

Better Collective Decisions via Uncertainty Reduction

Shiri Alouf-Heffetz^{1,4}, Laurent Bulteau², Edith Elkind³, Nimrod Talmon¹ and Nicholas Teh³

¹Ben-Gurion University, Israel

²LIGM, CNRS, Université Gustave Eiffel, France

³Department of Computer Science, University of Oxford, UK

⁴TwoFiveOne Research

shirihe@post.bgu.ac.il, laurent.bulteau@u-pem.fr, elkind@cs.ox.ac.uk, talmonn@bgu.ac.il,
nicholas.teh@cs.ox.ac.uk

Abstract

We consider an agent community wishing to decide on several binary issues by means of issue-by-issue majority voting. For each issue and each agent, one of the two options is better than the other. However, some of the agents may be confused about some of the issues, in which case they may vote for the option that is objectively worse for them. A benevolent external party wants to help the agents to select the majority-preferred option for as many issues as possible. This party may have one of the following tools at its disposal: (1) educating some of the agents, so as to enable them to vote correctly on all issues, (2) appointing a subset of highly competent agents to make decisions on behalf of the entire group, or (3) guiding the agents on how to delegate their votes to other agents, in a way that is consistent with the agents' opinions. For each of these tools, we study the complexity of the decision problem faced by this external party, obtaining both NP-hardness and fixed-parameter tractability results.

1 Introduction

A rural community needs to decide whether to install wind turbines on a nearby hill. Some 40% of the residents are certain that this is a good idea, because they are committed to renewable energy. About 30% of the residents are firmly against the proposal, because of the construction noise or the impact on the views from their homes. However, the remaining 30% of the residents struggle to make up their minds: they are generally in favor of wind power, but worry that the location chosen for the project may be on the migration route of a rare bird species. A local environmental charity wants to help this community to make a decision that is consistent with the majority's preferences. By consulting with scientists, the charity concludes that the selected location is unlikely to present a danger for the birds, so building the turbines is the 'correct' decision. As the decision will be made by a majority vote at a community meeting, the charity needs to reach out to (some of) the undecided voters and explain that it is in their best interest to vote in favor of the wind turbines. Another approach would be to discourage the confused residents from

participating in the meeting, or to suggest to them that they may want to delegate their votes to more competent voters; these strategies may result in decisions that are better for the entire community, including the confused voters themselves.

Inspired by this scenario, in this work we study the problem faced by a benevolent party that wants to help a group of agents to make majority-preferred decisions on a number of binary issues. For each agent and each issue, one of the two options is 'correct', in the sense that this is the option the agent would have selected if they could invest time and effort into studying it. However, just as in our example, some of the agents may be uncertain about some of the issues, in which case they may vote against their best interest. Examples of such confusion abound in the real world and are well-documented, e.g., in the context of Brexit or nuclear power.

In our setting, there are multiple issues and the benevolent party does not have its own preferences over the outcomes. Rather, it wants to maximize the number of issues on which the group is guaranteed to make the decision that matches the true preferences of the majority. This party, which may be a charity, an impartial governmental organization, or an ad-hoc working group, may have several tools at its disposal. For instance, it may be able to reach out to a subset of the agents and offer them an opportunity to learn more about the issues (e.g., by presenting information in an accessible manner, or holding a Q&A session with an expert). Alternatively, it may discourage some of the agents from participating in the vote.

The benevolent party may also help the agents with delegation decisions. Specifically, we assume that the voting mechanism used by the community supports delegation: an agent may delegate her vote to another agent, who will vote on her behalf on all issues. The agents are assumed to be willing to delegate to other agents who have similar preferences (i.e., i would not delegate to j if they disagree on an issue they are both certain about), but are more knowledgeable. The benevolent party can suggest delegation options to some of the agents, in a way that respects these constraints and maximizes the number of good decisions.

1.1 Our Contribution

We develop a formal model that enables us to reason about voters' confusion and ways to mitigate it. In this paper, we focus on independent binary issues, but our ideas can be extended to more complex decision-making scenarios. Our ap-

proach is worst-case: we want to maximize the number of issues that are ‘safe’, in the sense that the majority-supported outcome is guaranteed no matter how the uncertain voters cast their ballots. We consider both a general discrete issue space and a one-dimensional setting, where the issues are ordered in such a way that, for each agent, the set of issues on which she would benefit from a positive decision forms an interval of this ordering, and she is only uncertain about issues that are close to the endpoints of her ‘positive interval’. For instance, consider a vote over tax rates, and a voter whose policy preferences are consistent with a tax rate in the [25%, 35%] range: this voter may confidently reject a proposal to set the tax rate to 15% or 45% and support a proposal to set it to 30%, but may struggle to make up her mind about the proposals in the [23%, 27%] range or the [34%, 38%] range.

We formulate three computational problems that model, respectively, educating the agents, preventing some agents from participating, and helping with delegation decisions. We establish that these problems are NP-hard even in the one-dimensional setting, but show that they are fixed-parameter tractable—even in the general setting—both with respect to the number of voters and the number of issues. Moreover, we consider a natural special case of the one-dimensional setting in which all three problems are polynomial-time solvable. We omit some proofs due to space constraints.

