

A Contribution to the Defense of Liquid Democracy

GREGORY BUTTERWORTH* and RICHARD BOOTH, Cardiff University School of Computer Science and Informatics, UK

Liquid democracy is a hybrid direct-representative decision-making process that provides each voter with the option of either voting directly or to delegate their vote to another voter, i.e., to a representative of their choice. One of the proposed advantages of liquid democracy is that, in general, it is *assumed* that voters will delegate their vote to others that are better informed, which leads to more informed and better decisions. Considering an audience from various knowledge domains, we provide an accessible high-level analysis of a prominent critique of liquid democracy by Caragiannis and Micha. Caragiannis and Micha's critique contains three central topics: 1. Analysis using their α -delegation model, which does not assume delegation to the more informed; 2. Novel delegation network structures where it is advantageous to delegate to the less informed rather than the more informed; and 3. Due to NP hardness, the implied impracticability of a social network obtaining an optimal delegation structure. We show that in the real world, Caragiannis and Micha's critique of liquid democracy has little or no relevance. Respectively, our critique is based on: 1. The identification of incorrect α -delegation model assumptions; 2. A lack of novel delegation structures and their effect in a real-world implementation of liquid democracy, which would be guaranteed with constraints that sensibly distribute voting power; and 3. The irrelevance of an optimal delegation structure if the correct result is guaranteed regardless. We conclude that Caragiannis and Micha's critique has no significant negative relevance to the proposition of liquid democracy.

Additional Key Words and Phrases: liquid democracy, delegative democracy, representative democracy

ACM Reference Format:

Gregory Butterworth and Richard Booth. 2023. A Contribution to the Defense of Liquid Democracy. In *Proceedings of 24th Annual International Conference on Digital Government Research (dg.o 2023)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The UK political system of *representative* democracy remains essentially the same as it was when *political representatives*, i.e., MP's (Members of Parliament), had to travel to the House of Commons by horse. In the UK and many other parts of the world, the numerous opportunities for technology to enhance the representative democratic process has not gone unnoticed.

For example, *liquid democracy* [4, 5, 16, 17] provides an opportunity to utilize the Internet to enable mass democratic participation. Notably, for 'voters' that do not have the time to participate, and/or know somebody that is better informed, participation is enhanced with the provision of an option to delegate their vote. It should be noted that one of the proposed and *assumed advantages* of liquid democracy is that, in general, voters will delegate to others that are better informed, which leads to more informed and better decisions. Liquid democracy may be implemented via e.g., *liquid feedback* [3], which enables the electorate to create polls, deliberate, delegate their vote or vote directly, and then potentially have their elected representative vote the same way e.g., in Parliament. Therefore, liquid democracy provides an opportunity to introduce an element of *direct democracy* into a *representative democracy* system to form a

*Gregory Butterworth is partially funded by UKRI.

modern *direct-representative hybrid* democracy. Notable examples of political parties that implement liquid democracy include the Pirate Party in Germany, Flux in Australia, and the Net Party (Partido de la Red) in Argentina.

However, the acceptability and uptake of the liquid democracy inspired direct-representative hybrid democracy proposal is hampered by various theoretical apprehensions. Recent research has explored such apprehensions and in some cases offered potential solutions. Apprehensions include, but are not limited to: the power of users that have amassed a large number of delegated votes [e.g. 5, 12]; the failure of assumed delegation to the more informed to collectively arrive at better decisions than direct or representative democracy [e.g. 6, 10]; voting security [e.g. 1, 15]; the practicalities of mass deliberation [e.g. 13]; etc. Therefore, the widespread adoption of liquid democracy is hampered until it can allay such concerns and can be shown to be the most obviously preferable alternative to the democratic system that it intends to replace.

From the above apprehensions, we focus on the subject of liquid democracy’s assumed delegation to the more informed to collectively arrive at a better decision than direct democracy. This subject was notably analyzed by Kahng, Mackenzie, and Procaccia [10, 11].

Kahng et al., investigate if there are ‘*delegation mechanisms that are guaranteed to yield more accurate decisions than direct voting?*’, and conclude negatively, but point out that a non-local, i.e., centralized mechanism, that imposes a cap on the voting power of any delegate, will adequately distribute power so that liquid democracy is at least as accurate as direct voting. Kahng et al. introduce their delegation mechanism, which includes the constant α to denote the minimum assumed increase in voter competence level when an agent delegates their vote to another agent. This recent work was then critically analyzed by Caragiannis and Micha.

