On the Potential and Limitations of Proxy Voting: Delegation with Incomplete Votes

Georgios Amanatidis
University of Essex,
Archimedes / Athena RC
g.d.amanatidis@gmail.com

Aris Filos-Ratsikas
University of Edinburgh
arisfilosratsikas@gmail.com

Philip Lazos
Input Output Global (IOG)
philip.lazos@iohk.io

Evangelos Markakis
Athens University of Economics and Business, Archimedes / Athena RC, Input Output Global (IOG)
markakis@gmail.com

Georgios Papasotiropoulos
Athens University of Economics and Business, Input Output Global (IOG), University of Warsaw
gpapasotiropoulos@gmail.com

ABSTRACT

We study elections where voters are faced with the challenge of expressing preferences over an extreme number of issues under consideration. This is largely motivated by emerging blockchain governance systems, which include voters with different weights and a massive number of community generated proposals. In such scenarios, it is natural to expect that voters will have incomplete preferences, as they may only be able to evaluate or be confident about a very small proportion of the alternatives. As a result, the election outcome may be significantly affected, leading to suboptimal decisions. Our central inquiry revolves around whether delegation of ballots to proxies possessing greater expertise or a more comprehensive understanding of the voters' preferences can lead to outcomes with higher legitimacy and enhanced voters' satisfaction in elections where voters submit incomplete preferences. To explore this, we introduce a model where potential proxies advertise their ballots over multiple issues, and each voter either delegates to a seemingly attractive proxy or casts a ballot directly. We identify necessary and sufficient conditions that could lead to a socially better outcome by leveraging the participation of proxies. We accompany our theoretical findings with experiments on instances derived from real datasets. Our results enhance the understanding of the power of delegation towards improving election outcomes.

KEYWORDS

Computational Social Choice; Proxy Voting; Incomplete Preferences

ACM Reference Format:
Georgios Amanatidis, Aris Filos-Ratsikas, Philip Lazos, Evangelos Markakis, and Georgios Papasotiropoulos. 2024. On the Potential and Limitations of Proxy Voting: Delegation with Incomplete Votes. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

1 INTRODUCTION

Broadly speaking, an election refers to a voting system in which a set of participants express their preferences over a set of possible issues or outcomes, and those are aggregated into a collective decision, typically with a socially desirable objective in mind. Besides their "traditional" applications such as parliamentary elections or referenda, elections often underpin the livelihood of modern systems such as blockchain governance [11, 24] or participatory budgeting [4, 9]. Quite often, voters are called to vote on an extremely high number of issues, rendering the accurate expression of their preferences extremely challenging. For instance, the Cardano blockchain uses Project Catalyst (https://projectcatalyst.io) to allocate treasury funds to community projects, and routinely receives more than a thousand of proposals per funding round. Another application comes from platforms of civic participation, where the users express support on opinions or proposals [18]. An unfortunate consequence of these election scenarios is that the voters inevitably have a confident opinion only for a small number of issues (henceforth *proposals*), as investing enough time and effort to inform themselves on such a sheer volume of proposals is clearly prohibitive. In turn, the "direct voting" outcome, even under the best intentions, is likely to be ineffective in capturing the desires of the society.

A well-documented possible remedy to this situation is to allow for *proxy voting* [12], a system in which the voters may *delegate* their votes to *proxies*. The idea is that those proxies have the time and resources to study the different proposals carefully, and vote on behalf of the voters they represent. This in fact captures voting applications more broadly, where the reason for delegation might be a reluctance to express an opinion, lack of specialized knowledge, or even limited interest. When those proxies are part of an electorate together with other voters and proxies, the resulting system is known as *liquid democracy* [10, 21]. Liquid democracy has been scrutinized, with arguments presented in its favor [6, 17, 31] and against it [10, 21]. At the same time, it is being employed in real-world situations [29], including settings similar to the one studied in the current work, like the Project Catalyst mentioned before.

A takeaway message from the ongoing debate around delegative voting is that such processes might indeed be useful under the right circumstances. Extending this line of thought and motivated by the scenarios presented above, in this work we aim to identify the *potential* and *limitations* of proxy voting with regard to achieving socially desirable outcomes in settings with incomplete votes, under the classical *Approval Voting* rule. More precisely, we aim to characterize what can be theoretically made possible through delegation, and what remains impossible, even under idealized conditions.

1.1 Setting and Contribution

We focus on elections in which the aim is to choose one proposal to be implemented from a range of multiple candidate proposals. We introduce a model of proxy voting, where voters have intrinsic approval preferences over all proposals, which are only partially revealed or known to the voters themselves. A set of *delegation representatives (dReps)* can then advertise ballots over the proposals and the voters in turn may either delegate to a proxy, if there is sufficient agreement (i.e., over a certain *agreement threshold* between dRep's advertised ballot and the voter's revealed preferences), or vote directly. The outcome of the election is the proposal with the highest approval score, assembled by the score from the ballots of the dReps (representing voters who delegated their vote) and the voters that vote directly. The core question we pose follows:

> *"Is it possible for the dReps to advertise their preferences appropriately such that the outcome of the election has an approval score that is a good approximation of the best possible approval score; which would be achievable if all voters had full knowledge of their preferences?"*

We study the question under a *"Best-Case Scenario"* that involves the following assumptions, and, shortly, we will elaborate on why this scenario is meaningful for deriving results in the general setting.

(1) The dReps are fully informed about the preferences of the voters, i.e., they know exactly the vector of intrinsic preferences for each voter.
(2) The dReps themselves do not have actual preferences and their only goal is to achieve the best possible approximation of the optimal approval score. To do so they can coordinate with each other and advertise their types accordingly.
(3) When several proposals are tied for the winning position, ties are broken in favor of those with the highest intrinsic score.

One should not of course expect all of these assumptions to apply in practice: we would expect the dReps to be only partially informed about voters' preferences (e.g., via some probabilistic model) and to exhibit some sort of rational behavior (e.g., needing to be appropriately incentivized to advertise ballots that are aligned with socially desirable outcomes). Still, studying the "best-case scenario" is already instructive for results in all other regimes. In particular:

- Our *negative results (inapproximability bounds)* immediately carry over to other settings as well, regardless of the choice of the dReps' information model, their rationality model, or the choice of the tie-breaking rule. In other words, we show that certain objectives are *impossible*, even when a set of fully-informed dReps coordinate to achieve the best outcome, hence they are certainly impossible for any other meaningful setting.
- Our *positive results (approximation guarantees)* establish the limits of the aforementioned impossibilities: if something is not deemed impossible by our bounds, it should be the starting point of investigations for an information and rationality model chosen for the application at hand. Clearly, if our upper bounds establish that a certain number of dReps suffices to achieve a certain approximation in the "best-case scenario", one should expect a slightly larger number of dReps to be needed in practice.