1.2 Related Work

The analysis of voting outcomes under uncertainty is a prominent topic in computational social choice (see, e.g., [Walsh, 2007; Hazon *et al.*, 2008; Bachrach *et al.*, 2010; Wojtas and Faliszewski, 2012; Kenig and Kimelfeld, 2019]); however, it is usually assumed that it is an external party, rather than the voters themselves, who is uncertain about the votes. There is a large body of work in political science that considers voter education (see, e.g., [Lupia and McCubbins, 1998; Bullock, 2011; Boudreau *et al.*, 2019]) and the references therein); however, it does not engage with the associated algorithmic issues. The computational problem associated with removing voters is closely related to election control by deleting voters (see, e.g., the survey of Faliszewski and Rothe [2016]); however, in the control literature, it is assumed that the party engaging in voter deletion pursues its own goals rather than tries to implement the popular opinion. Finally, in the context of vote delegation, we mention works on liquid democracy from the perspective of discovering a good outcome [Brill and Talmon, 2018; Bloembergen *et al.*, 2019; Kahng *et al.*, 2021; Gözl *et al.*, 2021] in particular, Cohenius *et al.* [2017] and Green-Armytage [2015] consider vote delegation in the one-dimensional Euclidean domain.

2 Preliminaries

We will now define our model and formulate the computational problems that we are going to study.

Voting instances There is a *proposal space* P and a set of n agents N . For each agent $i \in N$, the space P is split as $P = P^+(i) \cup P^-(i)$ with $P^+(i) \cap P^-(i) = \emptyset$; the proposals in $P^+(i)$ are *beneficial* for i , so i would want to approve them, while proposals in $P^-(i)$ are not beneficial

for i . Also, for each agent $i \in N$, the space P is split as $P = P^1(i) \cup P^2(i)$ with $P^1(i) \cap P^2(i) = \emptyset$; agent i is *certain* about the proposals in $P^1(i)$ and *uncertain* about the proposals in $P^2(i)$. Thus, an agent i can be described by a tuple $\pi(i) = (P^+(i), P^1(i)) \in P \times P$, which we will call agent i ’s *belief*; we write $\Pi = (\pi(1), \dots, \pi(n))$ and refer to the triple $I = (P, N, \Pi)$ as a *voting instance*. For $i \in N$ and $\star \in \{+, -, \dagger, ?\}$, let $P^{\star\dagger}(i) = P^\star(i) \cap P^\dagger(i)$. When casting an approval ballot, agent i will vote for all proposals in $P^{+1}(i)$ and against all proposals in $P^{-1}(i)$. Her vote over proposals in $P^2(i)$, however, may be arbitrary; in particular, she may disapprove proposals in $P^{+2}(i)$ and approve proposals in $P^{-2}(i)$.

Correct and possible outcomes For each $p \in P$ and $\star \in \{+, -, \dagger, ?\}$, let $N^\star(p) = \{i \in N : p \in P^\star(i)\}$. We say that a proposal p is *good* if $|N^+(p)| \geq |N|/2$ and *bad* if $|N^-(p)| \geq |N|/2$. Note that a proposal can be both good and bad: this happens if it is beneficial for exactly $|N|/2$ agents. An outcome $z \in \{0, 1\}$ is *correct* for p if p is good and $z = 1$ (‘approve’), or if p is bad and $z = 0$ (‘disapprove’). We denote the set of correct outcomes for p by $\text{corr}(p, I)$. Because of agent uncertainty, the outcome of an approval vote is not always correct: we say that 1 (respectively, 0) is a *possible outcome* for p if $|N^+(p) \cup N^2(p)| \geq |N|/2$ (respectively, if $|N^-(p) \cup N^2(p)| \geq |N|/2$). Note that $\text{corr}(p, I) \subseteq \text{poss}(p, I)$ for each $p \in P$, but for some $p \in P$ we may have $\text{poss}(p, I) \setminus \text{corr}(p, I) \neq \emptyset$.

Safe proposals and safe zones Given a voting instance $I = (P, N, \Pi)$, we say that a proposal $p \in P$ is *safe* if $\text{poss}(p, I) = \text{corr}(p, I)$ and *unsafe* otherwise. The set of safe proposals is called the *safe zone* for I .

Example 1. Consider an instance I with $N = \{1, 2, 3\}$, $P = \{p_1, p_2, p_3, p_4, p_5\}$, and Π given by $P^+(1) = \{p_2, p_3, p_4\}$, $P^2(1) = \{p_2, p_4\}$, $P^+(2) = \{p_2, p_3\}$, $P^2(2) = \{p_1\}$, $P^+(3) = \{p_1, p_2, p_3\}$, and $P^2(3) = \emptyset$.

The first agent benefits from p_3 and knows this, so she is going to approve p_3 . She also benefits from p_2 , but is uncertain about this proposal, so she may disapprove it. We have $\text{corr}(p_1, I) = \{0\}$, since agents 1 and 2 do not benefit from p_1 (so $N^-(p_1) = \{1, 2\}$, $N^+(p_1) = \{3\}$). However, $\text{poss}(p_1, I) = \{0, 1\}$, because agent 2 is uncertain about p_2 (i.e., $p_2 \in P^2(2)$) and may vote either way on it (i.e., $|N^+(p_1) \cup N^2(p_1)| = 2$, $|N^-(p_1) \cup N^2(p_1)| = 2$). Hence, p_1 is not safe. The other four proposals, however, are safe.

Proposal Spaces

In general, our framework allows for infinite proposal spaces. However, in this work we will focus on settings where the proposal space is *finite*, i.e., $P = \{p_1, \dots, p_m\}$. A natural restricted case of our model is a one-dimensional setting, where P is ordered as $p_1 \prec \dots \prec p_m$, and agent preferences respect this order, as described below. We say that a set of proposals $P' \subseteq P$ forms an *interval* of \prec if $P' = \emptyset$ or $P' = \{p_\ell : j \leq \ell \leq k\}$ for some $1 \leq j \leq k \leq m$. In the one-dimensional setting we assume that, for each agent $i \in N$, the set $P^+(i)$ forms an interval of \prec . This approach takes inspiration from the notion of interval approval domains [Elkind and Lackner, 2015].

