and queriers. In that case, the punishment option will not help to incentivize honesty.

- When performing our analysis in the non-cooperative model we do implicitly assume that the querier is not a reporter. This simplifies the analysis in the non-cooperative model. We agree that this assumption should be made explicit.

- The case where reporters and queriers can form coalitions is handled in section 11 "Incentive Compatibility in the Cooperative Model". Note that this also covers the case where the querier *is* a reporter, because the utility of a "querier-reporter" is the sum of the utilities of the querier role and the reporter role. So a single querier-reporter is equivalent (with respect to utility) to a coalition consisting of the querier and a separate reporter.

- However, none of this was made clear in the paper.

- To rectify this we've revised the introduction to point this out. We also revised section 3 ("Assumptions") to point out that, for the majority of the paper, we do assume that the queriers and the reporters are disjoint sets, and that the case where the querier is a reporter (or, more generally, that queriers and reporters form a coalition) is handled in section 11. Finally, we added a note to the introduction of section 11 noting why the case where the querier is a reporter is covered by the cooperative model analysis.

(2) Effects of the on-chain token payout for reporters are not considered in the "economic soundness condition", although these are even discussed later as "open interest" on a betting market (it might be that the authors implicitly assume that the benefits of winning a bet due to having made an oracle misreport an event always comes in tokens that can be rendered valueless by queriers – but this is neither stated nor realistic).

- We do not assume that the benefits of winning a bet due to having made the oracle misreport would be measured only in reporting tokens. We agree with Reviewer B that such an assumption would not be realistic.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2019.166

- The value `I` is intended to capture all "extraneous benefit" of making the oracle lie. This includes, for example, all open interest on a betting market that uses the oracle. It also includes any secondary bets (or derivatives) made against those primary bets. Indeed, we expect that almost all benefit of making the oracle lie will be gained in tokens that cannot be made valueless. The fact that this wasn't clear to Reviewer B indicates to us that we did not present this idea well when introducing the economic soundness condition in section 5.2.

- To address this shortcoming we've revised section 5.2 to make it clearer for readers that the value `I` is intended to capture *all* benefit -- including any "extraneous" benefit -- of making the oracle lie.

Such effects could still render an attack economically viable. It is the old problem that arises when derivative markets become larger than their underlying market like in sports betting for low-class leagues (this problem is the major limitation to Augur, the company the authors work for).

- This is correct, and it is indeed one of the major limitations of not just this approach, but all public oracles. The derivatives markets becoming larger than their underlying markets is a great example of this. Another example is a centralized oracle that is being used to control an amount of open interest that is greater than the cost of bribing/hacking/coercing the centralized oracle to lie.

- One way to express this limitation of our approach is to say that there exists an inherent limit on how harshly future queriers are able to punish lying reporters, but there is no inherent limit on how much benefit a lying reporter can gain.

- To help make this limitation clearer to readers, we have revised section 5.6 ("Weaknesses") to explicitly point out this limitation. We also made a slight modification to section 5.2 (in the paragraph that colloquially explains the meaning of the Economic Soundness Condition) to be sure readers are aware of this limitation.

The basic idea of the paper is very interesting and helpful for the debate. The reviewer thus suggests that paper can potentially be published, but only after a strong revision. In this revision (1) the distinction between reporter and querier needs to be presented as the core of the paper

- As noted above, we took this into consideration and revised the abstract, introduction, and assumptions section of the paper.

and (2) the limitations of the model are clearly spelled out and either discussed and/or debunked.

ISSN 2379-5980 (online)
associated article DOI
10.5915/LEDGER.2019.166

- As noted above, we took this into consideration and revised section 5.6 ("Weaknesses") and section 5.2 accordingly.

At the same time, the authors could think of only presenting one mechanism instead of three (and effectively omit sections 7-10) and rather present the other mechanisms as an add-on in a follow up paper. This would give the paper a clearer focus and stronger focus but is not a must.

- This was a tough decision. On the one hand we agree that presenting a single mechanism would make the paper more succinct and easier to read. On the other hand, we think it will be instructive to future researchers to see all three mechanisms, because that makes it clear which components of the final mechanism give rise to which desirable properties and which additional necessary conditions. We think this will help future researchers get a better reductive understanding of the principles upon which this approach is built.

- Since this is the first formal paper discussing this approach, we think "reductive understanding" provides a greater benefit to the research community than would brevity (at least in this case). For that reason, we prefer to keep all three mechanisms in the paper. (We do think that any future published work can safely omit the first two mechanisms).

Two small things the reviewer stumbled upon: (1) why is the absolute value of R and T used in the pay definition? Can there be a negative number of tokens? If yes, how?

- Since R and T are sets (not numbers) the notation |R| and |T| refers to the cardinality the sets R and T (that is, the number of elements in the sets). If R and T were numbers instead of sets, the notation |R| and |T| would denote absolute value.

- This notation is standard; we think this was just a small misunderstanding on the part of the reviewer (simply mistaking R and T to be numbers rather than sets) so we've made no revisions of the notation.

(2) On page 10 it might need to be “weakly dominant” instead of “weakly dominate”.

- Thank you. We corrected this during the revision.

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.2019.166