Our Results. We firstly present a strong impossibility, namely that for any agreement threshold higher than 50%, the best achievable approximation ratio is linear in the number of voters. On the positive side, we show that for an appropriate *coherence* notion of the instance, which captures the commonalities of the set of proposals that sets of voters are informed about, meaningful approximations are possible. For the natural case of an agreement threshold of 50%, we show that a single dRep is capable of achieving an approximation factor of 3, whereas only 2 dReps are sufficient to elect the optimal proposal. Most significantly, we present general theoretical upper and lower bounds on the possible approximations, depending on the agreement threshold, the number of dReps, and the coherence of the instance. We complement our theoretical results with a set of experiments using the MovieLens dataset [19], to measure the effects of proxy voting on realistic incomplete preferences.

1.2 Related Work

We first comment on some works that are closer in spirit to ours. Meir et al. [28] propose a model with a similar objective, but focusing on the analysis of *sortition*, i.e., the approximation of the welfare achieved by selecting a random small-size committee. In a related direction, Cohensius et al. [12] analyze particular delegation mechanisms, under elections with samples of voters located randomly in a metric space, according to some distribution. Our approach does not consider any randomization, neither for the voting rule nor for the preferences. Finally, Pivato and Soh [30] also consider the performance of proxy voting, focusing on understanding when the proxy-elected outcome coincides with the outcome of direct voting. Again, the model of Pivato and Soh [30] is randomized, where the voters delegate based on the probability of agreement to a proxy, and not based on a deterministic distance function. Moreover, no analysis of approximation guarantees is undertaken in the work of Pivato and Soh [30]. Our work can be seen as one that contributes to the corpus of findings in favor of proxy voting [6, 17, 31], albeit in a markedly different manner.

There is significant work within the field of computational social choice on elections with incomplete votes. One stream has focused on the identification of possible and necessary winners by exploring potential completions of incomplete profiles; see the work of Lang [26] for an overview. Recent work has concentrated on the computational complexity of winner determination under various voting rules within the framework of incomplete information [5, 20, 34]. Another direction has studied the complexity of centralized interventions to reduce uncertainty [1] (e.g., by educating a selected set of voters or computing delegations via a centralized algorithm). Furthermore, there has been an exploration of the effect of minimizing the amount of information communicated [3, 22] as well as of the interplay between voters' limited energy and social welfare [32]. Considerable attention has been devoted to the exploration of efficient extensions of incomplete profiles to complete ones that satisfy desirable properties [14, 25, 33]. A conceptually related area focuses on distortion in voting, investigating the implications of applying rules to preferences that are less refined than the voters' intrinsic preferences [2]. Beyond voting scenarios, similar solution concepts have been explored in domains such as fair division [7], hedonic games [23], and non-cooperative games [8].

2 FRAMEWORK AND DEFINITIONS

In the current section we formally describe the main attributes of the election setting we study.

Proposals and Voters. Let $C = \{1, 2, \dots, m\}$ be a set of candidate proposals for a single-winner election, where, evidently, for each proposal there are exactly two options: to be elected or not. Let also $N = \{1, 2, \dots, n\}$ be a set of voters responsible for determining the elected proposal. Each voter $i \in N$ is associated with *approval preferences* $v_i \in \{0, 1\}^m$ over the set of proposals; we refer those as true or *intrinsic* preferences. Here, 1 and 0 are interpreted as "accept" (or "support") and "reject" (or "oppose") a proposal, respectively.

Crucial to our model is the fact that voters do not actually know their entire intrinsic preference vector, but only a subset of it; this could be due to the fact that they have put additional effort into researching only certain proposals to verify if they indeed support them or not, but not necessarily all of them. Formally, we will say that each voter i has *revealed* preferences $\widehat{v_i} \in \{0, 1, \perp\}^m$, where $\perp$ denotes that the voter does not have an opinion on the corresponding proposal. As such, we have the following relations between v_i and $\widehat{v_i}$: $(\widehat{v_i}(j) = v_i(j)) \vee (\widehat{v_i}(j) = \perp), \forall j \in C$. The collection of proposals for which a voter i has developed an opinion is referred to as their *revealed set*, denoted by $R_i := \{j \in C : \widehat{v_i}(j) = v_i(j)\}$. Let $m_i := |R_i|$. Each voter i also has an integer weight w_i; in our work, it is without loss of generality to assume that $w_i = 1$ for all $i \in N$, as we can simply make w_i copies of voter i (with the same preferences), and all of our results go through verbatim.

dReps. A *delegation representative (dRep)* is a "special" voter capable of attracting voters to delegate their votes to her, who participates in the election with the combined weight of those voters. Consistently with the "best-case scenario" motivation (see Section 1.1), we view dReps as unweighted agents, devoid of personal preferences over the proposals, responsible solely for facilitating the election of a proposal with substantial voter support. For any proposal $j \in C$, every dRep has an *advertised*, "intended" vote (or type), $t(j) \in \{0, 1\}$, which is visible by the voters. We assume here that dReps present votes for all proposals.[1] We will sometimes abuse the notation and refer by t both to the type vector of a dRep as well as to the dRep itself. The distance between a voter i and a dRep of type t is calculated using the Hamming distance function and is dependent on the revealed preferences of i and the advertised type of t in proposals that are revealed to i. Formally, let $t_{|i}$ be the projection of the type t to the proposals that belong to R_i; then we define the distance between a voter i and a dRep t as $d(i, t) := H(\widehat{v_i}, t_{|i})$.

Agreement Threshold. For a voter to delegate their vote, they have to agree with the dRep in a certain number of proposals. This is captured by a threshold bound, any agreement above which results in delegation. To make this formal, we will assume that voter i delegates their vote to a dRep when their distance to the dRep's type, taking into account only the voter's revealed preferences, is at most $\lfloor \frac{m_i - k_i}{2} \rfloor$, where k_i is a parameter quantifying the reluctance of voter i to entrust their voting power to a proxy. Obviously $k_i \le m_i$, for every i, and we will mainly focus on scenarios in which all voters have the same parameter, thus $k = k_i$, for every voter i. For example, when $k = 0$, a voter delegates their vote if they agree with the dRep in at least half of the proposals in their revealed set; we will refer to this case as *majority agreement* (see also [13, 16] for a use of a very similar threshold in a difference context). Given a dRep of type t, we say that t *attracts* a set of voters $A(t) := \{i \in N : d(i, t) \le \lfloor \frac{m_i - k_i}{2} \rfloor\}$. Additionally we define $A(D)$, for a set of dReps D as

$$A(D) := \{i \in N : \exists\, t \in D \text{ s.t. } d(i, t) \le \lfloor \frac{m_i - k_i}{2} \rfloor\}.$$

If a voter is attracted by multiple dReps, we assume they delegate to any of them arbitrarily; this choice makes our positive results stronger, whereas, notably, our negative results work for any choice (e.g., even for the more intuitive choice of the closest, in terms of Hamming distance, accepted dRep).