Furthermore, in the one-dimensional case we expect an agent to be certain that she dislikes extreme proposals and likes proposals in the center of her interval; however, she may be uncertain about proposals that are close to the endpoints of her interval. Hence, we assume that, for each $i \in N$ with $P^+(i) = \{p_j, \dots, p_k\}$, we have $P^?(i) = P_L^?(i) \cup P_R^?(i)$, where both $P_L^?(i)$ and $P_R^?(i)$ are intervals of \prec that satisfy the following conditions: (1) $P_L^?(i) = \emptyset$ or $P_L^?(i) \cap \{p_{j-1}, p_j\} \neq \emptyset$; (2) $P_R^?(i) = \emptyset$ or $P_R^?(i) \cap \{p_k, p_{k+1}\} \neq \emptyset$.

Example 2. Note that the instance I from Example 1 is one-dimensional. However, if we were to add either of the two agents 4, 5 with $P^+(4) = \{p_1, p_5\}$, $P^?(4) = \emptyset$, $P^+(5) = \{p_2, p_3, p_4\}$, $P^?(5) = \{p_3\}$, this would no longer be the case: $P^+(4)$ is not an interval, and $P^?(5)$ fails conditions (1) and (2).

As the one-dimensional setting is essentially a domain restriction, algorithms for general finite proposal spaces apply in the one-dimensional setting, while NP-hardness results for the one-dimensional setting imply hardness for general spaces. Therefore, in what follows, for all problems we consider, we prove hardness results for the one-dimensional case (Section 3) and develop FPT algorithms for the general case (Section 4). Additionally, in Section 5 we present polynomial-time algorithms for a special, more restricted case of the one-dimensional model.

Algorithmic Challenges

We consider three approaches to eliminating uncertainty: (1) *educating* agents, (2) *removing* agents, and (3) *delegating* votes (known as *liquid democracy*). In each case, the goal is to maximize the size of the safe zone; an important special case of this task is to ensure that all projects are safe.

Educating agents In the first approach we consider, at a unit cost, we can educate one agent so as to entirely remove her uncertainty: if i has been educated, then the set $P^?(i)$ becomes empty, so, when voting, agent i will approve proposals in $P^+(i)$ and disapprove proposals in $P^-(i)$.

EDUCATING AGENTS PROBLEM (EAP):

Input: A voting instance $I = (P, N, \Pi)$, a budget κ , and a parameter $\lambda \in \mathbb{N}$.

Question: Can we educate at most κ agents in N so that in the resulting instance I' the number of safe proposals is at least λ ?

Note that the input to EAP contains the list Π ; that is, we assume that, when deciding which agents to educate, we know which proposals are beneficial for them, i.e., we know the sets $P^+(i)$, $P^-(i)$ for all $i \in N$. Of course, this is not always the case; in particular, the agents themselves can be assumed to know the sets $P^{+1}(i)$, $P^{-1}(i)$ and $P^?(i)$, but not $P^+(i)$ and $P^-(i)$. However, we expect that in practice we can usually make good predictions based on, e.g., demographic criteria.

Denying access The next approach we consider is to prevent up to κ agents from participating in the vote (equivalently, this approach can be viewed as selecting a committee). However, we still want to make decisions in a way that is

beneficial to the majority of all agents, including those who are not invited to vote. That is, we consider the set of correct outcomes for each proposal with respect to the original voting instance I and the set of possible outcomes with respect to the modified instance I' , with up to κ agents removed. Note that it may be the case that $\text{corr}(p, I) = \{0, 1\}$, but $\text{poss}(p, I')$ is a singleton; we consider such outcomes as acceptable.

DENYING ACCESS PROBLEM (DAP):

Input: A voting instance $I = (P, N, \Pi)$, a budget κ , and a parameter $\lambda \in \mathbb{N}$.

Question: Can we remove at most κ agents from N to obtain an instance I' such that the number of proposals $p \in P$ with $\text{corr}(p, I) \supset \text{poss}(p, I')$ is at least λ ?

Note that, to solve DAP, we do not need full access to the ‘hidden’ sets $P^+(i)$, $P^-(i)$; it suffices to have access to the ‘known’ sets $P^{+1}(i)$, $P^{-1}(i)$ and $P^?(i)$ for each $i \in N$ as well as the correct outcome for each proposal (which we expect to be easier to estimate than individual preferences).

Allowing delegations The last approach we consider is to allow the agents to delegate their votes to other agents, transitively. An agent that accumulates t votes participates in the election with voting weight t . This framework is known as *liquid democracy* [Green-Armytage, 2015]. We assume that our agents are rather conservative in their delegation decisions: we say that agent i is *willing to delegate her vote to agent j* if (1) $P^{+1}(i) \cap P^{-1}(j) = \emptyset$ and $P^{-1}(i) \cap P^{+1}(j) = \emptyset$, and (2) $|P^?(i)| < |P^?(j)|$. Condition (1) indicates that agent i cannot observe any disagreement between herself and agent j , and condition (2) indicates that agent j is more knowledgeable than agent i . A *delegation graph consistent with an instance I* is a directed acyclic graph \mathcal{D} over the vertex set N such that the outdegree of each vertex is at most 1; it may contain an arc (i, j) only if agent i is willing to delegate to agent j . Given a graph \mathcal{D} , for each $i \in N$ we denote the unique sink reachable from i by $s(i)$; the agent $s(i)$ is the ‘guru’ of agent i , so that i will adopt the beliefs of $s(i)$. The graph \mathcal{D} corresponds to a modified instance $I'(\mathcal{D}) = (N, P, \Pi')$ such that $\pi'(i) = \pi(s(i))$ for each agent $i \in N$; intuitively, in $I'(\mathcal{D})$ agent i copies the ballot of her guru $s(i)$. We want to construct a delegation graph that results in decisions that benefit the majority of the agents for as many proposals as possible.

