



Flexible representative democracy

An introduction with binary issues

Ben Abramowitz¹  · Nicholas Mattei¹

Received: 15 May 2022 / Accepted: 25 July 2024 / Published online: 9 September 2024
© The Author(s) 2024

Abstract

We introduce Flexible Representative Democracy (FRD), a novel hybrid of Representative Democracy and Direct Democracy in which voters can alter the issue-dependent weights of a set of elected representatives. In line with the literature on Interactive Democracy, our model allows the voters to actively determine the degree to which the system is direct versus representative. However, unlike Liquid Democracy, Flexible Representative Democracy uses strictly non-transitive delegations, making delegation cycles impossible, and maintains a fixed set of accountable, elected representatives. We present Flexible Representative Democracy and analyze it using a computational approach with issues that are binary and symmetric. We compare the outcomes of various voting systems using Direct Democracy with majority voting as an ideal baseline. First, we demonstrate the shortcomings of Representative Democracy in our model. We provide NP-Hardness results for electing an ideal set of representatives, discuss pathologies, and demonstrate empirically that common multi-winner election rules for selecting representatives do not perform well in expectation. To analyze the effects of adding delegation to representative voting systems, we begin by providing theoretical results on how issue-specific delegations determine outcomes. Finally, we provide empirical results comparing the outcomes of various voting systems: Representative Democracy, Proxy Voting, and FRD with issue-specific delegations. Our results show that variants of Proxy Voting yield no discernible benefit over unweighted representatives and reveal the potential for Flexible Representative Democracy to improve outcomes as voter participation increases.

✉ Ben Abramowitz
babramow@tulane.edu

Nicholas Mattei
nsmattei@tulane.edu

¹ Tulane University, New Orleans, Louisiana, USA

1 Introduction

Direct Democracy, in which all citizens or stakeholders vote directly, is often held as an ideal form of governance (Dunn 1995). However, Direct Democracy can be impractical at scale because it places too great a burden on voters (Black 2012; Green-Armytage 2015). Voting on a large number of important, complex issues requires a great deal of time, effort, knowledge, attention, and communication. Given the prohibitive costs of implementing a large-scale Direct Democracy, institutions often resort to forms of representation, relying upon a set of representatives to make decisions on behalf of the population as a whole. Countries have parliaments, companies have shareholders, and even groups of robotic agents select leaders to represent them (Yu et al. 2010). In principle, the use of Representative Democracy eases the heavy burden of deliberation and decision-making for the general population.

Representative Democracy also scales better than Direct Democracy as the total cost of making each decision scales with the number of representatives rather than the size of the overall population. However, representation is not without its downsides. The set of representatives may make different decisions than the citizens would and it can be difficult to hold representatives accountable. Representative Democracy also fails to take advantage of full voter participation in decision-making when it is available.

The limitations and trade-offs of Direct Democracy and Representative Democracy have given rise to the study of various democratic decision-making systems under the umbrella of *interactive democracy* (Brill 2018). This includes variations of delegative voting or proxy voting. Transitive proxy voting, or Liquid Democracy, in particular has received significant attention in the political science (Green-Armytage 2015), AI (Kahng et al. 2021; Gölz et al. 2021; Brill et al. 2022) and multiagent systems (Brill and Talmon 2018) communities, and has been implemented in both corporate (Hardt and Lopes 2015) and political (Blum and Zuber 2016; Behrens et al. 2014) settings. However, the transitivity of vote delegations creates new problems like the possibility of delegation cycles (Brill and Talmon 2018). It is worth noting that many proposed mechanisms for large-scale collective decision-making are only made practical by modern communication technology.

We propose a class of mechanisms we call Flexible Representative Democracy (FRD) capable of interpolating between Direct Democracy and Representative Democracy at the discretion of the voters. Similar to a traditional Representative Democracy, in FRD voters elect a set of representatives to serve a term during which they (publicly) deliberate and make decisions. However, unlike traditional Representative Democracy, in FRD the representatives can be weighted. The weights of the representatives are determined by some default weighting scheme and can then be updated based on delegations by voters. The delegations by voters, and hence the weights of the representatives, can differ over time and across issues.

In FRD, voters have great flexibility in determining how they are represented. For example, the day after the election an inattentive voter might choose a few elected representatives they trust, divide their delegation (i.e., unit of voting power) among these representatives for all future issues, and pay no attention until the next election. A more attentive voter might alter their delegation on an issue-by-issue basis as issues arise, reacting to representatives' deliberations, pronouncements, and publicized votes.

Suppose the default weighting is uniform across all representatives and the weight of a representative is the sum of weights assigned to them by voters. If no voter uses their option to delegate on any issue, the system behaves exactly like a traditional Representative Democracy as all representatives' votes have the same weight. On the other hand, if every voter delegates on every issue to a representative who votes exactly as the voter would, then the system can precisely emulate Direct Democracy as each voter's vote is counted as if they had voted directly. Hence, FRD takes advantage of however much burden voters are willing to take on without raising the minimum burden required beyond that of Representative Democracy. The degree to which FRD emulates Direct Democracy depends on both the caliber of the representatives and the fastidiousness of the voters.

Contributions We introduce Flexible Representative Democracy (FRD), a new class of voting systems that gives voters greater power over how they are represented and influence decision outcomes.¹ Our proposal for FRD solves standing issues in the literature on interactive democracy including maintaining a fixed, elected committee to generate legislation and making delegation cycles impossible. We analyze an FRD mechanism and related voting systems using a toy model for deciding binary, symmetric issues where the objective is to achieve the decision outcomes that would occur in an ideal Direct Democracy with full participation. Our primary findings are: (1) electing an optimal set of representatives is (NP-)hard for any Representative Democracy using a multi-winner voting rule, (2) the choice of election rule appears less important than the quality of the electoral candidates and fastidiousness of voters, and (3) assigning representatives issue-specific weights based on voter delegations can yield significant improvements in the quality of decision outcomes, but giving representatives a constant weight across all issues does little to overcome the limitations of representation.²

2 Related work

Miller (1969), inspired by Tullock (1967) and shareholder proxy voting, suggested an interactive democratic system for legislation that could take place at scale using computers. Miller lamented the lack of *flexibility* in traditional Representative Democracy and sought to remedy this using a dynamic system of proxies, though admitting this was not conducive to creating legislation. Soon after, Shubik (1970) warned that electronic systems may accelerate the legislative process in undesirable ways and suggested holding every referendum twice to guarantee time for sufficient public deliberation. FRD's use of a fixed, elected set of representatives answers Miller's question of how to produce legislation, and rather than holding redundant referenda we propose to give the voters sufficient time to continue deliberation and alter their delegations after the representatives cast immutable public votes.

¹ The proposal for FRD and many of our results first appeared in Abramowitz and Mattei (2019).

² Code for our experiments can be found here: <https://github.com/BenAbramowitz/FRD>.

Just before the dawn of the internet, Tullock (1992) revisited these ideas in a proposal that motivates the default distribution and delegation mechanism in FRD. The notion of the default distribution is also similar to the electoral weighting scheme proposed by Alger (2006), in which the weights of representatives are based on the preferences of voters expressed in the election, but these weights are fixed during the representatives' term. By contrast, in FRD the weight of each representative on each issue is not strictly determined by the election.

The hallmark of an interactive democracy is that rather than adjudicating whether a direct or representative system is better for expressing the will of the voters and asserting it by fiat, the extent to which the system is direct or representative is itself a function of the will of the voters. Currently, the most well-known and well-studied form of interactive democracy is Liquid Democracy, which has been studied from an algorithmic perspective as a decision-making process in the AI and Computational Social Choice (COMSOC) literature (Brill and Talmon 2018; Kahng et al. 2021; Bloembergen et al. 2019; Christoff and Grossi 2017; Escoffier et al. 2019; Colley et al. 2021; Becker et al. 2021; Markakis and Papasotiropoulos 2021; Colley 2021; Colley et al. 2022) and elsewhere (Green-Armytage 2015; Ford 2002; Harding 2022; Blum and Zuber 2016; Brill 2018; Hardt and Lopes 2015; Harding 2019; Gersbach et al. 2022). Unlike Liquid Democracy, FRD does not allow transitive delegations nor delegations to another voter, thereby violating the second axiom proposed by Green-Armytage (2015). However, the notion of voluntary representatives can be maintained in FRD if desired. Fractional delegations in FRD serve a similar function to that of the virtual committees proposed by Green-Armytage (2015), although, in theory, FRD could incorporate virtual committees as well as many other mechanisms for delegating voting power. We note that Cohensius et al. (2017) took an analytical approach to studying a Proxy Voting model very close to that of Alger (2006) for decision-making with no election, infinite voters, spatial preferences, and assuming agents lie in a metric space.

The design of FRD is also largely based on work in probabilistic voting, binary aggregation, statistical decision theory, and computational social choice. In particular, work on the optimal weighting of experts (Baharad et al. 2012; Nitzan and Paroush 2017; Grofman and Feld 1983; Nitzan and Paroush 1982; Ben-Yashar and Nitzan 1997), the Condorcet Jury Theorem (Grofman et al. 1983), variable electorates (Feld and Grofman 1984; Smith 1973; Paroush and Karotkin 1989), and optimal committee sizes (Auriol and Gary-Bobo 2012; Karotkin and Paroush 2003; Magdon-Ismail and Xia 2018). In FRD, one can view the voter delegations as a pseudo-tie breaking mechanism for the representatives or, conversely, see the default distribution as a way to dampen the variance in the outcome which occurs in Direct Democracy when the sample of participating voters is small or biased. Another view is that electing representatives is analogous to a compression algorithm (Rodriguez and Steinbock 2004). In this view, the delegations in FRD are a decompression mechanism where a higher delegation rate reduces the “loss” of representation.

The recent works most similar to ours are that of Pivato and Soh (2020), Soh Voutsas (2020), and Meir et al. (2021), which each examine weighted representative voting systems in terms of their agreement with the voter majority. However, each considers weightings of representatives that are fixed across a set of binary issues. Our evalua-

tions are also similar to those of Skowron (2015), however, in their approval model the quality of the committee is measured as the sum of the voter proportion being represented for each issue, while we focus only on the total number of issue outcomes in alignment with the voter majority.

In our work we compare representative voting systems according to the outcomes they produce, but in practice it is hard to evaluate the quality of decisions, especially when this involves reasoning about counterfactuals. It is often useful to use a surrogate measure for the quality of a representative body based on how the attributes of this set of representatives relates to the set of voters. John Adams's intuition was that the representatives should be a microcosm of the population (Adams 1979), but he did not formally define what this meant. Many different axioms for multi-winner voting have been proposed in recent years (Revel et al. 2023; Lackner and Skowron 2023). Some popular lines of thought are that committees should be proportional and/or have justified representation of the voters (Aziz et al. 2017; Elkind 2023) or that representation can be based on spatial preference models (Anslevich et al. 2021). However, when representatives are weighted by delegation, any group of voters can ensure a form of proportional representation through delegation if they have at least one representative to represent them, thereby reducing the need for the election rule alone to achieve these forms of proportionality and representation.

Within computer science, many applications face the task of selecting representatives for downstream decision-making. In portfolio selection, a particular set of algorithms and hyper-parameters are selected from a large pool of candidates and then used as representatives for later problems (KhudaBukhsh et al. 2016). In multi-agent systems, the role assignment problem uses distributed voting to decide on tasks for agents (Zhu et al. 2012). And in group recommendation settings, elections correspond to picking a set of experts to make decisions later. The COMSOC community has produced a large body of research on how to select and weight representatives (Brandt et al. 2016). Indeed, using multi-winner voting (Skowron et al. 2016; Revel et al. 2023), we can view the winners as a set of exemplars that may be used to decide some downstream application—e.g., we select a set of points in space and then aggregate these points (votes) over the set. Finally, we note that while Direct Democracy may place too high a demand on the general population when there are too many complex issues, preference learning is an active research area in AI (Domshlak et al. 2011; Fürnkranz and Hüllermeier 2003).

3 Preliminaries

In this section we give an overview of voting systems, including our formal model for Flexible Representative Democracy, and our desiderata with binary issues.

3.1 Voting systems

Direct democracy In a Direct Democracy, every voter can express their preferences directly on every issue that is to be decided. Their preferences are then aggregated to

select an outcome for each issue. Let V be a set of voters and S a set of issues. Voters express their preferences for each issue $s \in S$. The collection of voter preferences regarding s is denoted P_V^s , and the collective profile of preferences over the issues as a whole is P_V . We denote by R the *decision rule* that aggregates voter preferences into outcomes. For brevity, we assume that the same decision rule R is applied separately for all issues. The outcome on issue s is $R(P_V^s)$, and the vector of outcomes is denoted by the shorthand $R(P_V) = (R(P_V^1), \dots, R(P_V^{|S|}))$.

Representative democracy In a Representative Democracy (RD), the voters do not report their preferences regarding each issue. Instead, the voters elect a set of representatives who then vote on behalf of the whole population for each issue. Initially, there is a slate of candidates C who vie for election. The voters express preferences over the candidates, collectively denoted by the profile P_{VC} . An *election rule* E then selects a subset of candidates as representatives: $E(P_{VC}) = D \subseteq C$. Once elected, the representatives cast their votes on each issue in a set S , forming profile $P_D = (P_D^1, \dots, P_D^{|S|})$. Finally, the set of outcomes is determined by applying a decision rule $R(P_D) = (R(P_D^1), \dots, R(P_D^{|S|}))$.

Weighting representatives Many decision rules treat the votes from different agents equivalently. In other words, the decision rules are *anonymous* in that it does not matter which agent each vote comes from. One way to relax anonymity is to weight the agents and use a weighted decision rule. Our focus is on voting systems that weight representatives. A representative could have the same weight across all issues or a different weight for each issue. Letting $W = (W^1, \dots, W^{|S|})$ represent the weighting of the representatives, the decision outcomes are computed individually $R(P_D^s, W^s)$, with the vector of outcomes denoted in shorthand by $R(P_D, W)$.