Preference Profiles. Let $V = (v_i)_{i \in N}$ and $\widehat{V} = (\widehat{v_i})_{i \in N}$. We call *intrinsic preference profile* $P = (N, C, V)$ a voting profile that contains the intrinsic preferences of the voters in N on proposals from C. Similarly, we call *revealed preference profile* $\widehat{P} = (N, C, \widehat{V})$ the voting profile that contains their revealed preferences. Finally, $\widehat{P}_D = (N, C, \widehat{V} \cup D)$ refers to the preference profile on proposals from C, that contains the revealed preferences of the voters in N as well as the advertised types of the dReps in D.

Approval Voting Winners. The winner of the election is the proposal with the highest (weighted) approval score. Formally, let $sc(j)$ denote the score of a proposal $j \in C$ in the profile P, i.e., the total weight of the voters $i \in N$ such that $v_i(j) = 1$. A proposal $j \in C$ is the winning proposal in the profile P if $sc(j) \ge sc(j'), \forall j' \in C$. Similarly, we define $\widehat{sc}(j)$ and $\widehat{sc}_D(j)$ to be the score of a proposal $j \in C$ in the profile $\widehat{P}$ and $\widehat{P}_D$, respectively. Note that $\widehat{sc}(j)$ represents the score that proposal j would attain if all voters were to vote directly and $\widehat{sc}_D(j)$ comprises the scores of the dReps (whose weight is the total weight of the voters they have attracted) and the scores of the voters that have not delegated their votes to any dRep, i.e., that are voting directly. Let $\texttt{win}(P) := \arg\max\{sc(j), j \in C\}$ and $\texttt{opt}(P) := sc(\texttt{win}(P))$. The same notions can be extended to profiles $\widehat{P}$ and $\widehat{P}_D$.

Tie-breaking for the Winner: Given a profile P, we assume that $\arg\max\{sc(j), j \in C\}$ returns a single winning proposal rather than a set, according to a tie-breaking rule. Consistently with our discussion on the "best-case scenario" (see Section 1.1), we assume that the tie-breaking is always in favor of the proposal with the maximum intrinsic score. In the full version, we briefly explore alternative tie-breaking rules, noting a significant strengthening of our negative findings.

Our goal is to select a set of λ dReps that will collectively (by participating in the election and representing voters according to the submitted thresholds) ensure that a proposal of high intrinsic approval score will be elected, as formally presented in the definition of the PROXY SELECTION problem, that follows.

[1] We could also allow dReps to abstain in some proposals, and this would not make any difference in our setting.

PROXY SELECTION$(P, \widehat{P}, k, \lambda)$

Input: An intrinsic voting profile P and a revealed voting profile $\widehat{P}$ on a set C of m proposals and a set N of n weighted voters; the voters' true (resp. revealed) preferences V (resp. $\widehat{V}$); a parameter k, so that a voter i is attracted by a dRep with type t if $d(i,t) \leq \lfloor \frac{m_i - k}{2} \rfloor$; an upper bound λ on the number of dReps.

Output: Specify type vectors for all dReps in D, with $|D| \leq \lambda$, such that $\mathtt{win}(P) = \mathtt{win}(\widehat{P}_D)$.

The performance of a suggested set of dReps is measured in terms of how well the intrinsic score of the winning proposal under their presence approximates the highest intrinsic approval score.

Definition 2.1. Let $\rho \geq 1$. We say that a set of dReps D achieves a ρ-approximation if for every intrinsic and revealed preference profiles P and $\widehat{P}$, it holds that $sc(\mathtt{win}(\widehat{P}_D)) \geq \frac{1}{\rho}\, sc(\mathtt{win}(P))$.

One might be inclined to believe that attracting a sufficiently large set of voters is enough to achieve an analogous approximation ratio guarantee. The following proposition establishes that this is not the case, demonstrating that the attraction part is only one component towards solving PROXY SELECTION. Therefore, achieving good approximations requires further insights. Additionally, its proof serves as a smooth warm-up to the studied framework.

Proposition 2.2. *It is possible for a single dRep to attract half of the voters without achieving a 2-approximation.*

PROOF. Consider an instance of majority agreement on four proposals. Voters' intrinsic preferences are presented in the table below, where the revealed preferences are given in white background.

	I_1	I_2	I_3	I_4
voter 1	1	1	0	0
voter 2	1	0	1	1
voter 3	1	0	1	1

It is evident that $\mathtt{win}(P) = I_1$, and $\mathtt{opt}(P) = 3$. At the same time the dRep that votes in favor of all proposals attracts voter 2 and voter 3. This means that $\widehat{sc}_D(I_1) = |A(D)| = 2$, while $\widehat{sc}_D(I_2) = 3$, since both the dRep and voter 1 vote in favor of I_2. However, $sc(I_2) = 1 = \frac{1}{3}\mathtt{opt}(P)$, yielding an approximation ratio of 3. □

Coherent Voters' Sets and Instances. We conclude the section with the definition of an important notion for our work, that of a *coherent set of voters*, i.e., sets of voters with the same revealed sets. Several of our positive results will be parameterized by properties of those sets, such as the size of the largest coherent set.

Definition 2.3. A set of voters $N' \subseteq N$ is called coherent if $R_i = R_{i'}, \forall i, i' \in N'$. An instance of PROXY SELECTION is called coherent if N is coherent.

We remark that given an instance of PROXY SELECTION, it is computationally easy to verify if it is coherent, or to find the largest coherent set of voters.

3 THEORETICAL FINDINGS

We start our investigation with the case of a single dRep ($\lambda = 1$). Our main result here is rather negative, namely that no matter how the dRep chooses their vote, the approximation ratio falls short of being deemed adequate, confined by a linear dependence on the number of voters.

Theorem 3.1. *For a single dRep and any $k > 0$, the approximation ratio of PROXY SELECTION is $\Omega(n)$.*

PROOF. Consider an instance with an odd number of $m > 3$ proposals and $n = m - 1$ voters, where $k_i > 0, \forall i \in [n]$, such that:

- For every voter $i \in [n]$, their preferences with respect to proposal m are as follows: $v_i(m) = 1$ and $\widehat{v}_i(m) = \perp$.
- The remaining $m-1$ proposals are partitioned in $\frac{m-1}{2}$ pairs, say $\{1,2\}, \{3,4\}, \ldots, \{m-2, m-1\}$ and for each one of these $\frac{m-1}{2}$ pairs of proposals, say $\{j, j+1\}$, there are two distinct voters, namely voters j and $j+1$, where voter j votes in favor of both proposals j and $j+1$ whereas voter $j+1$ votes in favor of proposal j but against proposal $j+1$; for every other proposal j', the preferences of these voters are $v_i(j') = 0$ and $\widehat{v}_i(j') = \perp$, where $i \in \{j, j+1\}$.