CONSTRAINED LIQUID DEMOCRACY (CLD):

Input: A voting instance $I = (P, N, \Pi)$, and a parameter $\lambda \in \mathbb{N}$.

Question: Is there a delegation graph \mathcal{D} consistent with I such that the number of proposals $p \in P$ with $\text{corr}(p, I) \supset \text{poss}(p, I'(\mathcal{D}))$ is at least λ ?

Importantly, in a way, our model implicitly assumes that the delegation decisions are made in a centralized manner; however, no agent can be asked to delegate their vote in a way that is unacceptable to them.

Example 3. Consider again the instance I of Example 1.

Note that, in I , we can make all proposals safe by educating agent 2. If we can only remove agents, then the only way to ensure the correct outcome on p_1 is to remove agents 2 and 3; however, if we do so, we may get an incorrect outcome on p_2 and p_4 , so no removal strategy can ensure the correct outcome on all proposals.

Finally, observe that agent 1 can delegate to agent 2 (they agree on the two proposals on which they are both certain, i.e., p_3 and p_5 , and $|P^1(2)| > |P^1(1)|$), and agent 2 can delegate to agent 3, but agent 1 cannot directly delegate to agent 3. However, if 1 delegates to 2 and 2 delegates to 3, then agent 3 becomes the guru for agent 1. No delegation strategy ensures a correct outcome on all proposals.

3 Intractability Results

It follows from our definitions that checking whether a given proposal is safe or determining the number of safe proposals can be done in polynomial time. Next we show that EAP, DAP, and CLD are NP-complete. Containment in NP follows immediately; thus, in what follows, we focus on showing that these problems are NP-hard. Perhaps surprisingly, our hardness results hold even for the one-dimensional setting.

Theorem 1. EAP, DAP and CLD are NP-complete, even in the one-dimensional proposal space.

Proof. We provide a proof for CLD; the proofs for EAP and DAP are omitted due to space constraints. We reduce from the NP-hard problem VERTEX COVER ON CUBIC GRAPHS (VC-CG) [Garey *et al.*, 1976]. An instance of this problem consists of an undirected cubic graph $G = (V, E)$ (i.e., the degree of each vertex $v \in V$ is exactly 3) and an integer κ ; it is a yes-instance if G admits a vertex cover of size κ (i.e., a subset $V' \subseteq V$ that contains at least one endpoint of each edge), and a no-instance otherwise.

Consider an instance of VC-CG given by a graph $G = (V, E)$ and an integer κ . Fix an arbitrary order \triangleleft_E on E . For each edge $e \in E$ with endpoints u and v , we introduce proposals $s_e, r_{e,u}, p_e, r_{e,v}$ (jointly denoted as the *edge gadget* for e). For each vertex u , we introduce proposals s_u, p_u , and q_u (jointly denoted as the *vertex gadget* for u). Additionally, we introduce two proposals, denoted by $s_\$, p_\$$ (the *sink gadget*). The proposals are ordered so that (1) all edge gadgets precede all vertex gadgets, and all vertex gadgets precede the sink gadget; (2) the edge gadgets appear in the order induced by \triangleleft_E ; (3) the vertex gadgets appear in an arbitrary order; (4) within an edge gadget for $e = \{u, v\}$, the proposals are ordered as $s_e \prec r_{e,u} \prec p_e \prec r_{e,v}$; (5) within a vertex gadget for u , the proposals are ordered as $s_u \prec p_u \prec q_u$; (6) within the sink gadget, the proposals are ordered as $s_\$ \prec p_\$$.

For each vertex $u \in V$, we consider the edges a, b, c incident on u (assume $a \triangleleft_E b \triangleleft_E c$), and introduce eight agents, denoted $v_{a,1}^u, v_{a,2}^u, v_{b,1}^u, v_{b,2}^u, v_{c,1}^u, v_{c,2}^u, v_1^u, v_2^u$ (they are said to be *related to* u) with the following beliefs:

- the sets $P^+(v_{a,1}^u)$ and $P^+(v_{a,2}^u)$ consist of all proposals in the shortest interval containing $p_a, r_{a,u}$ and $r_{b,u}$, and $P^+(v_{a,1}^u) = \{p_a, r_{a,u}\}$, whereas $P^+(v_{a,2}^u) = \{r_{b,u}\}$.

- the beliefs of $v_{b,1}^u$ and $v_{b,2}^u$ (resp., $v_{c,1}^u$ and $v_{c,2}^u$) are defined similarly, replacing p_a by p_b (resp., p_c), $r_{a,u}$ by $r_{b,u}$ (resp., $r_{c,u}$) and $r_{b,u}$ by $r_{c,u}$ (resp., p_u).
- the sets $P^+(v_1^u)$ and $P^+(v_2^u)$ consist of all proposals in the interval from p_u to $p_\$,$ and $P^+(v_1^u) = \{p_u, q_u\}$, $P^+(v_2^u) = \{p_\$\}$.