Caragiannis and Micha critically analyze and expand on the work of Kahng et al. with their: ‘*A Contribution to the Critique of Liquid Democracy*’ [6], which is apprehensive of liquid democracy when compared to direct democracy. Their critique of liquid democracy focuses on three main areas: 1. Challenging liquid democracy’s assumption of delegation to the more informed, and a critique based on their proposed alternative α -delegation model of vote delegation, where delegation confirms intuition; 2. Providing novel delegation network structures that benefit from delegation to the less informed; and 3. Suggesting due to the theoretical NP hardness of finding an optimal delegation network structure, that real-world implementations would also be flawed due to their assumed inability to also discover such an optimal delegation network structure.

We suggest that in the real world, Caragiannis and Micha’s critique of liquid democracy has little or no relevance. Respectively, our critique is based on: 1. The identification of incorrect α -delegation model assumptions. Specifically, that intuition is not the dominant factor that determines a delegation choice. Rather, delegation is determined predominantly by trust. Most importantly, most delegation choices fall outside the scope of α -delegation because they are predominantly of a preset nature and involve no issue knowledge or associated ‘intuition’, and are based on reputation and trust; 2. A lack of novel delegation structures in a real-world delegation network, and the negation of any significant negative effect by implementing constraints that sensibly distribute voting power; and 3. The irrelevance of an optimal delegation structure if the correct result is guaranteed regardless.

1.1 Contribution

Caragiannis and Micha’s prominent ‘*A Contribution to the Critique of Liquid Democracy*’ [6], is cited as a valid critique of liquid democracy. Our contribution is an accessible analysis that reveals that their critique has, arguably, little to no significance in the real world.

1.2 Related Work

Our work is a direct response to the Caragiannis and Micha paper and aims to counter all their conclusions that are negative towards liquid democracy. To the best of our knowledge, there is no other paper that attempts to make such a complete, accessible, and high-level counterargument to their conclusions. Therefore, our work takes on a more philosophical nature, and has a focus of relating theory to the real-world, with the aim of creating a fuller understanding of the relevance of Caragiannis and Micha’s conclusions, and the discussion of the proposition of liquid democracy in general. To obtain the widest possible audience we deliberately avoid any demanding rigorous mathematical analysis.

Our work is most closely related to the very recent, independently conducted, work of Halpern et al. [9], which also aims to counteract Caragiannis and Micha’s negative perspective of liquid democracy. We consider our work to be complementary. Halpern et al., although also taking a real-world perspective, take a mathematical focus on the potential for delegated voting power to be distributed to too small a number of voters. Their mathematical delegation model is focused on other areas than ours.

We discuss at a general high level, issues surrounding Caragiannis and Micha’s ODP (optimal delegation problem). At a more formal and mathematical level, the ODP topic has recently been investigated by Becker et al. [2], and Zhang and Grossi [18].

2 PRELIMINARIES

We consider elections with a *ground truth*, i.e., with two voting options, one of which is objectively correct. There is a social network consisting of a set N of n voters, which are nodes connected by directed edges E on a graph G , stated as $G(N, E)$. Voters are also called agents. Agents may vote directly or delegate their vote to a connected agent. Delegations may be transitive but do not form cycles. Agents that both receive delegated votes *and* vote directly are often referred to as being a *guru* or *expert*. A guru or expert that has acquired a large number of votes is often referred to as being a *super-voter*. The election result is calculated using weighted majority, with each guru having a weight equal to one plus the total number of delegations it receives. Each agent i has a competence level p_i that indicates their probability of voting correctly. It is assumed there is some fixed parameter α such that agent i may only delegate to agent j if $p_i + \alpha < p_j$. This assumption is questioned by Caragiannis and Micha, leading them to propose their refined model, i.e., α -delegation, which is described in the next section.

We use colloquial concepts of *reputation* and *trust*. Based on an agent’s observation of another agent’s past behaviour or performance, *reputation* is defined as the subjective qualitative belief that the observing agent has regarding the observed agent. *Trust* is defined as a measure of how much an agent (trustor) is willing to rely on the actions of another agent (trustee), and the situation is directed to the future. The quality of the reputation that the trustor has of the trustee, determines the amount of trust that the trustor has in the trustee.