Say that P and $\widehat{P}$ are the intrinsic and revealed profiles of the created instance, respectively. Clearly, $\mathtt{win}(P) = m$ and $\mathtt{opt}(P) = n$. We claim that a single dRep, called t, regardless of their advertised type, will contribute to electing a proposal i that satisfies $sc(i) \leq 2$, leading to an inapproximability of $\frac{n}{2}$. Towards this, first notice that t cannot attract both voters of any pair. This is easy to see, as a distance of $\max\{0, \lfloor \frac{m_i - k_i}{2} \rfloor\}$ for $m_i = 2$ and $k_i \geq 1$ means agreement on both revealed proposals. As a result, for any such pair of voters, t will either attract zero or one voter(s), given that any two voters do not share the same preferences with respect to both their revealed proposals.

Consider first the scenario where $A(\{t\}) = \emptyset$. In this case, since $\widehat{sc}_D(m) = 0 < \widehat{sc}_D(j)$ for any proposal $j \neq m$, it holds that there exists a proposal j' such that $\mathtt{win}(P_D) = j'$ for which $sc(j') = 2$, or in other words, the direct voting will lead to the election of an outcome that is being accepted by exactly two voters. On the other hand, if t attracts at least one voter, say i where i is odd (resp. i is even), then t must have voted in favor of at least one proposal apart from proposal m, namely for proposal i (resp. for proposal $i-1$), or else i would not fully agree with t. But then, voter $i+1$ (resp. $i-1$), who is not attracted by t, is also voting in favor of proposal i (resp. for proposal $i-1$). This results to a proposal i such that $\widehat{sc}_D(i) = \widehat{sc}_D(m) + 1$. Therefore, the winning proposal is not m but a proposal i for which $sc(i) \leq 2$. □

For the case of a single dRep, Theorem 3.1 should be interpreted as a very strong impossibility result since it holds even in the "best-case scenario" (see the discussion in Section 1.1). A natural follow-up question is whether some meaningful domain restriction can circumvent this impossibility. For this, we will appeal to the notion of coherent sets of voters, defined in Section 2, and we will show a bounded approximation guarantee that degrades smoothly as the size of the largest coherent set grows and as the agreement threshold decreases.

Theorem 3.2. *For a single dRep,* PROXY SELECTION *admits an approximation ratio of* $\min\left\{n, \frac{3n(k+2)}{2|S|}\right\}$*, where S is the largest coherent set of voters in the instance.*

An immediate, but noteworthy, corollary of Theorem 3.2 concerns coherent instances and majority agreement.

Corollary 3.3. *For a single dRep, the approximation ratio for coherent instances and majority agreement becomes* 3.

Certainly, for majority agreement and coherent instances, the impossibility result of Theorem 3.1 does not apply. In that case, one might wonder what the best achievable approximation ratio is. To partially answer this question we offer the following result.

Theorem 3.4. *Let $\varepsilon > 0$. For a single dRep,* PROXY SELECTION *does not admit a $(1.6 - \varepsilon)$ approximation, even for coherent instances and majority agreement.*

Proof. Consider an instance in which $N = \{v_1, v_2, \dots, v_8\}$ and $C = \{c_1, c_2, c_3, c_4\}$. Suppose that $v_i(c_1) = 1$ and that $\widehat{v_i}(c_1) = \perp$, for every voter $i \in N$. Furthermore, there is exactly one voter whose preferences with respect to proposals c_2, c_3, c_4 belongs in $\{110, 101, 011, 100, 010, 001\}$ and two voters that are voting for $\{111\}$. Note that in the current proof, for simplicity, we abuse the notation and use strings instead of ordered tuples to indicate voters' preferences. In this instance, $\mathtt{opt}(P) = sc(c_1) = 8$ and $\widehat{sc}(j) = 5, \forall j \in C \setminus \{c_1\}$. Therefore, direct voting cannot result in an approximation factor that is better than 1.6. We will prove that for any possible choice t of advertised ballot of a dRep, and if $D = \{t\}$, then $sc(\mathtt{win}(\widehat{P}_D)) \leq 5$, which again results to the claimed factor.

- If $t = 111$, then $A(D)$ equals the set of voters whose preferences belong in $\{111, 110, 101, 011\}$. Therefore $|A(D)| = 5$ and hence $\widehat{sc}_D(c_1) = 5$. However c_2 is both approved by the dRep and by a voter that doesn't belong to $A(D)$, consider, e.g., the voter who is voting for 100, which leads to $\widehat{sc}_D(c_2) = 6$ and hence to a winning proposal $\mathtt{win}(\widehat{P}_D)$ such that $\widehat{sc}_D(I') \geq 6$. Hence, $\mathtt{win}(\widehat{P}_D) \neq c_1$, and, consequently, $sc(\mathtt{win}(\widehat{P}_D)) = 5$.
- If $t \in \{110, 101, 011\}$, then $A(D)$ equals the set of voters whose preferences belong in $\{111, 110, 100, 010\}$. Therefore $|A(D)| = 5$ and hence $\widehat{sc}_D(c_1) = 5$. Using the same rationale to before, one can observe that, again, $sc(\mathtt{win}(\widehat{P}_D)) = 5$.
- If $t \in \{100, 010, 001\}$, then $A(D)$ equals the set of voters whose preferences belong in $\{100, 110, 101\}$. Therefore $|A(D)| = 3$ and hence $\widehat{sc}_D(c_1) = 3$. However c_2 is both approved by the dRep and by the voter whose preference vector is 111, who does not belong to $A(D)$, which leads to $sc(\mathtt{win}(\widehat{P}_D)) = 5$.
- If $t = 000$, then $|A(D)| = 3$, but $\widehat{sc}_D(c_2) = 4$, which again leads to a winning proposal of $sc(\mathtt{win}(\widehat{P}_D)) = 5$. □

To understand the power of coherence towards achieving good approximations, it is instructive to explore the limitations of the best possible dReps also on coherent instances. To this end, we provide a couple of results: the first generalizes Theorem 3.1 to be parameterized by the size of the largest coherent set, and the second shows robustness to coherent instances, as long as the agreement thresholds are sufficiently high. The core message of those results is that coherent sets are not a panacea, and can result in meaningful approximations only under further conditions.