We then add polynomially many *balancing agents* so that we have $2K + 1$ agents in total for some $K \in \mathbb{N}$ and:

- $|N^+(p_e)| = K + 2$, $|N^{+?}(p_e)| = 2$ for each $e \in E$;
- $|N^+(r_{e,u})| = K + 2$, $|N^{+?}(r_{e,u})| = 1$ or $|N^+(r_{e,u})| = K + 3$, $|N^{+?}(r_{e,u})| = 2$ for each $e \in E, u \in e$;
- $|N^+(p_u)| = |N^+(q_u)| = K + 3$, $|N^{+?}(p_u)| = 2$, $|N^{+?}(q_u)| = 1$;
- $|N^+(s_e)| = |N^+(s_u)| = |N^+(s_\$)| = K + 1$, $N^?(s_e) = N^?(s_u) = N^?(s_\$) = \emptyset$, for each $e \in E, u \in V$; and
- $|N^+(p_\$)| = K + |V| + \kappa + 1$, $|N^{+?}(p_\$)| = |V|$.

To this end, we add agents that benefit from a single proposal only (and are certain about it), until the desired differences in support for each proposal are accomplished, and then add sufficiently many agents who benefit from all proposals, to get the target numbers. Let N_V denote the set of all vertex-related agents and let N_B be the set of balancing agents.

We now consider the delegation graph. Note first that for each $u \in V$ agent v_1^u may delegate to v_2^u , and for each edge e incident on u agent $v_{e,1}^u$ may delegate to $v_{e,2}^u$. We claim that the delegation graph cannot contain any other arcs.

To see this, observe first that for each $i \in N_B$ we have $P^?(i) = \emptyset$, $|P^+(i)| = 1$. Moreover, for each agent $j \in N_V$ we have $|P^+(j)| \geq 2$. Hence, no agent in N_B can delegate her vote to another agent, and no agent in N_V can delegate her vote to an agent in N_B .

Further, let $S = \{s_\xi : \xi \in V \cup E \cup \{\$\}\}$ be the set of *separators*, and note that $N^?(s) = \emptyset$ for all $s \in S$. Hence, for the arc (i, j) to be in the delegation graph, i and j must agree on all separators. This can only happen if the positive intervals for i and j start in the same gadget as well as end in the same gadget, which establishes our claim.

It follows that all paths in the delegation graph consist of at most one arc. Hence, any delegation scenario may only change the number of uncertain agents for each proposal, e.g., for each $p \in P$, no agent in $N^{+1}(p)$ can delegate (either directly or indirectly) to an agent in $N^{-1}(p)$ or vice versa.

Hence, we can make all proposals safe if and only if the following conditions are satisfied:

- For each $e \in E$, $|N^?(p_e)|$ decreases by at least 1.
- For each $e \in E$ and $u \in e$, $|N^?(r_{e,u})|$ does not increase.
- For each $u \in V$, $|N^?(p_u)|$ does not increase.
- $|N^?(p_\$)|$ increases by at most κ .

We will now prove that there exists a delegation graph \mathcal{D} that guarantees correct decisions on all proposals if and only if G has a vertex cover of size κ .

For the ‘if’ direction, let X be a size- κ vertex cover of G . Then, for each $x \in X$, consider the following delegations: $v_{e,1}^x \rightarrow v_{e,2}^x$ for each e incident on x , and $v_1^x \rightarrow v_2^x$. As each

edge e is incident on some $x \in X$, these delegations reduce by 1 the size of each set $N^?(p_e)$, $e \in E$, and increase by κ the size of the set $N^?(p_\S)$. The only other proposals affected by these delegations are proposals of the form $r_{e,x}$ and p_x, q_x , where $x \in X$. But then it can be checked that the sizes of the sets $N^?(r_{e,x})$, $N^?(p_x)$ and $N^?(q_x)$ either decrease by one or do not change. Specifically, if $a \triangleleft_E b \triangleleft_E c$ are the three edges incident on x then (1) $v_{a,1}^x \rightarrow v_{a,2}^x$ decreases $|N^?(r_{a,x})|$ by 1 and increases $|N^?(r_{b,x})|$ by 1; (2) $v_{b,1}^x \rightarrow v_{b,2}^x$ decreases $|N^?(r_{b,x})|$ by 1 and increases $|N^?(r_{c,x})|$ by 1; (3) $v_{c,1}^x \rightarrow v_{c,2}^x$ decreases $|N^?(r_{c,x})|$ by 1 and increases $|N^?(p_x)|$ by 1; (4) $v_1^x \rightarrow v_2^x$ decreases $|N^?(p_x)|$ and $|N^?(q_x)|$ by 1. Overall, the number of uncertain agents for each proposal falls within the desired bounds, and all proposals are safe.

For the ‘only if’ direction, consider a delegation graph \mathcal{D} that guarantees correct decisions on all proposals. Fix a node $u \in V$, and let $a \triangleleft_E b \triangleleft_E c$ be the edges of G incident on u . If \mathcal{D} contains the arc $v_{c,1}^u \rightarrow v_{c,2}^u$, then it must also contain the arc $v_1^u \rightarrow v_2^u$ (or else $|N^?(p_u)|$ would increase by 1). Similarly, if \mathcal{D} contains the arc $v_{b,1}^u \rightarrow v_{b,2}^u$, it must also contain $v_{c,1}^u \rightarrow v_{c,2}^u$ (or else $|N^?(r_{c,u})|$ would increase by 1) and hence $v_1^u \rightarrow v_2^u$. Finally, if \mathcal{D} contains the arc $v_{a,1}^u \rightarrow v_{a,2}^u$, it must also contain $v_{b,1}^u \rightarrow v_{b,2}^u$ (or else $|N^?(r_{b,u})|$ would increase by 1) and hence $v_{c,1}^u \rightarrow v_{c,2}^u$ and, in turn, $v_1^u \rightarrow v_2^u$. That is, if \mathcal{D} contains a delegation between two agents related to u , it must contain the arc $v_1^u \rightarrow v_2^u$.