3 THE REAL-WORLD RELEVANCE OF THE ‘INEFFECTIVENESS OF LOCAL MECHANISMS’ AND THE CONCEPT OF α -DELEGATIONS

Caragiannis and Micha [6] analyze the liquid democracy delegation model proposed by Kahng et al. [10], which includes a constant parameter, α , so that delegations from an agent i to a neighboring agent j can only occur if $p_j > p_i + \alpha$. Caragiannis and Micha object to liquid democracy’s assumed delegation to the more informed, i.e., to those ‘*neighbors that have strictly higher competence level*’.

Caragiannis and Micha's main objection to the assumed delegation to the more informed, is based around their suggestion that it is unlikely that agents with low competence levels, e.g., $p_i = 10\%$, would delegate to those with high competence levels, e.g., $p_j = 70\%$, and we assume *vice versa*. They discount this possibility on the grounds that, because a competence level can be interpreted as an indication of an agent's confidence in their choice - '*A competence level of 10% for agent i can be interpreted as a 90% confidence*' - it follows that agent i would not delegate to agent j , because agent i would assume agent j is '*highly misinformed*'. Caragiannis and Micha then go on to refine the assumed liquid democracy delegation model by adding constraints that force delegations so that '*a voter can choose to delegate to someone who confirms her intuition and looks more expert.*', i.e., '*We refine the sufficient condition for a delegation from agent i to agent j to be $p_j > p_i + \alpha$ if $p_i \geq \frac{1}{2}$ and $p_j < p_i - \alpha$ if $p_i \leq \frac{1}{2}$* ', and they then coin the term ' α -delegation' to describe delegations that follow their adapted model.

In essence, Caragiannis and Micha's adaptation of Kahng et al.'s model essentially forces *correct* agents, i.e., $p_i > 0.5$, to delegate to *more correct* agents, i.e., $p_j > p_i + \alpha$, and forces *incorrect* agents, i.e., $p_i < 0.5$, to delegate to *more incorrect* agents, i.e., $p_j < p_i - \alpha$. We differentiate and refer to these two forms as respectively *positive- α -delegation*, or *+ α -delegation*, and *negative- α -delegation*, or *- α -delegation*.

3.1 Incorrect assumptions of the α -delegation model

We suggest that Caragiannis and Micha are incorrect when they assume that delegation agent choice is determined by confirmation of intuition. Instead, we suggest that trust is the dominant determinant on delegation agent choice. Therefore, the α -delegation model, which is based on delegation that confirms intuition, has little to no relevance to delegation within a real-world implementation of liquid democracy.

Let us remind ourselves that in a scenario of direct democracy it is neither desirable nor practical to have the entire electorate deliberate on every issue. Liquid democracy's *raison d'être* is its ability to circumvent this problem by offering a delegation option to those voters that do not have the time to participate in the democratic process. Therefore, by design, liquid democracy referendums will often have the vast majority of votes cast *automatically* and therefore *passively* via *preset* delegation options. It is important to note that these passive automatic delegations will have been set in advance of future issues, and with no awareness of any of those issue's specific details. Note that, automatic passive preset delegations are distinct from *manual active dynamic* delegations made *manually* during *active* participation in the deliberation phase of an issue: consider an agent that during the process of deliberation *dynamically* decides to either delegate their vote rather than vote directly, or decides to redelegate their vote to another delegate that they think is preferable to their current delegate.

Having supplied some context, our next point of analysis requires an estimate for the percentage of votes that are made via automatic preset delegations.

However, we are not aware of a resource that provides such data for a real-world large-scale, demographically representative, mature and stable liquid democracy implementation – in fact we are not aware of such an idealized real-world implementation of liquid democracy. Regardless, observing the relatively small scale German Pirate Party via analysis by Kling et al. [12] and the German Pirate Party themselves [7], we can surmise and extrapolate an approximation. Kling et al.'s extensive analysis of the German Pirate Party focuses on the power of super-voters, and as a result does include a suitable temporal analysis of the ratio of direct to delegated votes, or what percentage of the delegated votes are preset or active. But, Kling et al. do show that over time, users continue to make more delegations, and also that there is a reduction in the number of active users to approximately 10% of the total number of registered users: users are labelled inactive after 180 days without login. Also, considering the section of the German Pirate