Electoral weighting One way to weight representatives is based on the electoral profile P_{VC} . Specifically, we take there to be some function f that maps electoral profiles to weights: $f(P_{VC}) = W$, and these weights are used to compute the decision outcomes $R(P_D, W)$. In this sort of weighting scheme, the weights of representatives remain constant across issues, and we consider this a sub-type of RD.

Proxy voting In proxy voting, there is a set D of potential proxies for each voter to choose from. The proxies may or may not be elected. Each voter selects a single proxy, and the proxies are given weights based on how many voters choose them. As with RD, we will identify a subroutine f as translating proxy choices into weights. If we denote the proxy choices of the voters as profile P_{VD} , then with some weighting function $f(P_{VD}) = W$, the outcomes are $R(P_D, W)$. Proxy voting can be generalized by allowing each voter to select more than one proxy or assigning scores to the proxies. Unlike RD with electoral weighting, in proxy voting the weights are independent of how the set of representatives D is constructed, and depends instead on preferences over the representatives expressed after the set of representatives is determined. Proxy voting is frequently used on an issue-by-issue basis, where voters can report P_{VD}^s for each issue s rather than using a single weighting of representatives across all issues.

However, in proxy voting systems, voters typically do not have the ability to alter their proxy choices after the proxies vote.

Flexible representative democracy In a Flexible Representative Democracy (FRD), the voters begin by electing a set of representatives $D = E(P_{VC})$. Voters then have the option to express preferences over the elected representatives, e.g., choose proxies, on an issue-by-issue basis (P_{VD}^s). The outcomes are then determined by $R(P_{VC}, P_{VD}, P_D)$.

We focus on FRD voting systems that weight representatives and use a weighted voting rule. Namely, the outcome of each issue s is then determined by $R(P_D^s, W^s)$ where $W^s = f(P_{VC}, P_{VD}^s)$ for some weighting protocol f . Once again, the vector of outcomes across the issues is denoted $R(P_D, W)$.

While we do not address the temporal aspects of the voting systems in our work, one feature of FRD to keep in mind is that if the representatives fix their votes publicly on each issue before the decision rule is applied, this affords the voters an opportunity to update P_{VD}^s with knowledge of P_D^s before each decision is made.

Default and delegation The focus of our analysis is on a form of FRD in which the weight of each representative for each issue is the sum of weights assigned to them by the voters. The weights from each voter are either determined passively by default or actively by delegation. We only consider a version of FRD that maintains one-person-one-vote among voters, so each voter assigns a single unit of voting weight distributed over the representatives. The default weighting scheme we focus on is the uniform distribution. After the election, the unit of voting weight from each voter is spread equally among the representatives. Therefore, each representative receives $\frac{1}{|D|}$ weight from each voter on each issue by default, and all representatives have a total weight of $\frac{|V|}{|D|}$. Since all representative weights are equal, if no voter were to use their delegation option, then it would be equivalent to an unweighted RD. When agents delegate, they update their assignment of weight to the representatives, though the weights they assign must still sum to unity. Thus if every voter delegates all of their weight to representatives who vote identically to them on every issue, then the behavior is exactly equivalent to Direct Democracy.

Assumptions and simplifications For our analysis, we assume for brevity that all voters take part in the election, so that P_{VC} is complete. However, we do not assume that voters are active on every issue, so they do not necessarily make use of their option to delegate. The weighting function f accounts for lack of delegation using a default weighting. While one can think of the weighting function f as being a subroutine of the decision rule R , it simplifies our definitions and notation to separate them. Note that while we focus on voting rules that use weighting of representatives, this is not strictly required as part of the general definition of FRD, which can make decisions of the general form $R(P_{VC}, P_{VD}, P_D)$ without constructing an explicit weighting.

3.2 Binary issues model

Let $V = \{v_1, \dots, v_n\}$ be an indexed set of voters, and $S = \{s^1, \dots, s^r\}$ be an indexed set of binary issues. The two possible outcomes for each issue are denoted $\{0, 1\}$. On each issue s^i , each voter v_j has a preference $v_j^i \in \{0, 1\}$. The full preferences of voter v_j across the issues are denoted $\mathbf{v}_j = \{v_j^1, \dots, v_j^r\}$ and $P_V = \{\mathbf{v}_j : v_j \in V\}$ denotes the profile of voters' preferences as a whole. Without loss of generality, we relabel the alternatives so that the outcome preferred by the voter majority is 1 and the other is 0 on every issue.

Let $D = \{d_1, \dots, d_k\}$ be an indexed set of representatives with preferences $\mathbf{d}_l = \{d_l^1, \dots, d_l^r\} \in \{0, 1\}^r$ comprising the preference profile $P_D = \{\mathbf{d}_l : d_l \in D\}$. We will assume that the number of representatives k is always odd to eliminate the need for tie-breaking when representatives are unweighted in RD or equally weighted in FRD.

We hold Direct Democracy with majority voting and full participation as our ideal for comparison. Therefore, after relabeling the preference of the voter majority to be 1 on every issue, the ideal outcome is **1**. We will be evaluating the performance of FRD and other voting systems in terms of the fraction of issues on which they produce the outcome corresponding to the voter majority.

Definition 1 (Agreement) Let \mathcal{S} be the set of all binary vectors of length r . For any two vectors $\mathbf{x}, \mathbf{y} \in \mathcal{S}$, we measure their similarity or *agreement* $L(\mathbf{x}, \mathbf{y})$ by the fraction of their entries that are the same.³ In other words, the agreement between two vectors is inverse to the normalized correct only Hamming Distance between them:

$$L(\mathbf{x}, \mathbf{y}) = 1 - \frac{1}{r} \sum_{i=1}^r |x^i - y^i|$$

If \mathbf{x} is a vector of outcomes yielded by some voting procedure, we measure its *agreement* with the voter majority $L(\mathbf{x}, \mathbf{1})$.

Definition 2 (Majority Agreement) The majority agreement between a set of representatives D and the set of voters V is the fraction of issues on which the majority of representatives agrees with the voter majority, i.e., $L(R(P_V), R(P_D))$.

Note that with R as majority rule and an odd number of unweighted candidates (to prevent tie-breaking),

$$L(R(P_V), R(P_D)) = \frac{1}{|S|} |\{s^i \in S : \sum_{d_l \in D} d_l^i > \frac{|D|}{2}\}|.$$

We also define a weaker notion than majority agreement that we call *coverage*. An issue is covered if at least one representative agrees with the voter majority. Majority agreement on any issue implies coverage of that issue, but coverage does not imply

³ Recall that the order of the binary issues does not matter in our model and the issues are treated identically, so other common measures of distance/similarity used in the Social Choice literature are either equivalent to the Hamming Distance or not applicable.

majority agreement. Coverage characterizes the potential for delegation to yield an improvement over majority agreement. We will generally talk about coverage in terms of the fraction of issues that are covered by a set of representatives.

Definition 3 (Coverage) A set of representatives D covers an issue s^i if $\exists d_l \in D$ such that $d_l^i = 1$.

4 Difficulties of representation

In the real world, the outcomes preferred by the voter majority are not known. However, even if it were known, electing a good set of representatives for RD would still be hard due to computational complexity. This complexity barrier shows that even if voter preferences can be sampled to ascertain the majority opinion with high probability (i.e., by polling), RD is still limited in how well it can recover the “will of the voters”.

4.1 Complexity with full information

In this section we prove that for a set of candidates C with binary preferences over the issues S , it is NP-Hard to compute the subset of k candidates that maximizes majority agreement. To make matters worse, even maximizing coverage is NP-Hard. In practice we care about maximizing these values than checking if they meet certain thresholds, and therefore cast the problems as optimization problems rather than decision problems when evaluating their complexity. We refer to the problems of electing k candidates to maximize coverage and majority agreement when the preference of the voter majority is known for each issue as *Max k-Coverage* and *Max k-Majority Agreement*, respectively.

We begin by considering Max k -Coverage, and will use the NP-hardness of Max k -Coverage to prove the NP-hardness of Max k -Majority Agreement. In later sections, when we consider mechanisms that weight representatives, coverage of an issue will become the minimum necessary condition for there to exist a weighting that leads the outcome to agree with the voter majority.

Problem 1 (Max k -Coverage) Let $S = \{s^1, \dots, s^r\}$ be a set of binary issues and $C = \{c_1, \dots, c_m\}$ a set of candidates where candidate c_l has preference $c_l^i \in \{0, 1\}$ on issue s^i . The problem of Max k -Coverage is the problem of computing a subset of $k \leq m$ representatives $D \subseteq C$ that maximizes the number of covered issues, where issue $s^i \in S$ is covered if $\sum_{d_l \in D} d_l^i > 0$.

Theorem 1 *Max k -Coverage is NP-Hard.*

Proof Our proof of the hardness of Max k -Coverage is a Karp reduction via the NP-Hard problem of MAX K-COVER (Feige 1998). The inputs to MAX K-COVER are a set $S = \{s^1, \dots, s^r\}$ of r points, a collection $C = \{c_1, \dots, c_m\}$ of subsets of S , and an integer k . The objective of MAX K-COVER is to select k subsets from C such that their union has maximum cardinality. Given an instance (S, C, k) of MAX K-COVER, we create an instance $(\tilde{S}, \tilde{C}, \tilde{P}_C, k)$ of Max k -Coverage as follows.

For every point $s^i \in S$ create an issue $\tilde{s}^i \in \tilde{S}$, and for every subset $c_l \in C$ create a candidate $\tilde{c}_l \in \tilde{C}$. For all points $s^i \in S$ and subsets $c_l \in C$, if $s^i \in c_l$ then let $\tilde{c}_l^i = 1$, otherwise let $\tilde{c}_l^i = 0$. Let k be the number of representatives to be elected. There is a one-to-one correspondence between the number of issues covered by our k representatives and the cardinality of the corresponding subsets in the original MAX K-COVER instance. Therefore, any set of k candidates that maximizes coverage corresponds exactly to a collection of k subsets in our MAX K-COVER instance whose union has maximum cardinality. This concludes our proof of Theorem 1. \square

Problem 2 (Max k -Majority Agreement) Let $S = \{s^1, \dots, s^r\}$ be a set of binary issues and $C = \{c_1, \dots, c_m\}$ a set of candidates where candidate c_l has preference $c_l^i \in \{0, 1\}$ on issue s^i . The problem of Max k -Majority Agreement is the problem of computing a subset of $k \leq m$ representatives $D \subseteq C$, where k is odd, that maximizes the number of issues on which the majority of representatives prefers 1, i.e.

$$D = \arg \max_{C' \subseteq C, |C'|=k} \left\{ s^i \in S : \sum_{c_l \in C'} c_l^i > \frac{|D|}{2} \right\}$$

Theorem 2 *Max k -Majority Agreement is NP-Hard.*

Proof Our proof of the hardness of Max k -Majority Agreement is a Karp reduction via the problem of Max k -Coverage that we proved to be NP-Hard in Theorem 1. We will take an instance of Max k -Coverage and add $r + 1$ additional issues to the original r issues. In addition, we augment the candidate set with $k + 1$ additional candidates in such a way that any subset of $2k + 1$ candidates that maximizes agreement must contain the $k + 1$ added candidates. The remaining k winning candidates will correspond to the winning candidates for the original Max k -Coverage instance.

Suppose we have an instance of Max k -Coverage (S, C, P_C, k) with $r = |S|$ issues and $m = |C|$ candidates. We construct an instance of Max k -Majority Agreement $(\tilde{S}, \tilde{C}, \tilde{P}_C, |\tilde{D}|)$ as follows. Let $\tilde{S} = (\tilde{s}^1, \dots, \tilde{s}^{2r+1})$ be a set of $2r + 1$ binary issues. Let $\tilde{C} = (\tilde{c}_1, \dots, \tilde{c}_{m+k+1})$ be a set of $m + k + 1$ candidates. \tilde{C} is made up of three types of candidates based on how we construct their preferences $\tilde{C} = (\tilde{c}_1, \dots, \tilde{c}_m) \cup (\tilde{c}_{m+1}, \dots, \tilde{c}_{m+k}) \cup (\tilde{c}_{m+k+1})$. The first m candidates have the same preferences as the original m candidates on the first r issues and prefer 0 for the added $r + 1$ issues. That is, $\forall l \leq m, \forall i \leq r, \tilde{c}_l^i = c_l^i$ and for all $i > r, \tilde{c}_l^i = 0$. The next k candidates unanimously prefer 1 on all issues: $\forall l$ such that $m < l \leq m + k, \forall i \leq 2r + 1, \tilde{c}_l^i = 1$. The final candidate, \tilde{c}_{m+k+1} , prefers 0 on the first r issues and 1 on the remaining $r + 1$ issues. The construction of the candidate profile is illustrated in Table 1. Lastly, let $|\tilde{D}| = 2k + 1$.

Observe that any set $\tilde{D} \subseteq \tilde{C}$ of $2k + 1$ candidates that maximizes majority agreement must contain candidates $(\tilde{c}_{m+1}, \dots, \tilde{c}_{m+k})$ and \tilde{c}_{m+k+1} . The majority of \tilde{D} prefers 1 on issues $r + 1$ through $2r + 1$ if and only if it contains $(\tilde{c}_{m+1}, \dots, \tilde{c}_{m+k})$ and \tilde{c}_{m+k+1} . If even one of these $k + 1$ candidates is not included in \tilde{D} , then the majority of \tilde{D} will not prefer 1 on any of these $r + 1$ issues. Since this constitutes more than half of the total issues, all of those candidates must be included in \tilde{D} .