Let X be the set of vertices x with $v_1^x \rightarrow v_2^x$. Each such delegation increases $|N^?(p_\S)|$ by 1, so $|X| \leq \kappa$. Now, consider an edge $e = \{u, w\}$. Since $|N^?(p_e)|$ must decrease by at least one, at least one of the delegations $v_{e,1}^u \rightarrow v_{e,2}^u$ or $v_{e,1}^w \rightarrow v_{e,2}^w$ must take place; assume without loss of generality that \mathcal{D} contains the arc $v_{e,1}^u \rightarrow v_{e,2}^u$. As argued above, this means that \mathcal{D} also contains the arc $v_1^u \rightarrow v_2^u$, i.e., $u \in X$. Hence, X is a vertex cover for G of size at most κ . \square

4 Fixed-Parameter Tractability in n and m

In this section, we show that all three of our problems—EAP, DAP, and CLD—are fixed-parameter tractable (FPT) with respect to the number of voters (n), and independently, with respect to the number of proposals (m). These results hold even for the maximization version of these problems, where the goal is to maximize the number of proposals on which the correct outcome is obtained, and also do not require the assumption that the proposal space is one-dimensional.

For the number of agents n , we can use brute-force algorithms, as there are 2^n possible subsets of agents to educate or delete, and at most n^n different delegation graphs. For the number of proposals m , our approach is based on integer linear programming.

Proposition 1. *EAP, DAP, and CLD are FPT for n as well as FPT for m .*

Proof. We show the proof for DAP and m . First, note that we can remove all proposals p with $|N^+(p)| = |N^-(p)|$ from the description of our voting instance (as we are satisfied with either outcome on any such proposal). Thus, from now on we

will assume that, for each $p \in P$, we have $|N^+(p)| > n/2$ or $|N^-(p)| > n/2$.

Given an agent i , construct a string $\mathbf{t} = (t_1, \dots, t_m)$ over $\{+, -, ?\}$, where, for each $j \in [m]$, we set $t_j = +$ if $p_j \in P^+(i)$, $t_j = -$ if $p_j \in P^-(i)$, and $t_j = ?$ if $p_j \in P^?(i)$. We will refer to \mathbf{t} as the *type* of i . Let $T = \{+, -, ?\}^m$. By construction, there are at most 3^m distinct agent types.

For each $\mathbf{t} \in T$, let $y_{\mathbf{t}}$ denote the number of agents of type \mathbf{t} in I , and let $x_{\mathbf{t}}$ denote the number of agents of type \mathbf{t} in the instance I' obtained after some agents have been deleted. The constraint that we can remove at most κ agents is encoded as

$$0 \leq x_{\mathbf{t}} \leq y_{\mathbf{t}} \text{ for all } \mathbf{t} \in T, \quad \sum_{\mathbf{t} \in T} x_{\mathbf{t}} \geq n - \kappa. \quad (1)$$

For each proposal $p_j \in P$, let $T^{+1}(j) = \{\mathbf{t} \in T : t_j = +\}$, $T^{-1}(j) = \{\mathbf{t} \in T : t_j = -\}$, $T^?(j) = \{\mathbf{t} \in T : t_j = ?\}$. Given a proposal p_j , let

$$\begin{aligned} a_j &= \sum_{\mathbf{t} \in T^{+1}(j)} x_{\mathbf{t}} - \sum_{\mathbf{t} \in T^{-1}(j)} x_{\mathbf{t}} - \sum_{\mathbf{t} \in T^?(j)} x_{\mathbf{t}}; \\ b_j &= \sum_{\mathbf{t} \in T^{-1}(j)} x_{\mathbf{t}} - \sum_{\mathbf{t} \in T^{+1}(j)} x_{\mathbf{t}} - \sum_{\mathbf{t} \in T^?(j)} x_{\mathbf{t}}. \end{aligned} \quad (2)$$

Note that $a_j, b_j \leq n$. If p_j is good, then we would like a_j to be positive, and if p_j is bad, then we would like b_j to be positive. Thus, for each $j \in [m]$ we introduce a variable z_j that takes values in $\{0, 1\}$, and add the constraint

$$n \cdot z_j \geq a_j, \quad n \cdot (1 - z_j) \geq 1 - a_j \quad (j\text{-good})$$

if p_j is good and the constraint

$$n \cdot z_j \geq b_j, \quad n \cdot (1 - z_j) \geq 1 - b_j \quad (j\text{-bad})$$

if p_j is bad. The reader can verify that condition (j-good) ensures that $z_j = 1$ if and only if $a_j > 0$, whereas condition (j-bad) ensures that $z_j = 1$ if and only if $b_j > 0$.