Party’s statistics that are based on their ‘01/28/2014’ data [7], there is a general trend of an increasing proportion of delegated votes to approximately 80% of the total vote. It should be pointed out that the German Pirate Party data does not distinguish between the delegated votes being preset or active, but, considering that users would have to login to actively delegate their vote, and there is a reduction of active logins to approximately 10% of all users, it seems reasonable to assume that the vast majority of delegated votes are preset delegated votes. Taken together, this seems to suggest there is a trend of increasing preset delegations, which would not be unexpected considering that this is exactly how liquid democracy voting behaviour is expected to evolve, i.e., over time, voters will increasingly identify other voters that they are willing to trust with their preset delegated vote.

We suggest that this trend should be extrapolated to the point that, in a stable implementation of liquid democracy, the number of delegations would on average make up approximately 90% of the vote, and at least 95% of these delegated votes would be preset. We also suggest this is a reasonable figure for a large-scale implementation of liquid democracy, because this would represent a realistic design target of 10% of the agents being actively involved in the thorough deliberation of an issue, and/or voting directly.

The German Pirate Party statistics, based on their ‘08/11/2012’ data, show that the ratio of direct to delegated votes varies widely, from approximately 1:4 to 4:1, or put another way, the percentage of direct and delegated votes both vary from approximately 20% to 80%. This variation is expected because some issues are much more important or interesting than others. Obviously, when an issue has minimal importance, the percentage of the agents being actively involved in deliberation and/or voting directly would be lower than a 10% design target, and when an issue is very important with very wide-ranging effects, the percentage of the agents being actively involved in deliberation and/or voting directly would be much higher than 10%. It is conceded that in cases where a large proportion of the agents are involved in the deliberation process, and there is a high proportion of direct voting, it seems probable that this scenario is conducive to more dynamic delegation. However, in general, we suggest that for a stable large-scale implementation of liquid democracy, an average of approximately 90% of the delegated votes will be preset.

Having, hopefully, justified an approximate average automatic preset delegation rate of 90%, we now make the key point that these preset delegated votes, which are made in advance of an issue with the aim of avoiding issue awareness and deliberation, fall outside of the α -delegation model because these voters cannot be attributed to have any sense of issue intuition. In such cases of preset delegation, voters will have no voting intuition and will have solely made their delegation choices previously, based on reputation and trust.

For the active delegating voters, delegation would more often than not, also be decided on reputation and trust, rather than exclusively via confirming intuition. Caragiannis and Micha dismiss the possibility of a very incorrect agent delegating to a very correct agent, because the very incorrect agent would consider the very correct agent to be ‘*highly misinformed*’. Instead, they suggest delegation would be based on confirming intuition. Crucially, Caragiannis and Micha ignore or dismiss that the delegating influence of being ‘*highly misinformed*’ could be overcome by trust. We suggest that a very incorrect agent could very feasibly delegate to a very correct agent if there was a strong enough trust relationship that overcame the influence of the appearance of being ‘*highly misinformed*’. In such a case it is conceded that there is a requirement of a strong relationship of trust. For example, consider an issue with a ground truth that relies on the understanding of a complex mathematical issue, and three socially connected agents with an associated p_i measure: an undergraduate Computer Science student, with $p_i = 0.1$; a Mathematics PhD student, with $p_i = 0.3$; and a Mathematics Professor, with $p_i = 0.9$. Now, in this scenario, if the Mathematics PhD student decides to delegate their vote, to whom would they delegate? We suggest, that based on trust and reputation, the Mathematics PhD student would be more likely to delegate to the Mathematics Professor rather than the undergraduate Computer

Science student. Also, we suggest that the undergraduate Computer Science student should also be aware of their own relative shortcomings and it would seem logical that if they decided to delegate their vote it would also be to the Mathematics Professor.

However, α -delegation excludes the possibility of an incorrect agent ever delegating to a slightly more correct agent. But in these cases, delegation to the ‘*slightly more correct*’ would now appear as a delegation to the ‘*slightly misinformed*’. In these cases, delegation to a more correct agent is much more likely because there is a significantly reduced requirement for the level of trust to have a dominant influence on delegation.