The question remains how to select an additional k candidates such that \tilde{D} maximizes majority agreement on the first r issues. The k candidates $(\tilde{c}_{m+1}, \dots, \tilde{c}_{m+k})$

Table 1 Construction of Max k -Majority Agreement instance from candidates' profile in original Max k -Coverage instance

| | s^1 | ... | s^r | s^{r+1} | ... | s^{2r+1} |
|-------------|----------|----------|----------|-----------|----------|------------|
| c_1 | c_1^1 | ... | c_1^r | 0 | ... | 0 |
| c_2 | c_2^1 | ... | c_2^r | 0 | ... | 0 |
| \vdots | \vdots | \vdots | \vdots | \vdots | \vdots | \vdots |
| c_m | c_m^1 | ... | c_m^r | 0 | ... | 0 |
| c_{m+1} | 1 | ... | 1 | 1 | ... | 1 |
| \vdots | \vdots | \vdots | \vdots | \vdots | \vdots | \vdots |
| c_{m+k} | 1 | ... | 1 | 1 | ... | 1 |
| c_{m+k+1} | 0 | ... | 0 | 1 | ... | 1 |

prefer 1 on the first r issues but \tilde{c}_{m+k+1} prefers 0, so to create a majority within \tilde{D} on any of those issues requires one or more of the other k winning candidates to prefer 1. In other words, it is exactly the problem of using k out of the original m candidates to cover as many of the original r issues as possible. Therefore, our solution to Max k -Majority Agreement on our constructed instance gives us the solution to the original Max k -Coverage problem. This concludes our proof of Theorem 2. \square

We can see that no polynomial-time algorithm can maximize majority agreement, even with complete information about voters' issue preferences. Thus, election rules, which typically have only second-hand information in the form of electoral preferences, can only hope to offer approximations for majority agreement.

4.2 Majority agreement from elections with partial information

We have shown that even when the preferences of the voter majority are known for every issue, it is computationally hard to select the optimal set of representatives of a fixed size from a set of candidates. While solving Max k -Majority Agreement may be tractable for small enough instances, this relies on knowing the voter majority on every issue, which is typically not the case. We now turn our attention to the question of how well election rules can approximate coverage and majority agreement in the partial information setting when only voter preferences over candidates are known.

We give the election rules the best chances we can by assuming that voters' preferences over candidates are derived exactly from their agreement on the issues. For brevity, we assume where necessary that the sets of voters and candidates are disjoint because this assumption has little to no impact on our results. Voters can express their electoral preferences P_{VC} as either approvals or complete rankings over the candidates. When voters report approvals over the candidates, we assume that v approves of c if and only if they agree on more than half of the issues, i.e., $L(v, c) > \frac{1}{2}$. Observe that we do not restrict the number of approvals a voter gives, so a voter may approve all or none of the candidates. When voters provide rankings of the candidates, we assume that v prefers c to \hat{c} if they agree on a greater number of issues, i.e., $L(v, c) > L(v, \hat{c})$. If $L(v, c) = L(v, \hat{c})$, then v breaks the tie privately (e.g., randomly). When issue prefer-

ences are uniformly random, these induced rankings with random tie-breaking follow the impartial culture model because all electoral profiles are equally likely (Tsetlin et al. 2003).

A Condorcet-consistent election rule elects a candidate as a representative if they beat all other candidates in pairwise competition. Unfortunately, regardless of whether agents report approvals or rankings, no Condorcet-consistent election rule can approximate majority agreement.

Theorem 3 *No Condorcet-consistent election rule for selecting k candidates in which voters report complete rankings over the candidates induced by their preferences over issues can provide a bounded worst-case approximation of majority agreement or coverage.*

Theorem 3 is a consequence of the Ostrogorski paradox (Kelly 1989; Rae and Daudt 1976). The relationship between the existence of Condorcet-winners, issue-wise majority voting, and presence of the Ostrogorski paradox has been previously established (Laffond and Lainé 2009; Bezembinder and Van Acker 1985), and we reformulate part of this result in terms of the impossibility of approximation algorithms in Theorem 3. The Ostrogorski paradox can occur with as little as 5 voters and 3 issues, as shown in Table 3, but we choose to use an instance derived from the foundational work of Anscombe (1976) for the purposes of our proof.

Proof The example in Table 2 shows the Ostrogorski paradox with 11 voters, 11 issues, 2 candidates, and $k = 1$ and proves Theorem 3. We have two candidates; the ideal candidate $\mathbf{c}_1 = \mathbf{1}$ and the worst conceivable candidate $\mathbf{c}_2 = \mathbf{0}$. Each row represents the preference vector of a voter or candidate, with a column for each issue. The last column shows the electoral preferences of the voters over the candidates. Since the two candidates have exactly opposite preferences, no voter can approve both, and thus the approvals and rankings can both be described by giving just the voters' top candidates. The last row shows the preference of the voter majority. If we sum the columns to get the voter majority (maj) we see that the voter majority supports 1 on every issue. However, if we look at each voter's preferences over candidates we can see that 7 out of 11 voters prefer c_2 over c_1 . Since c_1 and c_2 have opposite preferences on every issue, no voter will ever approve of both candidates. Any Condorcet-consistent election rule must select the candidate preferred by the majority, which is c_2 , leading to a majority agreement of 0. This pathology arises because the majority of voters are in the minority on the majority of issues. Thus, if we have any number of duplicates of these two candidates for any k , there are cases in which the voters will elect representatives who achieve a majority agreement of exactly 0, even when a set of candidates exists who would achieve an agreement of 1. This concludes our proof of Theorem 3. \square

A smaller example is given below.

While Condorcet-consistent rules may not be able to provide us with any bounded worst-case approximation of majority agreement, not all election rules are Condorcet-consistent. Furthermore, we may also care about the expected agreement rather than the worst case. We proceed with an empirical investigation by simulation of the expected majority agreement for a number of election rules.

Table 2 Example of the Ostrogorski paradox based on Anscombe's paradox

| | s^1 | s^2 | s^3 | s^4 | s^5 | s^6 | s^7 | s^8 | s^9 | s^{10} | s^{11} | P_{VC} |
|----------|-------|-------|-------|-------|-------|-------|-------|-------|-------|----------|----------|----------|
| v_1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | c_2 |
| v_2 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | c_2 |
| v_3 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | c_2 |
| v_4 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | c_2 |
| v_5 | 1 | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | c_2 |
| v_6 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | c_2 |
| v_7 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | c_2 |
| v_8 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | c_1 |
| v_9 | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | c_1 |
| v_{10} | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 0 | 1 | c_1 |
| v_{11} | 0 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | c_1 |
| c_1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | |
| c_2 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | |
| maj | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | c_2 |

Table 3 Simplified example based on Anscombe's paradox

| | s^1 | s^2 | s^3 | P_{VC} |
|-------|-------|-------|-------|----------|
| v_1 | 0 | 0 | 1 | c_2 |
| v_2 | 0 | 1 | 0 | c_2 |
| v_3 | 1 | 0 | 0 | c_2 |
| v_4 | 1 | 1 | 1 | c_1 |
| v_5 | 1 | 1 | 1 | c_1 |
| c_1 | 1 | 1 | 1 | |
| c_2 | 0 | 0 | 0 | |
| maj | 1 | 1 | 1 | c_1 |

5 Simulated elections

We now investigate majority agreement as a function of the number of voters, issues, candidates, and representatives, as well as the choice of election rule. For all issues, the preference of every voter and candidate is 1 with probability $\frac{1}{2}$. In other words, the preferences of every agent on every issue are determined by the flip of a fair coin. Thus, if the representatives were selected from the candidates uniformly at random, the expected majority agreement would be 0.5. This assumption isolates the effects of the election, since any observed majority agreement beyond random chance is due entirely to the election process. We will relax the assumption of uniform random preferences later in Sect. 7.

Our assumption of uniformly random voter preferences also implies that the size of the voter majority relative to the voter minority on an issue tends to be small. Let V_1 be the voters who vote for 1 on an issue and V_0 be those who vote for 0. If we

think of the voters as voting in some arbitrary sequence then the value of $|V_1| - |V_0|$ is modeled by an unbiased random walk, and so $\mathbb{E}[(|V_1| - |V_0|)^2] = |V|$.

As before, we assume that voters' preferences over candidates are induced by their agreement on the issues. Our assumption that the electoral preferences are derived without uncertainty is a best-case assumption for the election rules. In addition to approvals and rankings, we add the option for voters to weight the candidates using normalized weights. These weights provide strictly greater information than rankings. With weights, v_j assigns each candidate c_l a weight of $w_{jl} = \frac{L(v_j, c_l)}{\sum_{c' \in C} L(v_j, c')}$. The weight attributed to each candidate is the sum of weights assigned to them, and the k candidates with the greatest total weight are elected, with ties broken lexicographically. Note that one can infer the approvals or rankings of voters from these weights, up to tie-breaking, but not vice versa, and so these types of cardinal preferences contain greater information than the simpler categorical and ordinal preferences.

5.1 Election rules definitions

We compare election rules computable in polynomial time under different parameterizations to see how well they approximate coverage and majority agreement. While there are many different ways to elect bodies of representatives, we take seven election rules from the literature, and compare them with electing representatives based on cardinal weights (Zwicker 2015; Aziz et al. 2014). We consider two approval-based rules (Max Approval and RAV), three ordinal rules (Borda, Purity, and IRV), a rule using cardinal weights (Max Agreement), and ignoring any electoral preferences by selecting the representatives uniformly at random (Random Winners, i.e., sortition Ebadian et al. (2022)).

Let $C_j \subseteq C$ denote the subset of candidates which voter v_j approves when reporting approval preferences. The two approval-based rules we consider are Max Approval and Re-weighted Approval Voting (RAV), whose definitions can be found in Aziz et al. (2014).⁴

Definition 4 (Max Approval) $D = \arg \max_{C' \subset C: |C'|=k} \sum_{v_j \in V} |C_j \cap C'|$ with ties broken randomly.

Definition 5 (Re-weighted Approval Voting (RAV)) In each of k rounds we select an unelected candidate in $C \setminus D$ to add to the iteratively constructed winning set D with the highest sum of weighted approvals. In each round we weight each agent's approved candidates by $w_j = \frac{1}{1+|D \cap C_j|}$ for v_j . We update $D \leftarrow D \cup \{c^*\}$ where $c^* = \arg \max_{c \in C \setminus D} \sum_{v_j \in V: c \in C_j} w_j$ with ties broken randomly.

Let σ_j denote the strict preference ranking of voter v_j when voters report rankings over candidates, and let $\sigma_j(l) \in \{1, 2, \dots, |C|\}$ be the position at which v_j ranks

⁴ Max Approval is often referred to as "Approval Voting" in the literature (Aziz et al. 2014), but we use Max Approval to reduce confusion because the term "approval voting" can also refer to approval-based voting more broadly.

candidate c_l . The rankings are assumed to be complete such that every voter ranks every candidate. The ranking-based election rules we consider are Borda, Plurality, and Instant Runoff Voting, whose definitions can be found in Zwicker (2015). For an axiomatic characterization of IRV see Freeman et al. (2014).

Definition 6 (Borda) $D = \arg \min_{C' \subset C: |C'|=k} \sum_{c_l \in C'} \sum_{v_j \in V} \sigma_j(l)$ with ties broken randomly.

Definition 7 (Plurality) $D = \arg \max_{C' \subset C: |C'|=k} \sum_{c_l \in C'} |\{v_j \in V : \sigma_j(l) = 1\}|$ with ties broken randomly.

We refer to the summand $|\{v_j \in V : \sigma_j(l) = 1\}|$ as the plurality score of candidate c_l .

Definition 8 (Instant Runoff Voting (IRV)) In each of $|C| - k$ rounds the candidate with the lowest plurality score is removed from the candidate set, with ties broken randomly, and all agent rankings are updated to have this candidate removed. The remaining k candidates are elected as representatives.

When every agent v_j provides a cardinal value w_{jl} for each candidate c_l in their electoral preferences, we can assign a score to each candidate equal to the sum of values assigned to them and elect the k candidates with the highest scores. In our model we have assumed that the values assigned by each voter must sum to one, in accordance with one-person-one-vote, so these values reflect normalized agreements between voters and candidates. We refer to the associated rule as Max Agreement.

Definition 9 (Max Agreement) $D = \arg \max_{C' \subset C: |C'|=k} \sum_{c_l \in C'} \sum_{v_j \in V} w_{jl}$ with ties broken randomly.

5.2 Comparing election rules

Figures 1, 2, and 3 examine the numbers of issues, candidates, and representatives as independent variables, respectively. For each plot, the number of voters $|V|$ is held fixed at 501 as we did not observe significant dependence of majority agreement on the number of voters.

For each of the following experiments, the mean agreement and coverage were averaged over 10,000 instances with uniformly random issue preferences. Turning first to Fig. 1 we hold $|C| = 60$, $|k| = 21$ and vary $|S| \in \{15, \dots, 150\}$ in increments of 15. We see that for a small number of issues Max Approval, RAV, Borda, and the weighted election rule (Max Agreement) select a committee in agreement with the majority 70–80% of the time. However, as we add issues to the docket, the voting rules seem to converge around 60%. In Fig. 2 we fix $|k| = 21$, $|S| = 150$ and vary the number of candidates between $|C| \in \{21, \dots, 101\}$ in increments of 5. We observe again that Max Approval, RAV, Borda, and Max Agreement are the best followed closely by IRV. As we increase the number of candidates it is possible for the system to more frequently recover the will of the majority but this number does not climb




Fig. 1 Majority agreement varying number of issues with $k = 21$, $|C| = 60$, $|V| = 501$




Fig. 2 Majority agreement varying number of candidates with $k = 21$, $|S| = 150$, $|V| = 501$




Fig. 3 Majority agreement varying number of representatives with $|C| = 100$, $|S| = 150$, $|V| = 501$

above 65% across all treatments. Finally, in Fig. 3 we hold $|C| = 100$, $|S| = 150$ and vary $|k| \in \{1, 5, 11, 15, \dots, 91, 95\}$.

While both Max k -Coverage and Max k -Majority Agreement are NP-Hard, majority agreement is a stronger condition as a majority agreement of $x \in [0, 1]$ implies coverage of at least x . From the experiments in Figs. 3, 2, and 1 we get a sense of just how much stronger of a condition majority agreement is. In all of our runs, coverage was 1.0 for all election rules and all combinations of parameters. With FRD, issue-specific weights enable a potential increase in agreement from the rate of majority agreement (with equally weighted representatives) up to the rate of coverage. As long as an issue is covered, there exists a weighting such that the weighted majority of representatives agrees with the voter majority, even if the majority of representatives does not. Intuitively, if we take a set of representatives D and add one or more representatives, coverage can only increase or stay the same, but agreement can potentially decrease. As our experiments reveal, the difference between majority agreement and coverage can be highly significant for many common voting rules. Delegative voting therefore has great potential to increase agreement. However, the degree to which this benefit is realized depends upon the delegation behavior of the voters.