Theorem 3.5. *Let $\varepsilon > 0$. For a single dRep and any $k > 0$,* PROXY SELECTION *does not admit an $\left(\frac{n}{|S|} - \varepsilon\right)$-approximation, where S is the largest coherent set of voters in the instance.*

Theorem 3.6. *Let $\varepsilon > 0$. For a single dRep and $k = m - 2c(m)$, with $c(m) \in o(m)$, the approximation ratio of* PROXY SELECTION *is $\Omega(n)$, even for coherent instances.*

We conclude our discussion for the case of one dRep with a complementary result of a computational nature: a theorem that establishes that finding a dRep to attract voters in a way that ultimately elects the optimal proposal is computationally hard. Consequently, PROXY SELECTION turns out to be challenging both from the standpoint of information theory and computational complexity.

Theorem 3.7. *The decision variant of* PROXY SELECTION *is NP-hard, even for majority agreement and a single dRep.*

Proof Sketch. We will establish NP-hardness for the decision version of PROXY SELECTION, where for some given integer parameter r, we want to answer if there exists an advertised type for a dRep, so that the intrinsic score of the elected outcome is at least r. We will present here the formal construction and we refer to the full version for an illustrative exposition of voters' ballots as well as for the proof of correctness. We will reduce from the NP-hard problem MINIMAX APPROVAL VOTING (MAV) [15, 27]. In MAV, we are given an instance I of m binary proposals and n ballots where $v_i \in \{0,1\}^m, i \in [n]$ and we are asked for a vector v for which it holds that $\max_{i \in [n]} H(v_i, v) \leq \theta$, where H is the Hamming distance between two vectors of the same size. The NP-hardness has been established for instances with m being even, and $\theta = m/2$. We create an instance I' of PROXY SELECTION as follows:

- We have $m' = m + 3$ binary candidate proposals: $\{c_1, \dots, c_m, c_{m+1}, c_{m+2}, c_{m+3}\}$, i.e., three additional proposals from I.
- We have n voters corresponding to the voters of I, and an additional number of $n + 1$ dummy voters.
- For every voter $i \in [n]$, belonging to the group of the first n voters, their preferences for the first m proposals in I' are just as they are in I, and they are all revealed, so that $m_i = m$. The remaining three proposals are not visible for these voters and their intrinsic preferences are that $v_i(c_{m+1}) = 1$, $v_i(c_{m+2}) = v_i(c_{m+3}) = 0$.
- For the dummy voters, none of them approve the first m proposals, which are also not revealed to them. As for the last three proposals, there are exactly two dummy voters, who will be referred to as the special dummy voters, who approve all three proposals, and all three are revealed to them. All the remaining $n - 1$ dummy voters approve only the proposals c_{m+2} and c_{m+3}, which are revealed to them, whereas c_{m+1} is disapproved, and also not revealed to them.
- We set $r = n + 2$ and $\lambda = 1$, i.e. we have only one dRep available. Hence we are looking for an advertised type of the dRep, so that the instrinsic score of the elected outcome is at least $n + 2$.

Note that the only proposal that has an intrinsic score of $n + 2$ is the proposal c_{m+1}, while all the others have lower scores. But c_{m+1} cannot be elected via only direct voting, since it is not revealed to the first n voters. Hence, the question is whether there exists an advertised type for the dRep that can make c_{m+1} elected. □

Importantly, we note that all of the approximation guarantees in our paper can be obtained in time polynomial in the input size.

3.1 Multiple dReps

Having examined the limitations of what a single dRep can achieve, we turn our attention to the case of multiple dReps as this is not captured by the impossibility of Theorem 3.1. Is it possible to achieve much better approximations by using sufficiently many dReps?

A simple but reinforcing observation is that for majority agreement, 2 dReps suffice to elect the optimal proposal; and this holds without the requirement of coherence.

Theorem 3.8. *When $\lambda = 2$, PROXY SELECTION for majority agreement can be optimally solved.*

But what about instances in which voters are more discerning, indicated by larger values of k? Whether good approximations with multiple delegation representatives are achievable in general is still to be determined. We begin with the case of coherent instances and we provide the following theorem, which suitably generalizes Theorem 3.8, under the assumption of coherence.

Theorem 3.9. *When $\lambda = \min\{n, 2^{k+1}\}$, PROXY SELECTION for coherent instances and any $k \geq 0$ can be optimally solved.*

Proof. We first show that we need no more than n dReps. Let $N' = \{i \in N : v_i(\mathtt{win}(P)) = 1\}$. If we set D to consist of one dRep t_i for every voter $i \in N'$ such that $t_i(j) = v_i(j), \forall j \in I$, then D will attract all voters from N' and possibly other voters outside N' too. Since all dReps in D are voting in favor of $\mathtt{win}(P)$, it holds that $\widehat{sc}_D(\mathtt{win}(P)) = A(D) \geq |N'|$. Suppose that $\mathtt{win}(\widehat{P}_D) \neq \mathtt{win}(P)$. Therefore, it should also hold that $\widehat{sc}_D(\mathtt{win}(\widehat{P}_D)) > \widehat{sc}_D(\mathtt{win}(P)) \geq |N'|$. However, by the definition of N', $sc(\mathtt{win}(P)) = |N'|$ and, by the construction of D, $sc(\mathtt{win}(\widehat{P}_D)) \geq \widehat{sc}_D(\mathtt{win}(\widehat{P}_D))$. Thus,

$$sc(\mathtt{win}(\widehat{P}_D)) = \widehat{sc}_D(\mathtt{win}(\widehat{P}_D)) > |N'| = sc(\mathtt{win}(P)),$$

which contradicts the optimality of $\mathtt{win}(P)$. Therefore, with $|N'|$ dReps, which are at most n, we can retrieve the optimal solution.