Considering such negative- α -delegation and regardless of intuition, for any delegation from the incorrect to the more incorrect to occur, there must be an element of trust, because otherwise an agent would not delegate their vote to an agent that they do not trust. This poses the question, how did the agent that is more incorrect acquire this trust? It is conceded that in a real-world liquid democracy implementation that has just been created, the social network would not exclusively contain merit-based relationships of trust based on historical reputation. Agents that claim to be experts could be charlatans. Therefore, it is possible that trust could be assumed without any merit, which is conducive to an amount of negative- α -delegation. However, over time, the liquid democracy social network will review past poll outcomes and the associated votes of the supposed experts. This will lead to the identification of those experts who can and cannot be trusted. Therefore, over time a social network will learn who are the true experts and who are the charlatans, and will correctly apportion trust to deserving ‘*correct*’ agents. It follows that over time, the number of incorrect agents that are also trusted will be significantly reduced, which will significantly reduce the likelihood of negative- α -delegation.

In summary, we suggest that trust is the dominant influence on delegation. Considering that 90% of delegations are preset and outside the scope of α -delegation, and that we suggest a similar percentage of the remaining 10% of delegations would also be based on trust, then we approximate that 99% of delegations are not applicable to α -delegation. Therefore, any α -delegation based critical analysis of liquid democracy is weakened by a similar amount, essentially rendering any α -delegation based critique to insignificance.

4 THE REAL-WORLD RELEVANCE OF ‘*THE CURSE OF α -DELEGATIONS*’

We next evaluate Caragiannis and Micha’s second objection to α -delegation, which they consider to be significantly more important, i.e., ‘*The Curse of α -Delegations*’, where they present structures where ‘*delegating to less-informed agents is highly beneficial*’.

‘*Example 1*’, provides what is on the surface, a compelling critique that contradicts a promoted assumed advantage of liquid democracy delegative voting, i.e., that delegative voting involves delegating to more competent voters, which results in more informed decision making with a higher chance of ground truth discovery. Caragiannis and Micha show that structures exist where delegating to less-informed agents is highly beneficial. The structure example consists of an odd number of agents, split into a minimal majority forming a network star shape, with the central agent having a competence level slightly less than the leaf agents, and the remaining maximal minority of agents being isolated. They state this as follows:

‘*Example 1. Consider a set N of n agents (with odd n), connected through a social network $G(N, E)$ as follows: $(n - 1)/2$ agents are isolated while the remaining ones form a star. The competence levels p are $p_c = 1 - 2\epsilon$ for the center and $p_l = 1 - \epsilon$ for each leaf agent, where $\epsilon > 0$ is a negligibly small constant. The isolated nodes have competence level $p_i = 0$.*’

Figure 1 illustrates such a structure with $N = 7$. Note the dotted arrows show potential delegations. The associated node values represent competence values, p_i , that predict the probability of the voter voting for the issue’s ground truth.




Fig. 1. Beneficial less informed delegation.

Caragiannis and Micha then consider the scenario in which *only delegations to more informed agents* are allowed and observe that this structure provides zero possible options for delegation to the more informed. Therefore, all agents must vote directly. It follows that a correct result is only possible if all star agents vote correctly, which is increasingly unlikely and approaches zero as N grows, i.e., more and more likely for a single star agent to vote incorrectly and tip the balance to an incorrect result. Caragiannis and Micha, observe that when the external nodes (4, 6, and 7, on the right in figure 1) delegate to the lower competence value central node (5), there is a higher probability of ground truth discovery, compared to when all these nodes vote directly. Therefore, ‘*delegating to less-informed agents is highly beneficial*’. A stinging attack indeed.

However, stepping back into the real world, some key observations can be made.

Considering that there is only a problem if a poll result does not discover the ground truth, it is natural that analysis includes scenarios where the vote is in the balance, such as illustrated in figure 1.

Taking a general perspective of Caragiannis and Micha’s observation, their observation has an effect when the *loss of accuracy* caused by agents *delegating to a less informed agent* is *less damaging* than the potential for a *single direct voting agent to vote incorrectly* and swing the result to an *incorrect* one, which is more likely as N is increased. But how often would a real-world scenario relevant to their observation occur?

Within a real-world implementation of liquid democracy, complete with sensible constraints, e.g. power distributed across numerous gurus, it is accepted that networks similar to the above structure could possibly exist, but only as a *small outlier number of subnetworks* within a much *larger network*. The rest of the network would consist of a much larger number of subnetworks, consisting of structures that are assumed to be typical of a liquid democracy implementation.