6 Adding flexibility

Weighting representatives can increase the agreement between their vote outcomes and the voter majority as long as issues are covered. As we saw in the previous section, a

coverage rate of 1.00 was achieved in all of our experiments. To take advantage of a coverage rate greater than the rate of majority agreement, we look at different ways of weighting the representatives. Along these lines, we extend the notion of majority agreement to include weighted majority voting.

Definition 10 (*Weighted Majority Agreement*) The weighted majority agreement between a set of voters V and a set of representatives D with weighting W is the fraction of issues on which the weighted majority of representatives agrees with the voter majority, i.e., $L(R(P_V), R(P_D, W))$.

The most basic way to weight the representatives would be based on the agents' preferences over the candidates during the election. Unfortunately, as we will show in Sect. 6.3, weighting the representatives based on the election profile (P_{VC}) offers little hope under our preference model. In hopes of recovering the "will of the voters" we build into our FRD mechanism the ability for voters to weight the representatives on an issue-by-issue basis, if they choose.


We look first at basic features of FRD in a deterministic setting before considering probabilistic participation by the voters. In our FRD mechanism, each voter gets a single, divisible voting unit, or token, on every issue. This token is distributed to representatives by either default or delegation. Through delegation, the voters can divide the token among the representatives as they please, with only the constraint that the units they assign across the representatives must sum to unity. The weight of a representative is the sum (or average) of tokens given to them by the voters. The distribution of a voter's token among the representatives is referred to as *delegation*, in line with the literature on "delegative voting." Of course, if on every issue every voter delegates their full token to a single representative who agrees with them, then a weighted majority vote among the representatives exactly emulates Direct Democracy. But what happens if a voter does not use their option to delegate? In standard proxy voting, if a voter does not choose a proxy, then they abstain. The distinguishing feature of our FRD mechanism is that it permits a default assignment of tokens from voters to representatives from which the voters may deviate. The default distribution can be the uniform distribution, or it can be based on election preferences. Since election preferences do not appear to offer any benefit under our preference model, we will focus on the uniform distribution as the default when issue-specific delegations are allowed. Under the uniform default distribution, the default evenly splits every voter's token among the representatives on every issue. If no agents shift their vote tokens on any issue, we have exactly a RD. Here, FRD interpolates directly between RD and Direct Democracy at the will of the voters.

6.1 Deterministic delegation

Consider a single issue s^i . Let $V_1 = \{v_j \in V : v_j^i = 1\}$ be the set of voters in the voter majority, and $V_0 = V \setminus V_1$ be the minority. Similarly, let $D_1 = \{d_l \in D : d_l^i = 1\}$, and $D_0 = D \setminus D_1$.

We denote by λ_1 the number of voters in V_1 who use their option to delegate, and by λ_0 the delegators in the minority. We say that a voter's delegation is *incisive* if the

Fig. 4 Example of delegation increasing agreement (right) and decreasing agreement (left)



voter delegates their entire voting unit to representatives who agree with them on that issue. For example a voter may delegate entirely to a single representative who agrees with them on an issue.

If the default is abstention (i.e., no weight assigned to reps) and delegations are incisive, then agreement is guaranteed if the number of delegators in the majority is greater than the number of delegators in the minority, just like traditional proxy voting. The minority can push the outcome in their favor by delegating at a higher rate than the majority.

The first question we ask is, with a uniform default weighting, how many representatives must agree with the voter majority for the agreement to be guaranteed. That is, if all minority voters delegate incisively and none of the majority voters delegate, when does the uniform default guarantee agreement?

The weighting of the representatives depends on whether the representatives also have voting units to delegate to themselves. For brevity, we will assume the representatives do not, so if a representative receives no voting units from the voters, they have a weight of zero. For consistency, we assume the sets of voters and representatives are disjoint, just as we assume the sets of voters and candidates are disjoint in our simulations.

Example 1 Consider two FRD instances with two issues $S = \{s^1, s^2\}$, three voters $V = \{v_1, v_2, v_3\}$, and three representatives $D = \{d_1, d_2, d_3\}$ as shown in Fig. 4. The solid arrows from voters to representatives indicate incisive delegations, and any voter without an arrow stays with the default uniform distribution on that issue. The voter and representative preferences are given in the tables above and below the agents. We denote by X_l^i the total voting units assigned to representatives who agree with the voter majority on issue s^i and by W_l^i the weight of representative d_l on s^i .

On issue s^1 , the representative majority agrees with the voter majority, so majority voting would yield agreement. However, since only the voter in the minority delegates (v_3^1), the weighted majority of representatives now decides the outcome in favor of the voter minority ($X_1^1 < \frac{|V|}{2}$). This reversal can occur if the number of voters in the minority is large enough, the number of representatives who agree with the voter minority is large enough, and the voters in the minority delegate at a substantially higher rate than the voters in the majority. We formalize this intuition in Proposition 1.

On issue s^2 the representative majority disagrees with the voter majority, so the majority voting outcome (without delegations) would be, regrettably, 0. Looking again at the Fig. 4 we see the delegations flip the result to what would be achieved by Direct Democracy; namely, from disagreement to agreement ($X_1^2 > \frac{|V|}{2}$). Hence, FRD can improve the outcomes over RD as measured by agreement with Direct Democracy. Fortunately, for both s^1 or s^2 , if any two, or all three, of the voters delegate incisively, the outcome will always agree with the voter majority. We formalize this intuition in Proposition 2. This concludes Example 1.

On issue s^1 on the left side of Fig. 4, delegation by the minority flips the outcome away from agreement with the voter majority. This flip happens because the voters in the majority do not delegate, and the default distribution of weight from the voter majority is not sufficient to prevent the inversion, even though the majority of representatives agrees with the voter majority. However, even when no voters in the majority delegate, it may be impossible for the voter minority to sway the outcome in their direction because of the size of the voter majority ($|V_1|$) and the number of representatives who agree with them ($|D_1|$). This contrasts with traditional proxy voting with no default distribution, in which the outcome is determined entirely by the votes of the representatives and those who delegate. In Proposition 1 we give a necessary and sufficient condition on the size of V_1 and D_1 relative to V and D to ensure agreement even when none of the voters in the majority delegate.

Proposition 1 *If the representatives are weighted uniformly by default, all voters in the minority delegate incisively ($\lambda_0 = |V_0|$), and none of the majority voters delegate ($\lambda_1 = 0$), agreement is guaranteed when the number of representatives ($D_1 \subseteq D$) who agree with the voter majority ($V_1 \subseteq V$) is greater than $\frac{|D|}{2} \cdot \frac{|V|}{|V_1|}$.*

Proof We want to know under what conditions $X_1 > X_0$. With a uniform default and only incisive delegations by the minority, the total amount of weight on the representatives on either side of the issue are $X_1 = |V_1| \cdot \frac{|D_1|}{|D|}$ and $X_0 = |V_1| \cdot \frac{|D_0|}{|D|} + |V_0|$.

$$\begin{aligned} X_1 &= |V_1| \cdot \frac{|D_1|}{|D|} > |V_1| \cdot \frac{|D_0|}{|D|} + |V_0| = X_0 \\ |V_1| \cdot \frac{|D_1| - |D_0|}{|D|} &> |V| - |V_1| \\ |V_1| \cdot \left(\frac{|D_1| - |D_0|}{|D|} + 1 \right) &> |V| \\ |V_1| \cdot \left(\frac{|D_1| - |D_0| + |D|}{|D|} \right) &> |V| \end{aligned}$$

$$|V_1| \cdot \frac{2|D_1|}{|D|} > |V|$$

$$|D_1| > \frac{|D|}{2} \cdot \frac{|V|}{|V_1|}$$

This concludes our proof of Proposition 1. \square

Returning to issue s^1 on the left side of Fig. 4, we can see that $|V| = |D| = 3$ and $|D_1| = |V_1| = 2$, and $2 \not> \frac{3}{2} \cdot \frac{3}{2}$, so $X_0 > X_1$ and the minority prevails in accordance with Proposition 1.

Proposition 1 gives conditions on the size of the majority under which no amount of delegation by the minority can yield disagreement, even when the majority does not delegate at all. However, if the minority is large enough, then their delegations induce disagreement. If $|D_1| > |D_0|$, then the minority voters must delegate at a sufficiently higher rate than the majority voters for the weighted majority to no longer agree with the voter majority. If $|D_0| > |D_1|$, the majority delegation rate must be sufficiently higher than that of the minority to guarantee agreement. We make this intuition more precise in Proposition 2, which gives a condition on the delegation rate of the majority to ensure agreement.

Proposition 2 *If the representatives are weighted uniformly by default, the weighted majority of representatives agrees with the voter majority after delegation whenever $\lambda_1 > \frac{2\lambda_0|D_1| - |V|(|D_1| - |D_0|)}{2|D_0|}$.*

Proof Recall that the total voting units assigned to representatives in D_1 and D_0 are X_1 and X_0 , respectively. For the representatives to agree with the voter majority it must be that:

$$\begin{aligned} X_1 &> X_0 \\ \lambda_1 + \frac{|D_1|}{|D|}(|V_1| - \lambda_1) + \frac{|D_1|}{|D|}(|V_0| - \lambda_0) &> \lambda_0 + \frac{|D_0|}{|D|}(|V_1| - \lambda_1) + \frac{|D_0|}{|D|}(|V_0| - \lambda_0) \\ \lambda_1 + \frac{|D_1| - |D_0|}{|D|}(|V_1| - \lambda_1) &> \lambda_0 + \frac{|D_0| - |D_1|}{|D|}(|V_0| - \lambda_0) \\ \lambda_1|D| + (|D_1| - |D_0|)(|V_1| - \lambda_1) &> \lambda_0|D| - (|D_1| - |D_0|)(|V_0| - \lambda_0) \\ 2\lambda_1(|D_0|) + |V_1|(|D_1| - |D_0|) &> 2\lambda_0|D_1| - |V_0|(|D_1| - |D_0|) \\ 2\lambda_1(|D_0|) &> 2\lambda_0|D_1| - (|V_1| + |V_0|)(|D_1| - |D_0|) \\ \lambda_1 &> \frac{2\lambda_0|D_1| - (|V_1| + |V_0|)(|D_1| - |D_0|)}{2|D_0|} \end{aligned}$$

This concludes our proof of Proposition 2. \square

Proposition 2 gives a necessary condition for agreement between the voters and representatives using the weighted majority decision rule. The condition is in the form of a lower bound on the delegation rate by the voters in the majority as a function of (1) the delegation rate of the minority, (2) the difference in size between voter minority and voter majority, and (3) the sizes of the subsets of representatives who agree and disagree with the voter majority.

Recall the second issue s^2 from Fig. 4. In this example, $|V_1| = 2$, $|V_0| = 1$, $|D_1| = 1$, $|D_0| = 2$, $\lambda_1 = 1$, and $\lambda_0 = 0$. We can see that $1 > \frac{0-(2+1)\cdot(1-2)}{2\cdot2} = \frac{3}{4}$, and therefore agreement is achieved after delegation even though $|D_0| > |D_1|$.

6.2 Probabilistic delegation

We now investigate what happens if each voter chooses to delegate with some fixed individual probability. These results give us insight into how motivated or attentive voters must be to improve the outcome of FRD over RD.

As all issues are independent, we will consider a single issue. Suppose each voter $v_j \in V$ chooses to delegate incisively with independent probability p_j and leaves their voting unit distributed uniformly with probability $1 - p_j$. Let $x_j \in [0, 1]$ be the total voting units v_j assigns to candidates in D_1 . If v_j does not delegate then $x_j = \frac{|D_1|}{|D|}$, if v_j delegates incisively and is in the voter majority then $x_j = 1$, and if v_j delegates incisively and is in the voter minority then $x_j = 0$. Let $X_1 = \sum_{v_j \in V} x_j$ be the total voting units assigned to D_1 via both delegation and the uniform default. Let $\mu = E[X_1] = \sum_{v_j \in V_1} \left(p_j + (1 - p_j) \frac{|D_1|}{|D|} \right) + \sum_{v_j \in V_0} (1 - p_j) \frac{|D_1|}{|D|}$ be the expected value of the total voting units assigned to representatives who agree with the voter majority.

Theorem 4 *Suppose there is an odd number of voters $|V|$, odd number of representatives $k = |D|$, a uniform default, no abstentions, and only incisive delegations. Suppose further that each voter $v_j \in V$ delegates with probability p_j on an issue such that $\mu > |V|/2$. Then the probability that the outcome agrees with the voter majority is bounded by $P(\text{weighted majority agreement}) \geq 1 - e^{-(|V|-2\mu)^2/4|V|}$.*

Proof Sketch. Let $y \in \{0, 1\}$ denote the outcome of the weighted majority vote by the representatives. The probability that the outcome agrees with the voter majority is $P(y = 1) = P(X_1 > |V|/2) + P(y = 1|X_1 = |V|/2) \cdot P(X_1 = |V|/2)$ where $P(y = 1|X_1 = |V|/2)$ is due to the tie-breaking mechanism. First we show that with odd voters, odd representatives and only incisive delegations, there can be no ties. Namely, $X_1 \neq X_0 = |V| - X_1$. This proof is due to parity and holds regardless of the delegation rate. Without ties, we simply need to determine $P(X_1 > |V|/2)$. We use a Chernoff inequality to provide a lower bound on this value based on the delegation probabilities p_j of all voters.

Proof Recall that $x_j \in [0, 1]$ is the amount of voting units (fraction of a token) voter v_j assigns to representatives who agree with the voter majority on an issue and $X_1 = \sum_{v_j \in V} x_j$. Given some tie-breaking rule, we have that $P(y = 1) = P(X_1 > |V|/2) + P(y = 1|X_1 = |V|/2) \cdot P(X_1 = |V|/2)$. First we show that $P(X_1 = |V|/2) = 0$, then we give a lower bound for $P(X_1 > |V|/2)$.