We will now proceed with proving that whenever $2^{k+1} < n$, we can use a (different) set of dReps D, where $|D| \leq 2^{k+1}$, to elect the optimal proposal. Let R be the set of commonly revealed proposals to the voters of the given instance. It is without loss of generality here to assume that $\mathtt{win}(P) \notin R$, or otherwise, the direct voting would result in the election of $\mathtt{win}(P)$. To create the set D, we fix an arbitrary set $S_k \subseteq R$, of k proposals, and for every possible binary vector on S_k, i.e., for every $\sigma \in 2^{S_k}$, we add to D exactly two dReps, namely $t_{\sigma,0}$ and $t_{\sigma,1}$, advertising the following, with respect to a proposal $j \in C$:

$$t_{\sigma,0}(j) = \begin{cases} \sigma(j), & \text{if } j \in S_k, \\ 1, & \text{if } j = \mathtt{win}(P), \\ 0, & \text{otherwise.} \end{cases} \quad t_{\sigma,1}(j) = \begin{cases} \sigma(j), & \text{if } j \in S_k, \\ 1, & \text{otherwise.} \end{cases}$$

To prove the statement, it is sufficient, to show that $A(D) = N$, i.e., that D can attract all voters from N. We fix an arbitrary voter $i \in N$. Definitely, there is a vector, say σ', that defines a pair of dReps in D, say $t_{\sigma',0}$ and $t_{\sigma',1}$ (henceforth denoted by t_1 and t_2), such that i totally agrees in all proposals of S_k both with t_1 and t_2. Formally, if for a given vector x and a set of proposals Y, we denote by $x_{|Y}$ the projection of x to the proposals in Y, the following holds:

$$\max\{H(\widehat{v}_{i_{|S_k}}, t_{1_{|S_k}}), H(\widehat{v}_{i_{|S_k}}, t_{2_{|S_k}})\} = 0 \tag{1}$$

Let $R'_i := R_i \setminus S_k = R \setminus S_k$ and for $z \in \{0, 1\}$ we define

$$R'_i(z) := |\{j \in R'_i : \widehat{v}_i = z\}|.$$

Then, either $R'_i(0) \geq R'_i(1)$, or $R'_i(1) > R'_i(0)$. Therefore,

$$\min\{H(\widehat{v}_{i\,|R'_i}, t_{1\,|R'_i}), H(\widehat{v}_{i\,|R'_i}, t_{2\,|R'_i})\} \leq \left\lfloor \frac{|R'_i|}{2} \right\rfloor. \tag{2}$$

Combining Equations (1) and (2), we have that voter i agrees either with t_1 or with t_2, in at least

$$k + \left\lceil \frac{|R'_i|}{2} \right\rceil = k + \left\lceil \frac{m_i - k}{2} \right\rceil = \left\lceil \frac{m_i + k}{2} \right\rceil$$

proposals. Consequently, every voter will delegate to a dRep from D and the optimal proposal will be elected. □

Theorem 3.9 provides a bound on the sufficient number of dReps required to make sure that the optimal proposal is elected and raises a question regarding positive results (both optimal and approximate) for not-necessarily-coherent instances. In Theorem 3.12 below we provide a generalized and more refined version of this result, that relates the achievable approximation with the required number of dReps and the parameters of the instance, but does not need to assume that the instances are coherent.

3.2 Positive Results Beyond Coherent Instances

Based on the results presented earlier in this section, coherence stands out as very useful towards achieving meaningful approximation guarantees. At what follows, we define a more refined notion, namely a quantified version of it, which provides further insights into the structure of instances and how these affect the achievable approximations. In particular, we use the notion of (x, δ)-coherent sets for sets of voters that have a common set of proposals of size x in their revealed sets, as well as at most δ additional proposals.

Definition 3.10. A set of voters $N' \subseteq N$ is called (x, δ)-coherent if there exists a set $X \subseteq C$ such that for every $i \in N$ the following hold: $X \subseteq R_i$, $|X| \geq x$, and $|R_i \setminus X| \leq \delta$.

Using Definition 3.10, we generalize the result of Theorem 3.2, with an additional loss in the factor that is dependent on the type of (x, δ)-coherent sets that an instance admits.

Theorem 3.11. *For a single dRep and for any $\delta > 0$, PROXY SELECTION admits an approximation ratio of $\min\left\{n, \frac{3n(k+\delta+2)}{2|S|}\right\}$, where S is the largest $(k + \delta, \delta)$-coherent set in the instance.*

While our positive result for the case of multiple dReps (Theorem 3.9) is notable, it comes with the drawback that it holds only under the assumption of coherent instances. Our main positive result that follows is a relaxation of Theorem 3.9, that does not require any structural assumptions, and relates the approximation with the number of dReps, the threshold bound and the structure of the instance in terms of approximate coherence.

Theorem 3.12. *When $\lambda = \min\left\{n, \zeta 2^{k+1}\right\}$, PROXY SELECTION admits an approximation ratio of $\frac{\gamma}{3\zeta}$, where γ is the minimum number of $(k, m-k)$-coherent sets that can form a partition of N, and $\zeta \leq \gamma$ with $\zeta \in \mathbb{N}$.*

An interesting corollary of Theorem 3.12 is the following: When aiming for an optimal solution with 2^{k+1} dReps, it's not a necessity for instances to be coherent; rather, the key factor is the existence of a set of k proposals commonly revealed to all voters.

Concluding our discussion on not-necessarily-coherent instances, we present a theorem for the natural case of majority agreement ($k = 0$), that extends Corollary 3.3 by establishing a connection between the achievable approximation and the structure of the revealed sets, without the requirement of coherence.

Theorem 3.13. *For a single dRep, PROXY SELECTION for majority agreement admits an approximation ratio of $\min\left\{n, \frac{3\alpha}{\beta}\right\}$, where $\alpha := |\cup_{i\in N} R_i|$ and $\beta := \min\{|R_i|, i \in N\}$.*

4 EXPERIMENTS

We complement our theoretical results with experiments on the performance of voting with proxies on realistic data sets where voters exhibit incomplete preferences. We highlight the effect that the number of dReps, revealed set sizes and thresholds have on the total number of voters who delegate instead of voting directly, as well as the approximation compared to the optimal outcome.

The first hurdle to overcome is that, by definition, it is not possible to find datasets containing both revealed and intrinsic preferences, since voters only submit the first of the two. We use the MovieLens dataset [19] to circumvent this issue, as there is enough contextual information to calculate plausible intrinsic preferences given the revealed ones. This set contains the reviews (with scores ranging from 0.5 to 5.0) given by 162.541 users to 62.423 movies. Of course, not every user has reviewed every movie. Each movie is characterized by a set of *genres* and also has a relevance score for 1.084 different *tags*, provided by the users. Using this information we calculate a plausible intrinsic user-specific "rating" $\mathtt{Rat}_i$ for any non-reviewed movie i:

$$\mathtt{Rat}_i = \frac{\sum_{j\in\mathcal{R}} \mathtt{Rat}_j \cdot \mathtt{sim}(i,j)}{\sum_{j\in\mathcal{R}} \mathtt{sim}(i,j)},$$

where $\mathtt{sim}$ is a similarity metric (taking into account tag and genre relevance) and $\mathcal{R}$ is the set of movies reviewed by the user, hence $\mathtt{Rat}_j$ is the rating that user gave to movie j. These movie ratings are between 0.5 and 5.0, and are then converted to approval preferences by comparing them with the average rating given by that user.