Fig. 2. Potential for a correct decision from a very incompetent majority.

Interestingly, the structure of figure 1, but with *inverted p_i* values, i.e., the isolated nodes having a competence level $p_i = 100$ etc., as illustrated in figure 2, exposes a desirable property of this structure - consisting of a *very competent*

slight minority and a *very incompetent slight majority* - which is that in the real-world, as ‘ n ’ increases so does the likelihood of direct voting by external nodes, and therefore of obtaining a *correct decision from a very incorrect majority*. If we assume that this type of subnetwork structure would appear, as an outlier, as often as the one cited by Caragiannis and Micha, then it follows that their effects would cancel each other out.

We next consider Caragiannis and Micha’s next example, ‘*Example 2*’, as illustrated in figure 3. Note, that the dashed arrows show the options for a possible delegation.




Fig. 3. Delegation network.

Caragiannis and Micha observe that considering all the possible delegation combinations, the possible competence levels for the network vary from 67.5% to 85%. We highlight the two most relevant combinations. Firstly, the highest probability occurs when agent 4 delegates to agent 2, and agent 5 delegates to agent 3, which highlights that ‘*delegating to the most informed agent possible does not always equate to the optimal chance of the network discovering a ground truth*’. Secondly, the combination in which agents *exclusively* delegate to agents with *higher competence* levels, i.e., agents 4 and 5 both delegate to agent 3, which results in a *lower network competence level* of 70%.

Consider that an implication of Condorcet’s jury theorem [14] is, assuming all voters have the same voting weight, and the average competence level of the voters is above 50%, that as the number of voters is increased, a correct vote becomes a certainty. Now, considering the general case, and making the above assumptions, Caragiannis and Micha’s above observation has an effect when the *loss of accuracy* caused by agents *delegating to a less informed agent* is *offset* by the *increase in accuracy* caused by shifting the network structure closer to one that resembles a jury, i.e., equal voting weights and all competence levels hopefully above 50%.

The relevance is that, in a real-world implementation of liquid democracy with sensible constraints, e.g., a minimum number of gurus with a voting weight limit, the network structure that Caragiannis and Micha describe would only exist at most as a small number of substructures in a much larger network. If the network structure was extended to represent a realistic number of voters, there would be no negative effect because the number of voters would become sufficient to ensure that a correct result is guaranteed via the Condorcet jury theorem.

In summary, it is our opinion that α -delegation and the presence of rogue local subnetworks have no significant practical relevance that is terminal to the real-world proposition of liquid democracy.

5 OPTIMAL DELEGATION PROBLEM (ODP)

Caragiannis and Micha state the following: ‘*We show that the optimal delegation problem (ODP) of maximizing the probability to find the ground truth by coordinating delegations is not only computationally hard but also hard to approximate*

with a substantial additive constant. The criticism by this complexity result should be clear: if it is hard to discover the ground truth in a coordinated way, why should the voters be expected to find it by acting independently?.

First, we make the general point that the proposition of liquid democracy does not fail because it cannot acquire an optimal chance of ground discovery via a perfect delegation structure, it fails if it cannot acquire a better chance of ground truth discovery than the incumbent system of democracy that it aims to replace, which is - if based on the motivations of the majority of real-world implementations of liquid democracy - representative democracy.

As mentioned in our ‘related work’ section, the ODP topic has recently been investigated in-depth by Becker *et al.* [2], and Zhang and Grossi [18]. Becker *et al.*, strengthen the existing hardness result, but counter this with an approximation result, which is dependent on assumptions relating to the competencies of the voters or the structure of the network. They also create simulations that show that on a large class of instances, liquid democracy outperforms direct democracy. Zhang and Grossi obtain an optimal group accuracy by central coordination, which uses ‘a variant of liquid democracy where proxies are weighted by a probability of delegation or, alternatively, by a share of the delegator’s voting weight’, but achieve this by dropping Caragiannis and Micha’s network constraints of restricted connectivity.

Considering the above, and that a sensible implementation of liquid democracy could utilize a centralized system to limit guru power, it seems the ODP problem is overblown.