Lemma 1 *If $|V|$ is odd, $|D|$ is odd, and all delegations are incisive, then no ties are possible.*

Proof Let $x'_j = |D| \cdot x_j$ where $x_j \in [0, 1]$ is the voting units (fraction of a token) voter v_j assigns to candidates who agree with the voter majority on an issue via default or delegation. If v_j does not delegate then $x'_j = |D_1|$, if v_j delegates incisively and is in the voter majority then $x'_j = |D|$, and if v_j delegates incisively and is in the voter minority then $x'_j = 0$. Therefore, $\forall j : x'_j \in \{0, |D_1|, |D|\}$. Let $X'_1 = \sum_{v_j \in V} x'_j$ and $X'_0 = \sum_{v_j \in V} (|D| - x'_j)$. Then X'_1, X'_0 are non-negative integers and $X'_1 + X'_0 = |D| \cdot |V|$.

Since $|D| \cdot |V|$ is odd, it must be that X'_1 and X'_0 have opposite parity and so they cannot be equal. Therefore $X_1 = \frac{X'_1}{|D|} \neq X_0 = \frac{X'_0}{|D|}$, meaning the total amounts of weight delegated to the representatives on either side of the issue cannot be equal, so no ties may occur. This concludes our proof of Lemma 1. \square

Given that no ties are possible, we have that $P(y = 1) = P(X_1 > |V|/2)$. Remember that $X_1 = \sum_{v_j \in V} x_j$ where x_j is the total weight that v_j delegates to representatives who agree with the voter majority. If $v_j = 1$ then $E[x_j] = p_j + (1 - p_j) \frac{|D_1|}{|D|}$, else if $v_j = 0$ then $E[x_j] = (1 - p_j) \frac{|D_1|}{|D|}$. Let $\mu = E[X_1]$ be the expected total weight assigned to representatives who agree with the voter majority.

$$\mu = \sum_{v_j \in V_1} \left(p_j + (1 - p_j) \frac{|D_1|}{|D|} \right) + \sum_{v_j \in V_0} (1 - p_j) \frac{|D_1|}{|D|}$$

We now use the fact that $P(X_1 > |V|/2) = 1 - P(X_1 \leq |V|/2)$. Let $\delta = (2\mu - |V|)/2\mu$. If $\mu > |V|/2$, then $\delta > 0$. This allows us to apply a Chernoff bound to derive our lower bound:

$$\begin{aligned} P(X_1 > |V|/2) &= 1 - P(X_1 \leq |V|/2) \\ &= 1 - P(X_1 \leq (1 - \delta)\mu) \\ &\geq 1 - e^{-\delta^2 \mu^2 / |V|} \\ &= 1 - e^{-(2\mu - |V|)^2 / 4|V|} \end{aligned}$$

This concludes our proof sketch of Theorem 4. \square

This bound depends on the condition that $\mu > |V|/2$. This condition requires that the delegation rate of the majority cannot be too small relative to the minority. Observe that this condition is satisfied when $\forall j : p_j = p$ and $|D_1| > |D|/2$. Furthermore, as an increasing number of voters delegate incisively, we expect $\mu \rightarrow |V_1| > |V|/2$ regardless of $|D_1|$. Naturally, as the delegation rate increases, we observe our lower bound approach the ideal $1 - e^{-(2\mu - |V|)^2 / 4|V|} \rightarrow 1$. To find $\mu \leq |V|/2$, the voter majority cannot be too large compared to the voter minority, $|D_1|$ must be smaller than or somewhat close to $|D_0|$, and/or the voters in the majority must be significantly more apathetic towards delegation than voters in the minority. In other words, μ encodes the relative proportions used in Propositions 1 and 2, probabilistically.




Fig. 5 Comparing majority agreement as a function of the delegation rate for different delegation behaviors in FRD with uniform default, Max Agreement election rule, and weighted majority decision rule

6.3 Simulated delegation

Finally, we provide an empirical comparison between RD and FRD with different issue-specific delegation behaviors by the voters. We simulate instances where all voters delegate with equal probability. We refer to the total fraction of voters who delegate as the *delegation rate* (α). Whereas the p_j denoted a voter's probability of delegating in the previous section, α is the realized fraction of voters who have delegated.

We use the same model to generate candidate and voter preferences as used in Sect. 5.2. For our simulated delegations we create instances with $|V| = 301$, $|C| = 60$, $|S| = 150$, and $k = 21$. We have lowered the number of voters compared to previous experiments (from 501 to 301) because we have found this makes no qualitative difference to the results and allows us to run a larger number of instances. In Fig. 5 representatives are elected using the Max-Agreement election rule. We vary $\alpha \in [0, 1.0]$ in increments of 0.01 and for each setting of α we run 50 iterations. We plot the means of weighted majority agreement in Fig. 5. A value of 1.0 means that the outcomes of all issues are the same as they would be in a Direct Democracy with full participation.

We compare RD against four different delegation schemes: (1) *Approve* where voters delegate evenly to the representatives whom they approved in the election across all issues; (2) *Best Rep* where voters delegate to their single most preferred representative (i.e., proxy) for all issues; (3) *Best-3 Rep* where voters delegate equally to their three most preferred representatives for all issues; and finally (4) *Incisive* where voters delegate incisively to representatives with whom they agree on each individual issue. In all cases we use a uniform default weighting.

Most surprising is that there is no observable benefit to any of the non-incisive delegation behaviors (Approve, Best Rep, and Best-3 Reps). When $\alpha = 1$, each of these

correspond to an electoral weighting scheme that could be used as a default in FRD. The *Approve* weighting is perhaps closest to the proposal of Proxy Voting espoused by Miller (1969) but does not improve weighted majority agreement in a meaningful way. Similarly interesting are the *Best Rep* and *Best-3 Rep* which correspond to proxy voting with 1 and 3 proxies, respectively. These forms of proxy voting do not move the outcome towards the ideal of Direct Democracy. Hence, we can see that issue-specific flexibility can be effectively used to improve agreement for voting systems in a way that other weighting schemes do not. Recall that the difference between FRD with incisive delegations and proxy voting with issue-specific proxies is the use of the default distribution. This is why the curve for FRD with little to no delegation ($\alpha \approx 0$) achieves an agreement equal to RD.

A notable feature of Fig. 5 is how much FRD improves agreement over RD when voters are highly attentive and how little a difference it makes when voters are not attentive. When the delegation rate is below 20% there is minimal benefit, and at 10% there is no observable benefit. In the range of 30–40% we start to see an increase in agreement. With about 60% of the population delegating we can improve agreement by 10% to a level not reached by any of the unweighted election rules in Figs. 1, 2, or 3 with 150 issues. When the delegation rate reaches 80% we see an almost 20% increase; eventually reaching 100% agreement when everyone delegates (because the issues are covered).

To verify that the lack of improvement from Best- x delegation is a general phenomenon and not a consequence of using the Max Agreement election rule or small values of x , we repeat the experiment for Best- x reps varying x for all of our election rules in Fig. 6. We use $|V| = 100$, $|C| = 60$, and $k = 21$, varying x from 0 to 11. For each experiment the mean agreement was averaged over 1000 instances. Observe that the agreement line is flat for all election rules.

7 Issue preference distributions

In our empirical investigations so far we have assumed that every agent's preference on every issue is decided by an unbiased coin flip. In other words, each issue preference has been drawn i.i.d. from a Bernoulli distribution with parameter $p = 0.5$. In Sect. 5.2, we compared the mean agreement achieved under RD with various election rules using this assumption. We saw that certain election rules lead to consistently higher mean agreement than others under various values of the numbers of issues, candidates, and representatives. We now investigate the degree to which our prior observations depend on the uniform randomness of issue preferences by varying the preference distributions.

7.1 Varying candidate preferences

We now show two experiments varying the preferences of the candidates, using $|V| = 100$, $|C| = 60$, $|S| = 200$, and $k = 21$. In Fig. 7 we draw all voter preferences from a Bernoulli distribution with $p = 0.5$, and show how agreement changes when




Fig. 6 Comparing majority agreement for various election rules with Best- k delegation in FRD with uniform default and weighted majority decision rule

we vary the candidates' Bernoulli parameter. What we observe is that agreement is highest when the candidates' also have $p = 0.5$, and this holds for all of the election rules we consider. In Fig. 8 we repeat the same experiment, but now the voters have Bernoulli parameter $p = 0.6$. The symmetry in the previous experiment disappears. Now agreement is an increasing function of the candidates' Bernoulli parameter for all rules. Again, we see a consistency of behavior across all of our election rules. This observations suggests that one may be able to concretely measure what makes for a "good" set of candidates irrespective of the election rule.

At first glance, one might expect a set of candidates to be "good" if the distribution of preferences of the candidates is similar to the preference distribution of the voters. When measuring by agreement, we can see this is not the case. What makes for high agreement is not that the candidates' preference distribution is most similar to the voters', but that the candidates' preference distribution is most similar to that of the voter majority. When $p > 0.5$ for the voters, the Condorcet Jury Theorem gives us that the probability of the voter majority voting for 1 (before relabeling) tends to unity as the number of voters increases (Austen-Smith and Banks 1996). Thus, it is better to have candidates whose Bernoulli parameter is as close as possible to 1, rather than being equal to the voters' Bernoulli parameter.




Fig. 7 Varying candidates' Bernoulli parameter when voters' Bernoulli parameter is $p = 0.5$




Fig. 8 Varying candidates' Bernoulli parameter when voters' Bernoulli parameter is $p = 0.6$




Fig. 9 Mean agreement varying number of issues in *similar* preference regime

7.2 Similarity and polarization

We now hold fixed the Bernoulli parameters for voters and candidates and vary the other parameters to repeat the experiments in Sect. 5.2. In the *similar* regime (Figs. 9, 10, 11), all agents have their issue preferences drawn i.i.d. from a Bernoulli distribution with parameter $p = 0.6$. In the *polarized* regime (Figs. 12, 13, 14), the voters' Bernoulli parameter is $p = 0.6$ while the candidates' have Bernoulli parameter $p = 0.4$. This gives us a sense of how the performance of RD depends on the quality of the candidate set.

What we see is that the order of the performance of the election rules is essentially preserved across regimes, but the mean agreement achieved by each rule is highly sensitive to which preference regime we are in. In the similar regime, when all agents' preferences are drawn from the same Bernoulli distribution, high agreement can be achieved with any election rule, even random selection. By contrast, in the polarized regime, all election rules we consider struggle to achieve agreement and perform worse than choosing issue outcomes at random. This leads us to conclude that in our model the choice of election rule is much less important than the quality of the candidate set.

In all experiments in the similar and polarized regimes we use $|V| = 100$, $|C| = 60$, $|S| = 200$, and $k = 21$ for each variable when it is not varied as the independent variable. In Figs. 9 and 12 we vary the number of issues $|S|$ from 20 to 200 in increments of 20, in Figs. 10 and 13 we vary the number of candidates $|C|$ from 30 to 100 in increments of 10, and in Figs. 11 and 14 we vary the number of representatives k from




Fig. 10 Mean agreement varying number of candidates in *similar* preference regime




Fig. 11 Mean agreement varying number of representatives in *similar* preference regime




Fig. 12 Mean agreement varying number of issues in the *polarized* preference regime




Fig. 13 Mean agreement varying number of candidates in the *polarized* preference regime




Fig. 14 Mean agreement varying number of representatives in the *polarized* preference regime

10 to 60 in increments of 10. For each experiment the mean agreement was averaged over 10,000 instances.

8 Preference intensity

So far, in our investigations of FRD, all agents have their preferences over issues drawn according to the same Bernoulli distribution and the subset of voters who choose to delegate has been determined randomly. We now consider a setting in which each voter v_j is assigned a preference intensity parameter $p_j \in [0, 1]$. Each voter's issue preferences are drawn from the Bernoulli distribution with parameter p_j and their probability of becoming a delegator is $2|p_j - 0.5|$. Voters with preference intensity close to 0 or 1 will delegate almost surely, while a preference intensity close to 0.5 means a voter is unlikely to delegate. In our simulation we draw each voter's preference intensity from the uniform distribution over $[0, 1]$, with $|V| = 100$, $|C| = 60$, and $|D| = 200$. For each experiment the mean agreement was averaged over 10,000 instances. We continue to assume that all candidates' have the same Bernoulli parameter determining their issue preferences and investigate how this parameter affects mean agreement.

The agreement curves are each symmetric about 0.5, looking qualitatively similar to Fig. 7, except that the rates of agreement are higher for every election rule. We can see that intensity-based delegation behavior leads to higher agreement than RD, but the difference is at most 10 percentage points. Remarkably, the ordering of the




Fig. 15 FRD with intensity-based voter preferences and delegation behavior, varying the Bernoulli parameter for candidates' issue preferences

election rules is consistent with our other experiments, where Max Agreement and Borda come first, followed by Max Approval and RAV, trailed by IRV and Plurality which only beat random selection.

9 Summary of analysis

We began by showing that even when the voter majority is known for all issues, electing a set of representatives to maximize majority agreement is NP-Hard. Therefore, any polynomial-time election rule can only hope to approximate agreement when the numbers of candidates and issues are large. One potential way to increase agreement is to weight representatives, but as our experiments show, weighting the representatives with a constant weight across all issues based on the election profile does not offer much hope in terms of increasing agreement in our model. More work is needed to determine whether a default weighting based on the electoral profile can yield benefits.

While maximizing coverage is also NP-Hard, coverage is a weaker requirement than majority agreement because majority agreement implies coverage but not vice versa. It follows that the rate of coverage is always greater than or equal to the rate of majority agreement. The difference between the rate of coverage and the rate of agreement is the measure of potential benefit from issue-specific delegation. Whether this potential benefit is realized depends on voters' delegation behavior.

In our simulations of FRD, we see that expected agreement increases monotonically as a function of the delegation rate with uniform default. With uniformly random issue

preferences the voter majority tends not to be much larger than the voter minority, leading to a slower increase in agreement with the delegation rate. On the other hand, the rate of increase in expected agreement is optimistic due to the assumption of incisive delegation.