To summarize, the set of variables of our ILP is $\{x_{\mathbf{t}} : \mathbf{t} \in T\} \cup \{a_j, b_j, z_j : j \in [m]\}$, and its objective function is $\max \sum_{j \in [m]} z_j$. The set of constraints consists of (1), (2), and, for each $j \in [m]$, one of the constraints (j-good) or (j-bad), depending on whether p_j is good or bad, together with the constraint $0 \leq z_j \leq 1$. Following Lenstra’s result [Jr., 1983], this ILP can be solved in time $\text{poly}\left((3^m)^{O(3^m)}, n\right)$. \square

5 The Radical One-Dimensional Domain

Here we consider a more restricted variant of the one-dimensional model in which, for each agent $i \in N$, the set $P^+(i)$ is a suffix of \prec , i.e., either $P^+(i) = \emptyset$ or $p_m \in P^+(i)$. Intuitively, this model reflects settings where the proposals can be naturally ordered from radical proposals to mild proposals, so that a typical agent disapproves the most radical proposal, but approves the mildest proposal, and switches from disapproval to approval at a certain point in which the proposals become less extreme. Consequently, in this model, referred to as the *radical 1D domain*, $P^?(i)$ is an interval of \prec , for each $i \in N$. For convenience, in what follows we assume that the number of agents n is odd and hence $\text{corr}(p, I)$

is a singleton for each $p \in P$ (but our results hold also for even number of agents); slightly abusing notation, we write $\text{corr}(p, I) = z$ instead of $\text{corr}(p, I) = \{z\}$.

First, we make the following two observations.

Observation 1. *For the radical 1D domain, the list $(\text{corr}(p_i, I))_{i \in [m]}$ is of the form $(0, \dots, 0, 1, \dots, 1)$.*

Proof. Suppose that $\text{corr}(p_j, I) = 1$ for some $j < m$. Then $|N^+(p_j)| > n/2$. Consider a proposal p_k with $k > j$. For each agent $i \in N$ with $p_j \in P^+(i)$ we have $p_k \in P^+(i)$ and hence $N^+(p_j) \subseteq N^+(p_k)$. It follows that $|N^+(p_k)| > n/2$, i.e., $\text{corr}(p_k, I) = 1$. \square

Observation 2. *For the radical 1D domain, if a proposal p_j is safe and $\text{corr}(p_j, I) = 0$, then each proposal p_k with $k \leq j$ is safe. Similarly, if p_j is safe and $\text{corr}(p_j, I) = 1$, then each proposal p_k with $k \geq j$ is safe.*

Proof. Fix a safe proposal p_j . Suppose $\text{corr}(p_j, I) = 0$; the case $\text{corr}(p_j, I) = 1$ is symmetric. As p_j is safe, we have $|N^{-1}(p_j)| > n/2$. Consider a proposal p_k with $k < j$. Note that $p_j \in P^{-1}(i)$ implies $p_k \in P^{-1}(i)$. Hence $N^{-1}(p_j) \subseteq N^{-1}(p_k)$ and $|N^{-1}(p_k)| \geq |N^{-1}(p_j)| > n/2$, so p_k is safe. \square

We will now show that for the radical 1D domain the problems EAP, DAP and CLD admit polynomial-time algorithms.

Proposition 2. *In the radical 1D domain, EAP, DAP, and CLD are all solvable in polynomial time.*

Proof. By Observation 2, unsafe proposals form an interval of the form $\{p_{j+1}, \dots, p_{k-1}\}$ for some $p_j, p_k \in P$ with $\text{corr}(p_j, I) = 0$, $\text{corr}(p_k, I) = 1$. We consider all $O(m^2)$ possible choices for j and k such that $k+1-j \leq \lambda$, and aim to modify the instance so as to guarantee correct outcomes on p_j and p_k ; this ensures that we obtain correct outcomes on all proposals, except possibly for the proposals in the set $\{p_{j+1}, \dots, p_{k-1}\}$.

For EAP and DAP, we can decide if a suitable modification exists by running our FPT algorithm (Proposition 1) on the set of proposals $\{p_j, p_k\}$; note that this algorithm runs in polynomial time when there are just two proposals.

Unfortunately, this strategy fails for CLD, because proposals other than p_j and p_k influence which edges can appear in a delegation graph. Hence, we use a different approach. We say that an agent i is *I-correct on proposal p_ℓ* if $\text{corr}(p_\ell, I) = 0$ and $p_\ell \in P^{-1}(i)$, or $\text{corr}(p_\ell, I) = 1$ and $p_\ell \in P^+(i)$. Also, we say that a delegation graph \mathcal{D} is (j, k) -good if in the instance $I'(\mathcal{D})$ more than $n/2$ agents are *I-correct on p_j* and more than $n/2$ agents are *I-correct on p_k* .

To check if there is a (j, k) -good delegation graph, we partition the set N as $N = N_A \cup N_B \cup N_C \cup N_D$ so that

- agents in N_A are *I-correct on both p_j and p_k* ;
- agents in N_B are *I-correct on p_j* , but not on p_k ;
- agents in N_C are *I-correct on p_k* , but not on p_j ; and
- agents in N_D are *I-correct on both p_j and p_k* .

Suppose there exists a (j, k) -good delegation graph \mathcal{D} . Then, in $I'(\mathcal{D})$, more than $n/2$ agents are *I-correct on p_j* and more than $n/2$ agents are *I-correct on p_k* . As these two sets of agents must overlap, there is an agent in $I'(\mathcal{D})$ that is *I-correct on both p_j and p_k* . Since delegation amounts to agents copying the preferences of their gurus, this means that there is some such agent in I , i.e., N_A must be non-empty.