Trivially, we make the point that, if the correct result is essentially guaranteed regardless of optimization, then optimization is not an issue. Specifically, if the result of obtaining the ground truth is essentially guaranteed regardless of the voters finding an optimal delegation voting structure, i.e., via the Condorcet jury theorem, when there are enough agents and gurus with competence levels above 50%, then the ODP is not an issue. The ODP would only be relevant if an optimal delegation structure was the difference between the ground truth being discovered or not, but how often would this be the case if the liquid democracy implementation included quality deliberation and sensible voting mechanism constraints?

Also, a non-optimal delegation structure provides an opportunity for the improvement of subsequent delegation structures. If we consider that gurus that voted *incorrectly* will be identified, and then over time will become less likely to receive delegated votes, then a non-optimal delegation structure should contribute to the *learning* of the social network and over time move it towards an optimal structure, or at least one that increasingly guarantees a correct result as the average competence levels of the gurus increase.

Importantly, considering that all agents are subject to the real-world constraints of time, personal knowledge, and personal networking, it is not realistic to suggest or expect that all agents were ever going to independently coordinate their efforts to investigate every single possible delegation combination to arrive at an optimal chance of ground truth discovery. The central crux of Caragiannis and Micha’s criticism by their complexity result is not relevant because it is based on a proposed process that will never be a requirement of any liquid democracy implementation.

From a real-world perspective, it could be said that any final delegation structure is optimal considering the accuracy is subject to the above constraints. Arguably, a delegation structure that has been derived totally independently is theoretically optimal from the perspective of trust: trust being acquired via the transparency and the ownership of an individual’s decision to vote directly or delegate to an agent of their choice. The value of trust in any democratic system should not be ignored or underestimated.

Based on all of the above points, we consider Caragiannis and Micha’s ODP, and criticism by their complexity result to have no negative bearing on the proposition of liquid democracy.

6 DISCUSSION

We make the general point that the negative effects that Caragiannis and Micha observe are associated with essentially small network structures that have power concentrated across a relatively small number of agents. Such network structures more closely resemble those found in representative democracies, e.g., UK government cabinet meetings [8], rather than those that would be found in large-scale liquid democracy implementations. We suggest that Caragiannis and Micha’s negative observations are arguably more applicable to representative democracy than liquid democracy.

Note that Caragiannis and Micha restrict their proofs in their *Sections 3 and 4* to proxy voting, i.e., non-transitive voting, and in *Theorem 1* highlight the effect of low values of α to reinforce the significance of their analysis. It is conceded that although a transitive delegation chain can be approximated by a series of proxy votes, it is not the same unless it accounts for the potential of multiple additions of α to be obtained via the delegation chain. Also, we point out that the value of α has a fundamental effect on the performance of liquid democracy, and unfair assumptions of the preponderance of low values of α , provide critics of liquid democracy with an opportunity that in the real world we assume would not exist.

7 CONCLUSION

First, we conclude that in a real-world scenario of liquid democracy, the amount of α -delegation is insignificant, because trust rather than intuition is the dominant influence over delegation agent choice. Secondly, network structures that are observed to benefit from delegation to the less informed would be restricted to substructures with a minimal effect if sensible power distributing constraints are implemented to ensure the effect of the Condorcet jury theorem. Lastly the ODP problem is overblown, because a liquid democracy system utilizing such previous constraints to ensure the effect of the Condorcet jury theorem, would guarantee a correct result regardless of an optimal delegation structure.

In summary, Caragiannis and Micha’s critique is restricted to a narrow theoretical perspective that is not valid in a sensible real-world implementation of liquid democracy. We conclude that Caragiannis and Micha’s critique is not deleterious to the real-world proposition of liquid democracy.

8 FUTURE WORK

This work is intended to serve as an introduction and motivation for a future rigorous and comprehensive temporal analysis of the delegation process in general.