Delegation is not guaranteed to increase agreement. Even if there is majority agreement on an issue, if the minority is more active in delegating than the majority, they can flip the outcome away from the voter majority. Whether a minority can turn the outcome of a decision in their direction depends on (1) the size of the minority, (2) how active the minority is in delegating relative to the majority, (3) whether the representative majority agrees with the voter majority, and (4) the size of the representative majority. We give specific bounds for when this can occur assuming a uniform default and incisive delegation.

One of the consistent features throughout our simulated results is that the order of performance for election rules is consistent across parameterizations and preference regimes. Borda and Max Agreement consistently achieve the highest majority agreement. The next two best rules are Max Approval and RAV, which only require approval votes. IRV and Plurality consistently have the worst performance, except for selecting representatives randomly. However, the performance of RD therefore depends most strongly on the quality of the candidate set. Majority agreement is much more sensitive to the agents' preference distributions than it is to the choice of election rule under our model. Since there was no significant difference in coverage achieved by these rules in our simulations, the choice of election rule becomes of secondary concern once delegations are permitted.

9.1 Limitations of analysis

Our analysis makes use of a toy model of agent preferences and behavior that may not accurately reflect preferences and behavior in the real-world. In reality, a single agent's preferences on different issues may be correlated or dependent, different agents' preferences may be correlated with one another, agent behavior may not be a simple deterministic function of their issue preferences, and/or many other factors. Here we have assumed all agent preferences over issues are drawn uniformly at random from the space of possible profiles, and this assumption of impartial culture may also be unrealistic Egecioğlu and Giritligil (2013); Abramson et al. (2022); Mattei and Walsh (2013, 2017). However, our simplistic model is sufficient to show (1) the worst-case complexity of maximizing majority agreement, and (2) that optimal coverage will tend to be much easier to approximate than majority agreement. These are the core observations that characterize the potential benefit of issue-specific delegations in our model, and we expect both of these phenomena to generalize to different, more sophisticated, preference models. We further note that correlation across voters and across issues can potentially reduce average-case complexity, which we do not analyze. The simplicity of our model means that the expected rates of agreement and coverage seen in our simulations should not be used as an estimate for real-world instances. Despite this, the fact that the ordering among election rules was consistent across different parameterizations and preference models suggests a direction for further work within

the literature on selecting multi-winner election rules (Faliszewski et al. 2017). Relaxing these assumptions and considering different preference distributions may change the performance of all voting systems we consider, and could potentially increase or decrease the performance benefit of FRD over RD.

In practice, agents do not always participate in every aspect of the decision-making process available to them, e.g., they may abstain. In Direct Democracy, individual voters may choose not to cast votes on some or all issues. In RD, voters may not participate in the election. And in proxy voting, voters may not select a proxy. In our simulations, we only considered incomplete participation in the delegation process. Our assumption of complete participation in the election process attempts to give the best possible conditions for RD, but should be relaxed in future work.

One of our simplifying assumptions was that the sets of voters and candidates are disjoint, and we subsequently seem to implicitly ignore the votes of failed candidates after the election. However, in our theoretical work this is largely a difference of notation that simplifies our presentation and in our empirical work it makes a negligible difference. For our complexity results on elections, the assumption of disjoint voters and candidates plays no role. The exact same results are achieved if we assume every candidate is duplicated in the voter set. Similarly, for the deterministic bounds in Propositions 1 and 2, which apply only after the election, if every representative is duplicated in the voter set, the same results hold. One simply needs to add the assumption that representatives always “delegate” incisively to themselves and cannot do otherwise, but the bounds are unchanged. For the probabilistic bounds in Theorem 4, this assumption also plays no role, because one can assume all representatives are also voters who delegate to themselves with probability 1. As for our empirical results, in large populations where the probability of a small number of voters being pivotal becomes vanishingly small, our assumptions do not qualitatively change the results in our simulated experiments when $|D| \ll |V|$.

Our analysis is limited to independent, symmetric binary issues with an explicitly majoritarian objective. Many real-world applications include asymmetric binary issues in which there is a distinguished status quo (e.g., legislative bills) (Shapiro and Talmon 2018) or more than two alternatives. In the real world there can be dependencies between issues, the order that issues are decided may matter, and other objectives may be prioritized besides emulating a Direct Democracy with full participation. However, the concept of majority agreement generalizes in the obvious ways to agreement between the voters and representatives under any preference model and decision rule. Coverage also generalizes in the sense that a voter’s preference is covered if there exists a representative who expresses that preference as their vote, or a subset of representatives such that some weighting of their votes is equivalent to the voter’s preference.

10 Discussion

While Direct Democracy with full participation may be held as an ideal, costs and accessibility problems can make the use of proxies or elected representatives a more pragmatic alternative. Unfortunately, handing all decision-making power to a small

subset of the population is not without its downsides. The collective choices of representatives, and their motivations, may deviate from the “will of the voters”. In a Representative Democracy, periodic elections and the ability to recall representatives during their term are meant to keep representatives accountable. But these methods are severely limited in their ability to keep representatives accountable on specific issues or to promote civic engagement outside the election cycle. Weighting the representatives, as with proxy voting, is similarly of limited benefit when the weights are fixed across all issues between elections. Flexible Representative Democracy provides an alternative taking the best of both worlds from Direct Democracy and Representative Democracy. Ultimately, there is no perfect replacement for direct voter involvement on issues, and while this type of voter participation is not always available, that does not mean we should ignore it when it is available. We conclude by discussing some of the more detailed aspects of FRD and its different variations.

Balance of power One of the salient features of FRD is that the closer the representatives are to being tied in their votes, the more power returns to the voters. For example, in our binary model if the representatives’ votes are split so that roughly equal numbers vote for 0 and 1 on an issue, then the outcome is more sensitive to changes in delegations. Given any current weighting of the representatives, the smaller the margin of victory of the weighted majority the fewer changes by delegation are needed to change the outcome. One can see immediately how this phenomenon generalizes to other decision rules—the less decisive the representatives, the more likely delegations are to be pivotal.

The closer the representatives are to unanimous agreement, the more it would take for the voters to alter the outcome by delegation. With binary issues, this would require a large majority of the representatives to be opposed to a large majority of the voters. Presumably, if this occurs too often it will motivate more active delegations to remedy the discrepancy and impact future elections. Representatives on the losing side of any issue have incentive to promote civic engagement in the form of delegation, particularly among voters who do not already support them.

There is a symmetric phenomenon here. For example, if the voters are split on a binary issue with roughly equal delegations from both sides, the decision is predominantly determined by the representatives. Similarly, if the general population of voters are largely ambivalent about an issue and delegations are sparse, the decision-making power resides with those who were elected.

Overall, the benefit of permitting delegations between elections is that voters do not have to wait until the next election cycle to shift the balance of power if they deem the representatives to be performing poorly. The voters always have the ability to pull back power from the representatives without recalling the representatives or waiting for the next election.

Accountability and feedback One of the challenges of any system of representation is holding the representatives accountable. In a traditional Representative Democracy the primary tools for this are the election and potential for recall. Recalling elected representatives can be costly and slow and require a coordinated effort, and there can be several years between elections. To a representative who does not intend to run

for re-election, e.g., due to term limits, the threat of the next election means nothing. Delegations create a tool for empowering and dis-empowering representatives between elections without the need for a coordinated recall effort.

As Miller (1969) argued in his proposal for proxy voting, on an individual level, voters can “recall” any of their chosen proxies to whom they have assigned weight by reassigning that weight. We note that if all voters were to delegate weight away from a representative that representative would be left with no decision-making power. Under a weighted majority rule, this can happen even if the representatives’ weight is positive but too small to change the outcome given the votes of the other representatives (Banzhaf III 1964).

For any large-scale representative system in which representatives or proxies are weighted, it is generally desirable that the weights be public information. For accountability, the representatives’ weights are necessary to understand the balance of power among the representatives which is critical for assigning praise and blame (Halpern and Kleiman-Weiner 2018). In general, whether agents view a decision as acceptable or legitimate may depend on observing mechanism as much as the outcome (Abramowitz and Mattei 2022).

The literature on representation distinguishes between trustees and delegates where trustees are representatives who rely on their own independent judgment while delegates defer to the judgments of their constituents (Pitkin 2016). There is a tension between these two conceptions of representation. In FRD, the delegations of voters are a feedback mechanism that regulates the degree to which representatives act as trustees versus delegates. This is another point briefly alluded to by Miller (1969) who was concerned with the freedom of representatives to change their minds and vote their conscience.

Majoritarianism and preference intensity In our analysis we consider the majoritarian objective of emulating what the voter majority would decide on every issue. We also consider the conditions under which a voter minority can prevail over the voter majority. In some cases, this may be seen as desirable. If the effort to engage in delegation reflects preference strength, then perhaps a passionate minority should be able to overcome an apathetic majority. Where this balance is struck is determined by the election rule, and decision rule (including the weighting scheme with default and delegation). Our model of preference intensity is used to derive voters’ issue preferences as well as delegation behavior, but other preference intensity models from the literature are worth examining (Abramowitz et al. 2019; Cook and Kress 1985; Gerasimou 2021).

Issue-specific weights The FRD mechanism we investigated weights the representatives on each issue. In practice, it may be simpler to think of each representative as having a single weight that changes over time, rather than a sequence of issue-specific weights. How voters register their delegations for future issues or times is a matter of implementation.

Vacation and rolling elections While a representative with low weight might be effectively recalled, there is the option to create an automated recall that fully removes

a representative from office based on their weight. For example, if their weight is too low for too long or on too many issues, they may be automatically recalled to create the opportunity for someone new to take their place. This is simpler to implement when weights are not issue-specific, but criteria for recall can be created in either case. The removal of a representative from office may not be necessary if they can be effectively recalled by having no decision-making power, unless there is a need to replace them with a new representative.

One open question is what to do in FRD if the office of a representative is vacated, e.g., by recall or resignation. If the representative is not replaced, what should happen to the weight that was assigned to them? If a representative is replaced, what should the weight of the replacement be upon entering office? There are many potential options, and we will only briefly discuss some of the aspects of this problem.

If a representative vacates their position and is not replaced, one option is to treat their weights as still being assigned to them until the voters re-assign it to other representatives through delegation. If proposals (e.g., legislative bills) need a certain supermajority of the voting weight to pass, it needs to be determined whether this weight is counted as an abstention or treated as if it isn't there.

If a representative is replaced, the replacement could inherit the weighting of their predecessor upon entering office. If their predecessor was automatically recalled based on their lack of weight, the replacement would need at least some chance to garner delegations to prevent another automatic recall. On the other hand, if their predecessor had a very high weight due to delegation, it may be undesirable for their replacement to inherit this advantage. This inheritance makes more sense if one considers the weight to be assigned to the office rather than the individual holding the office.

A similar problem arises if all representatives do not begin their terms at the same time. What weight should new representatives be given upon entering office and how does it depend on the weights of the other representatives? New representatives could inherit the weight of their predecessors, or the weight assignment might be reset for all voters, or the distribution of weight from default might be reset only for the voters who have not used their option to delegate. Again, there are many possibilities to consider. We draw attention to the issue but do not offer a universal resolution, as the choice of procedure may be application-specific.

Public and private information When specifying an FRD mechanism it is necessary to note what information is public and what information is private. One must be particularly concerned with what information voters gain and reveal when delegating. When the FRD mechanism uses a decision rule that weights the representatives as a subroutine, for a voter to know exactly the effect their delegation will have, they must know (1) the decision rule, (2) the current weights of the representatives, (3) how their delegation changes the weights of the representatives, and (4) the representatives' votes. If delegation must stop before the representatives cast their ballots on an issue, then voters must base their delegations solely on the professed intentions of the representatives. Permitting delegation after the representatives publicly cast immutable votes allows voters who take the opportunity to delegate incisively. Unfortunately, for a voter to check the assignment of their individual voting units (or

tokens) to representatives poses a privacy concern similar to the privacy concerns for voting receipts (Sako and Kilian 1995).

There are several benefits to publicizing the weights of representatives. These weights provide information about what decisions are likely to be made and the relative support for different representatives. If multiple representatives vote similarly (e.g., come from the same political party) but have different weights, the weights can indicate factors other than preference similarity (e.g., perceived trustworthiness). Representatives have an incentive to inform the public to increase their weight to achieve their desired outcomes on issues. We note that we have not discussed the practical implementation of FRD mechanisms and the various privacy, security, and technological concerns that arise, but these are important areas of concern for future work to address.

Direct voting In our binary model, if an incisive voter delegates the entirety of their voting units (token) exclusively to representatives who agree with them on an issue, then this would be equivalent to them voting directly in a Direct Democracy with a single unit of voting weight. A voter can always delegate incisively as long as there exists at least one representative who agrees with them on the issue. The problem of electing a set of representatives such that for every issue there is at least one representative taking either position is clearly at least as difficult as achieving coverage because it implies coverage, and we prove its complexity formally in Appendix 11.1.

While coverage is easy to achieve in our toy model, this does not guarantee that it is achieved in the real world, particularly when deciding between more than two alternatives. Hence, there is an argument to be made that voters should be able to vote directly. For every possible preference someone could give as an input to the decision rule R which is not given by any representative for some issue we can construct a “dummy” representative with that preference. The dummy representative would not receive any weight from the default distribution, and only from delegation. Delegating one’s entire voting weight to the dummy can be made equivalent to voting directly under the decision rule. Voters may prefer to vote directly even when there is a representative to whom they could delegate. We view the definition of FRD as including FRD voting systems with direct voting because there is no functional difference between direct voting and delegating incisively (to a dummy representative if necessary), which is why we can denote the general function of FRD succinctly by $R(P_{VC}, P_{VD}, P_D)$ rather than $R(P_{VC}, P_{VD}, P_D, P_V)$.