To calculate $\mathtt{sim}(\cdot,\cdot)$ we use two vectors per movie i:

- the *tag* vector: $\mathtt{tag}_i \in [0,1]^{1084}$,
- the *genre* vector: $\mathtt{genre}_i \in \{0,1\}^{20}$,

and the actual similarity is given by

$$\mathtt{sim}(i,j) = 1.2^{\|\mathtt{tag}_i - \mathtt{tag}_j\|} \cdot \left(0.5 + \frac{\mathtt{genre}_i \cdot \mathtt{genre}_j}{\|\mathtt{genre}_i\| \cdot \|\mathtt{genre}_j\|}\right).$$

In addition, we take a random sample of users and movies such that each user has reviewed at least 5% of the movies and no movie has been reviewed by more than 10% of the users. Specifically, we begin by sampling a large subset of 13000 users and 150 movies, and filter out the users that have reviewed fewer than 5% of those movies, as well as the movies that have been reviewed by more than 10% of the users. This makes the effect of delegation (and using multiple dReps in particular) more pronounced, as there are few completely uninformed users that would readily delegate and no clear pick for a best movie. To allow for meaningful comparisons, we add a movie that has not been reviewed by any user, but is approved by all in their intrinsic ratings. Note that if we generate the preferences for all movies as described earlier, it turns out that in the produced instances, voting directly (without using any dReps) results in good approximations. We would like to explore what happens in the most interesting cases, of which the one with a proposal with full intrinsic support is the most challenging.

Since finding the optimal set of dReps is computationally intractable in the worst case (see Theorem 3.7), we use a greedy heuristic to approximate it. Specifically, we can build the type of a dRep t incrementally, proposal by proposal, setting $t(j) = 1$ or $t(j) = 0$ depending on which attracts the most voters, assuming that the voters' revealed sets only include proposals up to proposal j. We can repeat this procedure to create multiple dReps, removing the users that have already delegated at the end of every iteration.

We first measure the *fraction* of users that opted to delegate, as a function of the number of dReps and either their k_i (for fixed m_i) (Figure 1, left) or their m_i (Figure 1, right). In the first case each k_i ranges from 0 up to $0.4\,m_i$, indicating users increasing levels of user agreement required before delegating. In the second, each m_i ranges from 0 to $0.8\,m$, capturing increasingly detailed revealed user preferences (while keeping $k_i = 0.2\,m_i$). To obtain these m_is we start from the calculated intrinsic preferences (of size m) and we then "hide" some coordinates, uniformly at random, yielding the sparser revealed preferences. Our results show that it is easier to attract delegation from voters with smaller revealed sets (leading to smaller coherent sets) or with lower k_i. It turns out, that this also does indeed translate to better approximations of the optimal approval score (see Figure 2; created similarly to Figure 1). These results suggest that, interestingly, the case of coherent instances studied in many of our theoretical results seems to be the most challenging in practice.

Notice that while the graphs in Figures 1 and 2 are qualitatively similar, there are certain important differences. Specifically, the clearly diminishing returns structure observed when the only objective is to attract delegation in Figure 1 is not present in Figure 2. This is because while it gets progressively more difficult to "cluster" voters, the quality of the outcomes increases in ever bigger jumps: at the very top, the difference between the best and second-best proposals (movies) will be greater than the second-best and third-best and so on, until their quality plateaus. This is also why the "jumps" in the approximation are steeper in some regions: a certain number of dReps does not have any effect (because even though they attract delegation, they cannot change the winner), but then suddenly this changes.

Each run is repeated 20 times for a more accurate empirical mean and standard deviation. All experiments ran on a 2021 M1 Pro MacBook with 16 gigabytes of memory and 10 high-performance cores. The code (available at www.plazos.me/code/DelegatedVoting/) is parallelized and requires a few hours to complete, including some initial preprocessing of the data.

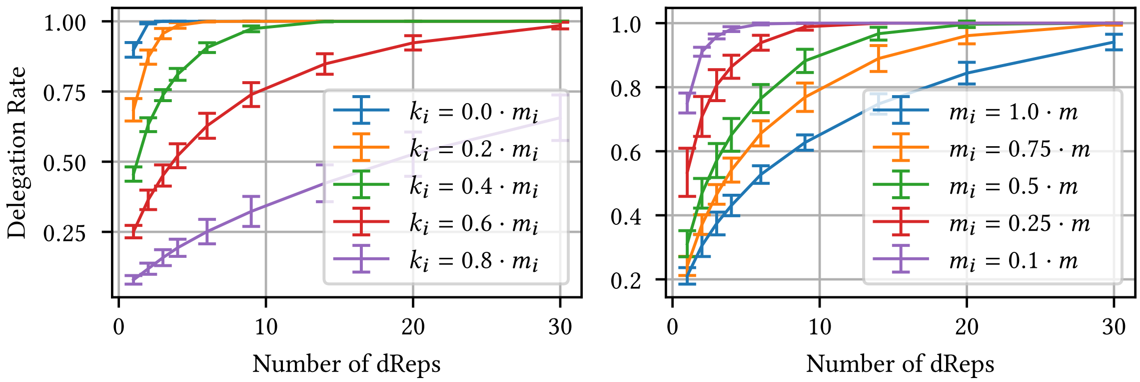

Figure 1: The fraction of voters that chose to delegate as a function of the total number of dReps.

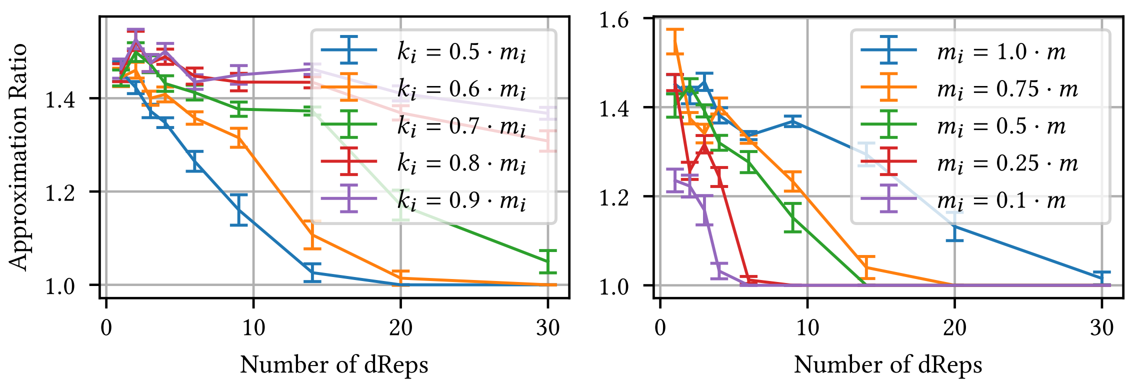

Figure 2: The quality of the approximation as a function of the total number of dReps.