REFERENCES

- [1] Ch Anwar ul Hassan, Muhammad Hammad, Jawaid Iqbal, Saddam Hussain, Syed Sajid Ullah, Hussain AlSalman, Mogeab A. A Mosleh, and Muhammad Arif. 2022. A Liquid Democracy Enabled Blockchain-Based Electronic Voting System. *Scientific programming* 2022 (2022).
- [2] Ruben Becker, Gianlorenzo D’Angelo, Esmail Delfaraz, and Hugo Gilbert. 2021. When Can Liquid Democracy Unveil the Truth? *CoRR* abs/2104.01828 (2021). arXiv:2104.01828 <https://arxiv.org/abs/2104.01828>
- [3] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. 2014. *The Principles of LiquidFeedback*. Interaktive Demokratie.
- [4] Jan Behrens, Axel Kistner, Andreas Nitsche, and Björn Swierczek. 2017. The origins of liquid democracy. *The Liquid Democracy Journal* 5 (2017), 2018–2024. <http://www.interaktive-demokratie.org/Allrightsreserved.http://www.liquid-democracy-journal.org/Archiveavailableat:http://www.liquid-democracy-journal.org/>
- [5] Christian Blum and Christina Isabel Zuber. 2016. Liquid Democracy: Potentials, Problems, and Perspectives. *Journal of Political Philosophy* 24, 2 (2016), 162–182. <https://doi.org/10.1111/jopp.12065> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/jopp.12065>
- [6] Ioannis Caragiannis and Evi Micha. 2019. A Contribution to the Critique of Liquid Democracy. In *Twenty-Eighth International Joint Conference on Artificial Intelligence*. 116–122. <https://doi.org/10.24963/ijcai.2019/17>
- [7] Piratenpartei Deutschland. 2014. Benutzer:SD/Delegationen in LQFB/Abstimmungen. https://wiki.piratenpartei.de/Benutzer:SD/Delegationen_in_LQFB/Abstimmungen#Abstimmungen_getrennt_in_eigene_Stimmen_und_delegierte_Stimmen Accessed: 2022-06-06.

- [8] Durrant, Tim. 2021. Cabinet | The Institute for Government. <https://www.instituteforgovernment.org.uk/explainers/cabinet>. Accessed: 2021-09-01.
- [9] Daniel Halpern, Joseph Y. Halpern, Ali Jadbabaie, Elchanan Mossel, Ariel D. Procaccia, and Manon Revel. 2021. In Defense of Fluid Democracy. arXiv:2107.11868 <http://arxiv.org/abs/2107.11868>
- [10] Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2018. Liquid Democracy: An Algorithmic Perspective. *32nd AAAI Conference on Artificial Intelligence (AAAI'18)* (2018).
- [11] Anson Kahng, Simon MacKenzie, and Ariel D. Procaccia. 2021. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research* 70 (2021), 1223–1252. <https://doi.org/10.1613/JAIR.1.12261>
- [12] Christoph Carl Kling, Jerome Kunegis, Heinrich Hartmann, Markus Strohmaier, and Steffen Staab. 2015. Voting Behaviour and Power in Online Democracy: A Study of LiquidFeedback in Germany's Pirate Party. *Proceedings of the International Conference on Weblogs and Social Media* (2015), 208–217. arXiv:1503.07723 <http://arxiv.org/abs/1503.07723>
- [13] Hélène Landemore. 2021. *Open democracy : reinventing popular rule for the twenty-first century*. Princeton University Press, Princeton.
- [14] M. J. A. N. de C. Marquis de Condorcet. 1785. *Essai sur l'Application de l'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. Imprimerie Royale, Paris.
- [15] Mahdi Nejadgholi, Nan Yang, and Jeremy Clark. 2021. Short Paper: Ballot Secrecy for Liquid Democracy. In *Financial Cryptography and Data Security. FC 2021 International Workshops*. Lecture Notes in Computer Science, Vol. 12676. Springer Berlin Heidelberg, Berlin, Heidelberg, 306–314.
- [16] Alois Paulin. 2020. An Overview of Ten Years of Liquid Democracy Research. In *The 21st Annual International Conference on Digital Government Research* (Seoul, Republic of Korea) (*dg.o '20*). Association for Computing Machinery, New York, NY, USA, 116–121. <https://doi.org/10.1145/3396956.3396963>
- [17] Chiara Valsangiacomo. 2022. Clarifying and defining the concept of liquid democracy. *Swiss Political Science Review = Schweizerische Zeitschrift für Politikwissenschaft* 28, 1 (March 2022), 61–80. <https://doi.org/10.1111/spsr.12486>
- [18] Yuzhe Zhang and Davide Grossi. 2021. Tracking Truth by Weighting Proxies in Liquid Democracy. arXiv:2103.09081 <https://arxiv.org/abs/2103.09081>