There may be circumstances under which direct voting is not desirable, particularly when issues are complex. It may be beneficial for the representatives to limit the number of options to enable deliberation and proper consideration. It may also be beneficial for the representatives’ votes to act as a filter to weed out extreme or inconsistent positions. For instance, suppose the representatives are voting by submitting rankings over a set of alternatives such that if all agents are properly informed we would expect the profile of rankings to be single-peaked. Suppose also that the intended voting rule is the median rule over the single-peaked domain (Young and Levenglick 1978). If the representatives are properly informed but voters are not, voters may submit votes that cause the profile not to be single-peaked and the median to be undefined. Restricting

voters to submitting weightings over the rankings cast by representatives can ensure single-peakedness and a well-defined median.

Voluntary representatives In our model we focused on election rules that are multi-winner voting rules, selecting a fixed number of representatives at once from a common pool of candidates (Faliszewski et al. 2017). In practice representatives can be selected in many different ways, including party systems and federations, and the notion of an election rule is much broader.

One of the central features of Liquid Democracy, including the voting systems proposed by Miller (1969) and Green-Armytage (2015), is that voters can delegate to whomever they choose. Voters are not restricted to delegating to a fixed set of elected representatives and can delegate to any other voter. However, allowing delegation to any other voter enables delegation cycles, unless those who receive delegations are forbidden from delegating. Use of voluntary representatives also leads to privacy concerns as agents may try to coerce delegations from others, or delegate to others to query their vote. In FRD there is no possibility of delegation cycles because representatives cannot delegate, and representatives will not know who has delegated to them. In a large population, knowing how a representative's weight changes at regular intervals should still not be enough information to infer whether a particular individual has delegated to them in FRD.

If we restrict who can be a voluntary representative, and forbid voluntary representatives from delegating, then this feature can be used in FRD. For example, we could allow all candidates for election, including those who are not elected as representatives, to receive delegations as long as they are committed to publicly casting their votes. However, only elected representatives would receive weight by the default mechanism, while the unelected representatives would receive weight solely from delegation. In this system, a party or minority group too small to get a representative elected can still ensure a form of proportional representation for themselves by actively delegating to unelected representatives. Whether this is desirable depends on whether it is desirable for the election rule to filter out extreme views so that extremist unelected candidates with small followings cannot receive any amount of voting power.

Concentration of power One concern with transitive delegations in Liquid Democracy is the concentration of power (Halpern et al. 2023; Zhang and Grossi 2021; Colley et al. 2023). By contrast, the default scheme in FRD is a ballast against a severe concentration of power, as long as there are sufficiently many representatives. The only way that power can concentrate in FRD is if active non-transitive delegations make it that way, in which case it may be justifiable. When delegations are transitive, a voter can change their delegation and still end up with their vote ultimately being cast by the same proxy. In FRD voters know exactly who will be casting a vote on their behalf when they delegate, without requiring any knowledge about the underlying (bipartite) delegation graph.

Districts and gerrymandering In some representative democracies, representatives represent a particular region or district. These regions can have different size populations, and the number of representatives from each region can depend on its population

size. However, since the number of representatives is necessarily a small integer, there may be no way to perfectly satisfy notions of proportionality without weighting the representatives. Moreover, political actors may try to change the defined boundaries of the regions to manipulate the election in an act of gerrymandering (Lewenberg et al. 2017). In FRD, the use of delegation to re-weight the representatives takes the sting out of gerrymandering when disenfranchised voters in a gerrymandered region can re-assign their voting weight to other representatives from their region or from other regions. The question remains what the default distribution should be across representatives from different regions given that the regions have different population sizes, and how it should be updated based on delegation. To this end we proposed Quadratic FRD.

Quadratic FRD There are two main uses of quadratics in the literature on voting. The first is the Penrose square root law regarding voting power (Penrose 1946), and the second is from proposals for Quadratic Voting (Lalley and Weyl 2018). We propose to combine aspects of each to illustrate the flexibility of FRD, but leave any analysis to future work.

With a set of districts of different population sizes, each district having a single representative, to ensure that all voters have the same power (according to the Penrose-Banzhaf power index), one must assign the representatives power in proportion to the square root of the size of their district before using a weighted majority vote (Penrose 1946). This is referred to as the Penrose method, and we use the Penrose method to construct a default weighting scheme. In Quadratic Voting, voters “pay” for votes where the number of votes they receive is the square root of how much they pay.

Assume voters are partitioned into districts and each district elects a single representative, and the representatives will vote using a weighted majority vote. In a simple extension of the Penrose method to allow for delegation, the weight of each representative in the weighted majority rule is simply the square root of the total weight assigned to them by the voters. While this is not exactly the same as setting each representative’s power equal to the square root of their weight, empirical work has shown it to be a close approximation (Zick et al. 2011). Formally, if we let w_{jl}^{def} be the weight delegated by voter v_j to representative d_l , and w_{jl}^{def} be the weight v_j assigns to d_l by default, then the weight of d_l becomes $w_l = \sqrt{\sum_{v_j \in V} (w_{jl}^{def}) + (w_{jl}^{def})}$.

An alternative is to blend the Penrose method with Quadratic Voting. Now, the weight of a representative will be the square root of the sum of weights assigned to them by default plus the sum of square roots of weights assigned to them by delegation: $w_l = \sqrt{\sum_{v_j \in V} (w_{jl}^{def})} + \sqrt{\sum_{v_j \in V} (w_{jl}^{del})}$ where $w_{jl}^{def} + w_{jl}^{del}$ is a constant for all voters. Thus, if no voter delegates we have precisely the Penrose method and if all voters delegate we have precisely Quadratic Voting, where voters expend their unit of voting weight to “buy” voting weight for representatives. Note that a voters’ delegation might be the same assignment of weight to representatives as the default, but the act of delegating with the effort it requires changes how this weight is aggregated.

Future work The binary issues model and the analysis we provide are only the tip of the iceberg for studying FRD. The most direct way to build on our results is to consider more sophisticated, and more realistic, models of preferences for binary issues and for election preferences. Assuming uniformly random preferences allowed us to isolate the effects of the election process and weighting schemes, but these types of preferences are rarely, if ever, observed empirically (Mattei and Walsh 2013, 2017). More sophisticated preference models, including spatial preferences (Enelow and Hinich 1984; Anshelevich et al. 2018) and conditional preference networks (Boutilier et al. 1999; Cornelio et al. 2013, 2014), should be considered.

Assuming that voter preferences over candidates were derived precisely from their agreement was meant to give Representative Democracy the best chance of achieving high majority agreement, but this too should be relaxed. Along the same lines, different participation rates, voter delegation behaviors, decision rules, and election rules remain to be examined.

Another line of inquiry would be to evaluate FRD mechanisms in terms of their ability to uncover ground-truths, similar to some of the literature on Liquid Democracy and Condorcet Jury Theorems (Zhang and Grossi 2022; List and Goodin 2001; Estlund 1994). Similarly, one can compare voting systems based on how they impact measures of social welfare. Our analysis also did not examine strategic behavior by the voters, candidates, or representatives, but this possibility raises many interesting questions. Problems related to pandering, or misrepresentation by candidates, in FRD have already begun to be examined (Sun et al. 2023).

Acknowledgements Portions of this work were completed while both Nicholas Mattei and Ben Abramowitz were at IBM Research, Yorktown Heights, New York. Ben Abramowitz was supported by the CIFellows Project. Nicholas Mattei was supported in part by NSF Awards IIS-RI-2339880, IIS-RI-2007955, IIS-III-2107505, and IIS-RI-2134857, as well as an IBM Faculty Award and a Google Research Scholar Award. This material is based upon work supported by the National Science Foundation under Grant #2030859 to the Computing Research Association for the CIFellows Project.

Data availability All results are reproducible by Monte Carlo simulation. No other data sets were used for this work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

11 Appendix

11.1 Full coverage

In Sect. 4.1, we prove that Max k -Coverage is NP-Hard. However, for all voters to be able to delegate incisively, more than coverage is required. It is necessary that the representatives are not unanimous. In other words, there must be at least one

representative who agrees with the voter majority and at least one representative who agrees with the voter minority on every issue. If we say that any issue on which the representatives are not unanimous is *fully covered* then we want to know if Max k -Full Coverage is NP-Hard as well. We provide an affirmative proof below.

Problem 3 (Max k -Full Coverage) Let $S = \{s^1, \dots, s^r\}$ be a set of binary issues and $C = \{c_1, \dots, c_m\}$ a set of candidates where candidate c_l has preference $c_l^i \in \{0, 1\}$ on issue s^i . The problem of Max k -Full Coverage is the problem of computing a subset of $k \leq m$ representatives $D \subseteq C$ that maximizes the number of fully covered issues, where issue $s^i \in S$ is fully covered if $0 < \sum_{d_l \in D} d_l^i < |D|$.

Theorem 5 *Max k -full coverage is NP-hard.*

Proof We now prove the complexity of Max k -Full Coverage by polynomial-time reduction from Max k -Coverage. To do this we construct an instance of Max k -Full Coverage by adding an additional candidate \hat{c} , adding $r + 1$ additional issues to the original r issues, and desire a set of $k + 1$ candidates. We show that in this new instance of Max k -Full Coverage the additional candidate must be selected in any optimal solution because they are uniquely required to cover the $r + 1$ added issues, and the remaining k candidates in the solution set correspond exactly to the optimal k candidates in the solution to our original Max k -Coverage instance.

Given an instance $(S = \{s^1, \dots, s^r\}, C = \{c_1, \dots, c_m\}, k)$ of Max k -Coverage we construct an instance of Max k -Full Coverage as follows. Create a set of binary issues \tilde{S} equal to S augmented with $r + 1$ additional binary issues so that $\tilde{S} = \{s^1, \dots, s^r, s^{r+1}, \dots, s^{2r+1}\}$. Create a set of candidates $\tilde{C} = C \cup \{\hat{c}\}$ where $\tilde{c}_l^i = c_l^i$ for all $c_l \in C$, for all $s^i \in \{s^1, \dots, s^r\}$, and $\hat{c}^i = 0$ for all issues $s^i \in \{s^{r+1}, \dots, s^{2r+1}\}$. Let $\tilde{c}_l^i = 0$ for all $\tilde{c}_l \in \tilde{C} \setminus \{\hat{c}\}$ for issues $s^i \in \{s^{r+1}, \dots, s^{2r+1}\}$ and let $\hat{c}^i = 1$ for all issues $s^i \in \{s^{r+1}, \dots, s^{2r+1}\}$. Our objective is to select a set $\tilde{D} \subseteq \tilde{C}$ of $k + 1$ candidates from \tilde{C} that maximizes full coverage. We will now prove that for all solutions \tilde{D} to our new Max k -Full Coverage problem, $\tilde{D} = \{\hat{c}\} \cup D$ where D is a set of k candidates whose corresponding counterparts maximize coverage over issues $\{s^1, \dots, s^r\}$ in our original Max k -Coverage instance.

Lemma 2 \tilde{D} must contain \hat{c} .

Proof Clearly, the set $\{\hat{c}\} \cup \{\tilde{c}_l\}$ achieves full coverage for issues $\{s^{r+1}, \dots, s^{2r+1}\}$ for any $\tilde{c}_l \in \tilde{C} \setminus \{\hat{c}\}$, and any set which does not contain \hat{c} cannot fully cover (or cover) $\{s^{r+1}, \dots, s^{2r+1}\}$. Since $\{s^{r+1}, \dots, s^{2r+1}\}$ comprises more than half the issues, any set of $k + 1$ candidates for $k \geq 1$ that maximizes the number of issues fully covered must contain \hat{c} . This concludes our proof of Lemma 2. \square

Given that $\hat{c}^i = 0$ for all issues $\{s^1, \dots, s^r\}$, the set of candidates $D = \tilde{D} \setminus \{\hat{c}\}$, which maximizes full coverage for issues $\{s^1, \dots, s^r\}$ is the set of k candidates which maximizes coverage over issues in $\{s^1, \dots, s^r\}$. Therefore, the k candidates corresponding to D are the solution to our original instance of Max k -Coverage and given the solution D to Max k -Coverage we simply add \hat{c} to find \tilde{D} . This concludes our proof of Theorem 5. \square