5 DIRECTIONS FOR FUTURE WORK

The upper and lower bounds presented in our work are not always tight, and future work could focus on sharpening these bounds. Perhaps more interesting is the migration from the "best-case scenario" that we study in this work. This would most probably entail the following two components:

- *An information model for the dReps*. It would be reasonable to assume that each dRep is correctly-informed about each voter i's approval preference of each proposal j with some probability p_{ij}, or that they are (perfectly or imperfectly) informed about a randomly-chosen set of proposals for each voter.
- *A rationality model for the dReps*. Delegate representatives might not have any incentives to coordinate towards the socially desirable outcome, and they would need to be properly incentivized to do that, e.g., via the form of payments. They could even have their own preferences regarding the proposals under consideration.

One could think of many other examples or refinements of the above, and the appropriate choice of information/rationality model for the dReps would depend on the application at hand. Tie-breaking rules that do not necessarily favor the optimal proposal also worth studying. Regardless of these choices however, the results of the "best-case scenario" should be the starting point of any investigation into those settings. Besides those extensions, other directions that we see as promising routes for further research on the topic include different distance metrics, different voting rules, the multi-winner setting, elections on interdependent proposals, as well as a partial delegation setting where voters can opt to delegate only on proposals for which they have no opinion.

ACKNOWLEDGMENTS

Georgios Amanatidis and Evangelos Markakis were partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program. Georgios Papasotiropoulos was supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.) under the "3rd Call for H.F.R.I. PhD Fellowships" (Fellowship Number: 5163) and by the European Union (ERC, PRO-DEMOCRATIC, 101076570). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the E.U. or the E.R.C. Neither the E.U. nor the granting authority can be held responsible for them.

REFERENCES

[1] Shiri Alouf-Heffetz, Laurent Bulteau, Edith Elkind, Nimrod Talmon, and Nicholas Teh. 2022. Better Collective Decisions via Uncertainty Reduction. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 24–30.

[2] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros Voudouris. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI) Survey Track*. 4294–4301.

[3] Manel Ayadi, Nahla Ben Amor, Jérôme Lang, and Dominik Peters. 2019. Single transferable vote: Incomplete knowledge and communication issues. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1288–1296.

[4] Haris Aziz and Nisarg Shah. 2021. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*. Springer, 215–236.

[5] Dorothea Baumeister and Sophie Dennisen. 2015. Voter Dissatisfaction in Committee Elections.. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1707–1708.

[6] Ruben Becker, Gianlorenzo D'angelo, Esmaeil Delfaraz, and Hugo Gilbert. 2021. Unveiling the truth in liquid democracy with misinformed voters. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*. 132–146.

[7] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. 2010. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*. 387–392.

[8] Markus Brill, Rupert Freeman, and Vincent Conitzer. 2016. Computing possible and necessary equilibrium actions (and bipartisan set winners). In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 418–424.

[9] Yves Cabannes. 2004. Participatory budgeting: a significant contribution to participatory democracy. *Environment and urbanization* 16, 1 (2004), 27–46.

[10] Ioannis Caragiannis and Evi Micha. 2019. A contribution to the critique of liquid democracy. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 116–122.

[11] Alfonso Cevallos and Alistair Stewart. 2021. A verifiably secure and proportional committee election rule. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*. 29–42.

[12] Gal Cohensius, Shie Mannor, Reshef Meir, Eli A. Meirom, and Ariel Orda. 2017. Proxy Voting for Better Outcomes. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 858–866.

[13] Andrei Constantinescu and Roger Wattenhofer. 2023. Computing the Best Policy That Survives a Vote. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2058–2066.

[14] Edith Elkind, Piotr Faliszewski, Martin Lackner, and Svetlana Obraztsova. 2015. The complexity of recognizing incomplete single-crossing preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*. 865–871.

[15] Moti Frances and Ami Litman. 1997. On Covering Problems of Codes. *Theory Computing Systems* 30, 2 (1997), 113–119.

[16] Robin Fritsch and Roger Wattenhofer. 2022. The Price of Majority Support. In *Proceedings of the 21st International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 436–444.

[17] Daniel Halpern, Joseph Y Halpern, Ali Jadbabaie, Elchanan Mossel, Ariel D Procaccia, and Manon Revel. 2023. In Defense of Liquid Democracy. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. 852.

[18] Daniel Halpern, Gregory Kehne, Ariel D Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. 2023. Representation with incomplete votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*. 5657–5664.

[19] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems* 5, 4 (2015), 1–19.

[20] Aviram Imber, Jonas Israel, Markus Brill, and Benny Kimelfeld. 2022. Approval-Based Committee Voting under Incomplete Information. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*. 5076–5083.

[21] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. 2021. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research* 70 (2021), 1223–1252.

[22] Meir Kalech, Sarit Kraus, Gal A Kaminka, and Claudia V Goldman. 2011. Practical voting rules with partial information. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 151–182.

[23] Anna Maria Kerkmann and Jörg Rothe. 2019. Stability in FEN-hedonic games for single-player deviations. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 891–899.

[24] Aggelos Kiayias and Philip Lazos. 2022. SoK: Blockchain Governance. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies (AFT)*. 61–73.

[25] Martin Lackner. 2014. Incomplete preferences in single-peaked electorates. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*. 742–748.

[26] Jérôme Lang. 2020. Collective decision making under incomplete knowledge: possible and necessary solutions. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 4885–4891.

[27] Rob LeGrand, Evangelos Markakis, and Aranyak Mehta. 2007. Some results on approximating the minimax solution in approval voting.. In *Proceedings of the 6th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 1185–1187.

[28] Reshef Meir, Fedor Sandomirskiy, and Moshe Tennenholtz. 2021. Representative committees of peers. *Journal of Artificial Intelligence Research* 71 (2021), 401–429.

[29] Alois Paulin. 2020. An overview of ten years of liquid democracy research. In *The 21st Annual International Conference on Digital Government Research*. 116–121.

[30] Marcus Pivato and Arnold Soh. 2020. Weighted representative democracy. *Journal of Mathematical Economics* 88 (2020), 52–63.

[31] Manon Revel, Daniel Halpern, Adam Berinsky, and Ali Jadbabaie. 2022. Liquid Democracy in Practice: An Empirical Analysis of its Epistemic Performance. In *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*.

[32] Zoi Terzopoulou. 2023. Voting with Limited Energy: A Study of Plurality and Borda. In *Proceedings of the 22nd International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 2085–2093.

[33] Zoi Terzopoulou, Alexander Karpov, and Svetlana Obraztsova. 2021. Restricted domains of dichotomous preferences with possibly incomplete information. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5726–5733.

[34] Aizhong Zhou, Yongjie Yang, and Jiong Guo. 2019. Parameterized complexity of committee elections with dichotomous and trichotomous votes. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. 503–510.