Thus, if $N_A = \emptyset$, there is no (j, k) -good delegation graph. Assume, then, that N_A is non-empty. Note that all agents in N_D must be uncertain about p_j and p_k (and hence about all proposals in between) and are therefore willing to delegate to agents in N_A , who are certain about both. Further, if there is a (j, k) -good delegation graph, then there is one in which all agents who can delegate to an agent in N_A do so (and thereby copy the preferences of their guru). Hence, as our next step, we repeatedly check whether $N_B \cup N_C \cup N_D$ contains an agent who is willing to delegate to an agent in N_A ; if so, then we move this agent to N_A . This process stops when no such agent remains. Observe that at this point $N_D = \emptyset$.

Now, if $|N_B| + |N_A| > n/2$, then there is a majority of agents who are *I-correct on p_j* , and if $|N_C| + |N_A| > n/2$, then there is a majority of agents who are *I-correct on p_k* ; so, if both of these conditions are satisfied, then we are done. Assume without loss of generality that $|N_C| \geq |N_B|$. As we have $|N_B| + |N_A| + |N_C| + |N_A| = n + |N_A| > n$, the condition $|N_C| + |N_A| > n/2$ is satisfied in this case.

Suppose, however, that $|N_B| + |N_A| \leq n/2$. Then, the only way to obtain the correct outcome on both p_j and p_k is to find $t = \lceil \frac{n}{2} \rceil - |N_A| - |N_B|$ agents in N_C who can delegate to agents in N_B ; indeed, in this case, after the delegation we will have $\lceil \frac{n}{2} \rceil$ agents who are *I-correct on p_j* and $|N_C| - t + |N_A| = n + |N_A| - \lceil \frac{n}{2} \rceil \geq \lceil \frac{n}{2} \rceil$ agents who are *I-correct on p_k* (where the last inequality holds because $|N_A| \geq 1$).

To determine whether these t delegating agents could be found, we construct a graph G on $N_B \cup N_C$ with an edge from agent i to agent i' if and only if i is willing to delegate to i' , and mark each agent in N_C that has a path to an agent in N_B ; the answer is positive if and only if at least t agents are marked. This completes the proof. \square

6 Outlook

We have considered three algorithmic approaches to coping with confused communities that wish to reach good joint decisions, and showed that, even though our problems are generally NP-complete, there are special cases for which efficient algorithms exist.

Avenues for future research include the following: (1) a study of other proposal spaces, including, in particular, continuous proposal spaces; (2) extending our analysis to settings where proposals may have different importance and/or issues are non-binary; (3) an exploration of different approaches to coping with voter uncertainty, including models in which the effect of education efforts is probabilistic and may propagate through an underlying social network; (4) the development of computer-based simulations for the practical evaluation of the effectiveness of our algorithmic approaches.

References

- [Bachrach *et al.*, 2010] Yoram Bachrach, Nadja Betzler, and Piotr Faliszewski. Probabilistic possible winner determination. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 697–702, 2010.
- [Bloembergen *et al.*, 2019] Daan Bloembergen, Davide Grossi, and Martin Lackner. On rational delegations in liquid democracy. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 1796–1803, 2019.
- [Boudreau *et al.*, 2019] Cheryl Boudreau, Christopher S El-mendorf, and Scott A MacKenzie. Roadmaps to representation: An experimental study of how voter education tools affect citizen decision making. *Political Behavior*, 41(4):1001–1024, 2019.
- [Brill and Talmon, 2018] Markus Brill and Nimrod Talmon. Pairwise liquid democracy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 137–143, 2018.
- [Bullock, 2011] John G Bullock. Elite influence on public opinion in an informed electorate. *American Political Science Review*, 105(3):496–515, 2011.
- [Cohensius *et al.*, 2017] Gal Cohensius, Shie Manor, Reshet Meir, Eli Meir, and Ariel Orda. Proxy voting for better outcomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 858–866, 2017.
- [Elkind and Lackner, 2015] Edith Elkind and Martin Lackner. Structure in dichotomous preferences. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2019–2025, 2015.
- [Faliszewski and Rothe, 2016] Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 7, pages 146–168. Cambridge University Press, 2016.
- [Garey *et al.*, 1976] Michael R. Garey, David S. Johnson, and Larry Stockmeyer. Some simplified NP-complete problems. *Theoretical Computer Science*, 1(3):237–267, 1976.
- [Gölz *et al.*, 2021] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D. Procaccia. The fluid mechanics of liquid democracy. *ACM Transactions on Economics and Computation*, 9(4):23:1–23:39, 2021.
- [Green-Armytage, 2015] James Green-Armytage. Direct voting and proxy voting. *Constitutional Political Economy*, 26(2):190–220, 2015.
- [Hazon *et al.*, 2008] Noam Hazon, Yonatan Aumann, Sarit Kraus, and Michael J. Wooldridge. Evaluation of election outcomes under uncertainty. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 959–966, 2008.
- [Jr., 1983] Hendrik W. Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.
- [Kahng *et al.*, 2021] Anson Kahng, Simon Mackenzie, and Ariel Procaccia. Liquid democracy: An algorithmic perspective. *Journal of Artificial Intelligence Research*, 70:1223–1252, 2021.
- [Kenig and Kimelfeld, 2019] Batya Kenig and Benny Kimelfeld. Approximate inference of outcomes in probabilistic elections. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 2061–2068, 2019.
- [Lupia and McCubbins, 1998] Arthur Lupia and Mathew D. McCubbins. *The democratic dilemma: Can citizens learn what they need to know?* Cambridge University Press, 1998.
- [Walsh, 2007] Toby Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 3–8, 2007.
- [Wojtas and Faliszewski, 2012] Krzysztof Wojtas and Piotr Faliszewski. Possible winners in noisy elections. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1499–1505, 2012.