References

- Abramowitz B, Mattei N (2019) Flexible representative democracy: an introduction with binary issues. In: Proceedings of the 28th international joint conference on artificial intelligence, pp 3–10
- Abramowitz B, Mattei N (2022) Social mechanism design: Making maximally acceptable decisions. arXiv preprint [arXiv:2211.08501](https://arxiv.org/abs/2211.08501)
- Abramowitz B, Anshelevich E, Zhu W (2019) Awareness of voter passion greatly improves the distortion of metric social choice. In: Web and internet economics: 15th international conference, WINE 2019, New York, NY, USA, December 10–12, 2019, proceedings 15, Springer, pp 3–16
- Abramson SF, Koçak K, Magazinnik A (2022) What do we learn about voter preferences from conjoint experiments? *Am J Polit Sci* 66(4):1008–1020
- Adams J (1979) II. To John Penn, 27 March 1776. In: Taylor RJ (ed) *The Adams papers, papers of John Adams*, vol. 4, February–August 1776, Harvard University Press, pp 78–86–482
- Alger D (2006) Voting by proxy. *Public Choice* 126(1/2):1–26. <http://www.jstor.org/stable/30026573>
- Anscombe GE (1976) On frustration of the majority by fulfillment of the majority's will. *Analysis* 36(4):161–168
- Anshelevich E, Bhardwaj O, Elkind E, Postl J, Skowron P (2018) Approximating optimal social choice under metric preferences. *Artif Intell* 264:27–51
- Anshelevich E, Fitzsimmons Z, Vaish R, Xia L (2021) Representative proxy voting. Proc AAAI Conf Artif Intell 35:5086–5093
- Auriol E, Gary-Bobo RJ (2012) On the optimal number of representatives. *Public Choice* 153(3–4):419–445
- Austen-Smith D, Banks JS (1996) Information aggregation, rationality, and the Condorcet jury theorem. *Am Polit Sci Rev* 90(1):34–45
- Aziz H, Gaspers S, Gudmundsson J, Mackenzie S, Mattei N, Walsh T (2014) Computational aspects of multi-winner approval voting. arXiv preprint [arXiv:1407.3247](https://arxiv.org/abs/1407.3247)
- Aziz H, Brill M, Conitzer V, Elkind E, Freeman R, Walsh T (2017) Justified representation in approval-based committee voting. *Soc Choice Welf* 48(2):461–485
- Baharad E, Goldberger J, Koppel M, Nitzan S (2012) Beyond Condorcet: Optimal aggregation rules using voting records. *Theor Decis* 72(1):113–130
- Banzhaf JF III (1964) Weighted voting doesn't work: a mathematical analysis. *Rutgers L Rev* 19:317
- Becker R, D'angelo G, Delfaraz E, Gilbert H (2021) Unveiling the truth in liquid democracy with misinformed voters. In: *Algorithmic decision theory: 7th international conference, ADT 2021, Toulouse, France, November 3–5, 2021, proceedings 7*, Springer, pp 132–146
- Behrens J, Kistner A, Nitsche A, Swierczek B (2014) The principles of liquidfeedback. *Interaktive Demokratie*
- Ben-Yashar RC, Nitzan SI (1997) The optimal decision rule for fixed-size committees in dichotomous choice situations: the general result. *Int Econ Rev* 175–186
- Bezembinder T, Van Acker P (1985) The Ostrogorski paradox and its relation to nontransitive choice. *J Math Sociol* 11(2):131–158
- Black D (2012) *The theory of committees and elections*. Springer Science & Business Media
- Bloembergen D, Grossi D, Lackner M (2019) On rational delegations in liquid democracy. Proc AAAI Conf Artif Intell 33:1796–1803
- Blum C, Zuber CI (2016) Liquid democracy: potentials, problems, and perspectives. *J Polit Philos* 24(2):162–182
- Boutilier C, Brafman RI, Hoos HH, Poole D (1999) Reasoning with conditional *ceteris paribus* preference statements. *UAI* 99:71–80
- Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) (2016) *Handbook of computational social choice*. Cambridge University Press
- Brill M (2018) Interactive democracy. In: *Proceedings of the 17th international conference on autonomous agents and multiagent systems*, pp 1183–1187
- Brill M, Talmon N (2018) Pairwise liquid democracy. In: *Proceedings of the twenty-seventh international joint conference on artificial intelligence (IJCAI-18)*, vol 18, pp 137–143
- Brill M, Delemazure T, George AM, Lackner M, Schmidt-Kraepelin U (2022) Liquid democracy with ranked delegations. Proc AAAI Conf Artif Intell 36:4884–4891
- Christoff Z, Grossi D (2017) Binary voting with delegable proxy: an analysis of liquid democracy. arXiv preprint [arXiv:1707.08741](https://arxiv.org/abs/1707.08741)

- Cohensius G, Manner S, Meir R, Meirom E, Orda A (2017) Proxy voting for better outcomes. In: Proceedings of the 16th conference on autonomous agents and multiagent systems, pp 858–866
- Colley R (2021) Multi-agent ranked delegations in voting. In: Proceedings of the 20th international conference on autonomous agents and multiagent systems, pp 1802–1804
- Colley R, Grandi U, Novaro A (2021) Smart voting. In: Twenty-ninth international joint conference on artificial intelligence (IJCAI 2020), pp 1734–1740
- Colley R, Grandi U, Novaro A (2022) Unravelling multi-agent ranked delegations. *Auton Agent Multi-Agent Syst* 36(1):1–35
- Colley R, Delemazure T, Gilbert H (2023) Measuring a priori voting power in liquid democracy. In: Proceedings of the thirty-second international joint conference on artificial intelligence, pp 2607–2615
- Cook WD, Kress M (1985) Ordinal ranking with intensity of preference. *Manage Sci* 31(1):26–32
- Cornelio C, Goldsmith J, Mattei N, Rossi F, Venable KB (2013) Updates and uncertainty in cp-nets. In: AI 2013: advances in artificial intelligence: 26th Australasian Joint Conference, Dunedin, New Zealand, December 1–6, 2013. Proceedings 26, Springer, pp 301–312
- Cornelio C, Grandi U, Goldsmith J, Mattei N, Rossi F, Venable KB (2014) Voting with cp-nets using a probabilistic preference structure. In: Proc. of the 5th international workshop on computational social choice (COMSOC)
- Domshlak C, Hüllermeier E, Kaci S, Prade H (2011) Preferences in AI: An overview. *Artif Intell* 175(7–8):1037–1052
- Dunn J (1995) Democracy: the unfinished journey. Oxford University Press
- Ebadian S, Kehne G, Micha E, Procaccia AD, Shah N (2022) Is sortition both representative and fair? *Adv Neural Inf Process Syst* 35:3431–3443
- Eğecioğlu Ö, Giritligil AE (2013) The impartial, anonymous, and neutral culture model: a probability model for sampling public preference structures. *J Math Sociol* 37(4):203–222
- Elkind E (2023) Proportionality in multiwinner voting: the power of local search. In: Proceedings of the 2023 international conference on autonomous agents and multiagent systems, pp 4
- Enelow JM, Hinich MJ (1984) The spatial theory of voting: an introduction. CUP Archive
- Escoffier B, Gilbert H, Pass-Lanneau A (2019) The convergence of iterative delegations in liquid democracy in a social network. In: International Symposium on algorithmic game theory. Springer, pp 284–297
- Estlund DM (1994) Opinion leaders, independence, and condorcet's jury theorem. *Theor Decis* 36:131–162
- Faliszewski P, Skowron P, Slinko A, Talmon N (2017) Multiwinner voting: A new challenge for social choice theory. In: Endriss U (ed) Trends in computational social choice, AI Access Foundation, chap 2
- Feige U (1998) A threshold of $\ln n$ for approximating set cover. *J ACM* 45(4):634–652
- Feld SL, Grofman B (1984) The accuracy of group majority decisions in groups with added members. *Public Choice* 42(3):273–285
- Ford B (2002) Delegative democracy, unpublished Manuscript. <http://brynosaurus.com/deleg/deleg.pdf>
- Freeman R, Brill M, Conitzer V (2014) On the axiomatic characterization of runoff voting rules. In: Proceedings of the AAAI conference on artificial intelligence, vol 28
- Fürnkranz J, Hüllermeier E (2003) Pairwise preference learning and ranking. In: European conference on machine learning, Springer, pp 145–156
- Gerasimou G (2021) Simple preference intensity comparisons. *J Econ Theory* 192:105199
- Gersbach H, Mamageishvili A, Schneider M (2022) Risky vote delegation. CEPR Discussion Paper No DP17044
- Gözl P, Kahng A, Mackenzie S, Procaccia AD (2021) The fluid mechanics of liquid democracy. *ACM Trans Econ Comput* 9(4):1–39
- Green-Armytage J (2015) Direct voting and proxy voting. *Const Polit Econ* 26(2):190–220
- Grofman B, Feld SL (1983) Determining optimal weights for expert judgment. In: Information pooling and group decision making, pp 167–72
- Grofman B, Owen G, Feld SL (1983) Thirteen theorems in search of the truth. *Theor Decis* 15(3):261–278
- Halpern D, Halpern JY, Jadbabaie A, Mossel E, Procaccia A, Revel M (2023) In defense of liquid democracy. In: Twenty-fourth ACM conference on electronic commerce
- Halpern J, Kleiman-Weiner M (2018) Towards formal definitions of blameworthiness, intention, and moral responsibility. In: Proceedings of the AAAI conference on artificial intelligence, vol 32
- Harding J (2019) Incorporating preference information into formal models of transitive proxy voting. MSc Thesis, University of Amsterdam
- Harding J (2022) Proxy selection in transitive proxy voting. *Soc Choice Welf* 58(1):69–99

- Hardt S, Lopes LC (2015) Google votes: a liquid democracy experiment on a corporate social network. Technical Disclosure Commons
- Kahng A, Mackenzie S, Procaccia A (2021) Liquid democracy: an algorithmic perspective. *J Artif Intell Res* 70:1223–1252
- Karotkin D, Paroush J (2003) Optimum committee size: Quality-versus-quantity dilemma. *Soc Choice Welf* 20(3):429–441
- Kelly JS (1989) 6. The Ostrogorski paradox. *Soc Choice Welf* 71–76
- KhudaBukhsh AR, Xu L, Hoos HH, Leyton-Brown K (2016) SATenstein: automatically building local search SAT solvers from components. *Artif Intell* 232:20–42
- Lackner M, Skowron P (2023) Multi-winner voting with approval preferences. Springer Nature
- Laffond G, Lainé J (2009) Condorcet choice and the Ostrogorski paradox. *Soc Choice Welf* 32:317–333
- Lalley SP, Weyl EG (2018) Quadratic voting: how mechanism design can radicalize democracy. In: AEA papers and proceedings, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol 108, pp 33–37
- Lewenberg Y, Lev O, Rosenschein JS (2017) Divide and conquer: using geographic manipulation to win district-based elections. In: Proceedings of the 16th conference on autonomous agents and multiagent systems, pp 624–632
- List C, Goodin RE (2001) Epistemic democracy: generalizing the condorcet jury theorem. *J Polit Philos* 9(3):277–306
- Magdon-Ismail M, Xia L (2018) A mathematical model for optimal decisions in a representative democracy. *Adv Neural Inf Process Syst* 31
- Markakis E, Papasotiropoulos G (2021) An approval-based model for single-step liquid democracy. In: International symposium on algorithmic game theory, Springer, pp 360–375
- Mattei N, Walsh T (2013) Preflib: a library for preferences <http://www.preflib.org>. In: Algorithmic decision theory: third international conference, ADT 2013, Bruxelles, Belgium, November 12–14, 2013, proceedings 3, Springer, pp 259–270
- Mattei N, Walsh T (2017) A preflib.org retrospective: lessons learned and new directions. *Trends Comput Soc Choice* 289–309
- Meir R, Sandomirskiy F, Tennenholz M (2021) Representative committees of peers. *J Artif Intell Res* 71:401–429
- Miller JC (1969) A program for direct and proxy voting in the legislative process. *Public Choice* 7(1):107–113
- Nitzan S, Paroush J (1982) Optimal decision rules in uncertain dichotomous choice situations. *Int Econ Rev* 289–297
- Nitzan S, Paroush J (2017) Collective decision making and jury theorems. In: Oxford Handbook of Law and Economics, pp 494–516
- Paroush J, Karotkin D (1989) Robustness of optimal majority rules over teams with changing size. *Soc Choice Welf* 6(2):127–138
- Penrose LS (1946) The elementary statistics of majority voting. *J R Stat Soc* 109(1):53–57
- Pitkin HF (2016) The concept of representation. In: Democracy: a reader. Columbia University Press, pp 155–158
- Pivato M, Soh A (2020) Weighted representative democracy. *J Math Econ* 88:52–63
- Rae DW, Daudt H (1976) The Ostrogorski paradox: a peculiarity of compound majority decision. *Eur J Polit Res* 4(4):391–398
- Revel M, Boehmer N, Colley R, Brill M, Faliszewski P, Elkind E (2023) Selecting representative bodies: an axiomatic view. arXiv preprint [arXiv:2304.02774](https://arxiv.org/abs/2304.02774)
- Rodriguez MA, Steinbock DJ (2004) Societal-scale decision making using social networks. In: North American association for computational social and organizational science conference proceedings
- Sako K, Kilian J (1995) Receipt-free mix-type voting scheme. In: International conference on the theory and applications of cryptographic techniques, Springer, pp 393–403
- Shapiro E, Talmon N (2018) Incorporating reality into social choice. In: Proceedings of the 17th international conference on autonomous agents and multiagent systems, pp 1188–1192
- Shubik M (1970) On homo politicus and the instant referendum. *Public Choice* 9:79–84
- Skowron P, Faliszewski P, Lang J (2016) Finding a collective set of items: from proportional multirepresentation to group recommendation. *Artif Intell* 241:191–216
- Skowron PK (2015) What do we elect committees for? A voting committee model for multi-winner rules. In: Twenty-fourth international joint conference on artificial intelligence (IJCAI 2015)

- Smith JH (1973) Aggregation of preferences with variable electorate. *Econ J Econ Soc* 1027–1041
- Soh Voutsas AC (2020) Approval voting & majority judgement in weighted representative democracy. Avail SSRN. <https://doi.org/10.2139/ssrn.3731754>
- Sun X, Masur J, Abramowitz B, Mattei N, Zheng Z (2023) Pandering in a (flexible) representative democracy. In: *Uncertainty in artificial intelligence*, PMLR, pp 2058–2068
- Tsetlin I, Regenwetter M, Grofman B (2003) The impartial culture maximizes the probability of majority cycles. *Soc Choice Welf* 21:387–398
- Tullock G (1967) *Toward a mathematics of politics*. University of Michigan Press
- Tullock G (1992) Computerizing politics. *Math Comput Modell* 16(8–9):59–65
- Young HP, Levenglick A (1978) A consistent extension of Condorcet's election principle. *SIAM J Appl Math* 35(2):285–300
- Yu CH, Werfel J, Nagpal R (2010) Collective decision-making in multi-agent systems by implicit leadership. In: 9th international joint conference on autonomous agents and multiagent systems 2010, AAMAS 2010, international foundation for autonomous agents and multiagent systems (IFAAMAS), pp 1189–1196
- Zhang Y, Grossi D (2021) Power in liquid democracy. *Proc AAAI Conf Artif Intell* 35:5822–5830
- Zhang Y, Grossi D (2022) Tracking truth by weighting proxies in liquid democracy. In: *Proceedings of the 21st international conference on autonomous agents and multiagent systems*, pp 1482–1490
- Zhu H, Zhou M, Alkins R (2012) Group role assignment via a Kuhn–Munkres algorithm-based solution. *IEEE Trans Syst Man Cybern Part A Syst Hum* 42(3):739–750
- Zick Y, Skopalik A, Elkind E (2011) The shapley value as a function of the quota in weighted voting games. *IJCAI* 11:490–495
- Zwicker WS (2015) Introduction to voting theory. In: Brandt F, Conitzer V, Endriss U, Lang J, Procaccia AD (eds) *Handbook of computational social choice*, chap 2. Cambridge University Press

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.